

MPC8315E PowerQUICC II Pro Integrated Host Processor Family Reference Manual

Supports
MPC8315E
MPC8315
MPC8314E
MPC8314

MPC8315ERM
Rev. 2
06/2010



How to Reach Us:

Home Page:

www.freescale.com

Web Support:

<http://www.freescale.com/support>

USA/Europe or Locations Not Listed:

Freescale Semiconductor, Inc.
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
+1-800-521-6274 or
+1-480-768-2130
www.freescale.com/support

Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
www.freescale.com/support

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064
Japan
0120 191014 or
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor China Ltd.
Exchange Building 23F
No. 118 Jianguo Road
Chaoyang District
Beijing 100022
China
+86 010 5879 8000
support.asia@freescale.com

For Literature Requests Only:

Freescale Semiconductor
Literature Distribution Center
+1-800 441-2447 or
+1-303-675-2140
Fax: +1-303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

Freescale, the Freescale logo, CodeWarrior, ColdFire, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc. Reg. U.S. Pat. & Tm. Off. CoreNet, QorIQ, QUICC Engine, and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org. IEEE 802, 802.1, 802.2, 802.3, 754, 1149.1, and 1588 are registered trademarks of the Institute of Electrical and Electronics Engineers, Inc. (IEEE). This product is not endorsed or approved by the IEEE.

© 2010 Freescale Semiconductor, Inc.



Contents

| Paragraph Number | Title | Page Number |
|------------------|-------|-------------|
|------------------|-------|-------------|

About This Book

Chapter 1 Overview

| | | |
|---------|---|------|
| 1.1 | MPC8315 PowerQUICC II Pro Processor Overview | 1-1 |
| 1.2 | MPC8315E Architecture Overview | 1-10 |
| 1.2.1 | Power Architecture Core | 1-10 |
| 1.2.2 | Security Engine..... | 1-13 |
| 1.2.3 | DDR Memory Controller..... | 1-13 |
| 1.2.4 | Dual Enhanced Three-Speed Ethernet Controllers..... | 1-14 |
| 1.2.5 | SerDes PHY | 1-14 |
| 1.2.6 | PCI Controller..... | 1-15 |
| 1.2.6.1 | PCI Bus Arbitration Unit..... | 1-15 |
| 1.2.7 | Universal Serial Bus (USB) 2.0..... | 1-16 |
| 1.2.7.1 | USB Dual-Role Controller | 1-17 |
| 1.2.8 | Enhanced Local Bus Controller (eLBC)..... | 1-17 |
| 1.2.9 | Integrated Programmable Interrupt Controller (IPIC)..... | 1-19 |
| 1.2.10 | Time Division Multiplexing (TDM) Interface..... | 1-19 |
| 1.2.11 | I ² C Interface..... | 1-20 |
| 1.2.12 | DMA Controller..... | 1-21 |
| 1.2.13 | Dual Universal Asynchronous Receiver/Transmitter (DUART)..... | 1-21 |
| 1.2.14 | System Timers | 1-22 |
| 1.3 | Applications | 1-22 |
| 1.3.1 | Media Server/NAS..... | 1-23 |
| 1.3.2 | Low-End Voice Gateway | 1-24 |
| 1.3.3 | 802.11n WLAN Access Point..... | 1-25 |

Chapter 2 Signal Descriptions

| | | |
|-----|---|------|
| 2.1 | Signals Overview | 2-1 |
| 2.2 | Output Signal States During Reset | 2-31 |

Chapter 3 Memory Map

| | | |
|-----|---|-----|
| 3.1 | Internal Memory-Mapped Registers | 3-1 |
| 3.2 | Accessing IMMR Memory From the Local Processor..... | 3-1 |
| 3.3 | IMMR Address Map..... | 3-1 |

Contents

| Paragraph Number | Title | Page Number |
|--|---|----------------|
| Chapter 4 | | |
| Reset, Clocking, and Initialization | | |
| 4.1 | External Signals | 4-1 |
| 4.1.1 | Reset Signals | 4-1 |
| 4.1.2 | Clock Signals | 4-3 |
| 4.2 | Functional Description | 4-4 |
| 4.2.1 | Reset Operations | 4-4 |
| 4.2.1.1 | Reset Causes | 4-5 |
| 4.2.1.2 | Reset Actions | 4-5 |
| 4.2.2 | Power-On Reset Flow | 4-6 |
| 4.2.3 | Hard Reset Flow | 4-8 |
| 4.2.4 | Soft Reset Flow | 4-9 |
| 4.3 | Reset Configuration | 4-9 |
| 4.3.1 | Reset Configuration Signals | 4-9 |
| 4.3.1.1 | Reset Configuration Word Source | 4-10 |
| 4.3.1.2 | SYS_CLK_IN Division | 4-11 |
| 4.3.1.3 | Selecting Reset Configuration Input Signals | 4-11 |
| 4.3.2 | Reset Configuration Words | 4-12 |
| 4.3.2.1 | Reset Configuration Word Low Register (RCWLR) | 4-13 |
| 4.3.2.1.1 | System PLL VCO Division | 4-13 |
| 4.3.2.1.2 | System PLL Configuration | 4-14 |
| 4.3.2.2 | Reset Configuration Word High Register (RCWHR) | 4-15 |
| 4.3.2.2.1 | PCI Host/Agent Configuration | 4-16 |
| 4.3.2.2.2 | Boot Memory Space (BMS) | 4-17 |
| 4.3.2.2.3 | Boot Sequencer Configuration | 4-17 |
| 4.3.2.2.4 | Boot ROM Location | 4-18 |
| 4.3.2.2.5 | eTSEC1 Mode | 4-20 |
| 4.3.2.2.6 | eTSEC2 Mode | 4-21 |
| 4.3.2.2.7 | e300 Core True Little-Endian | 4-21 |
| 4.3.2.2.8 | LALE Configuration | 4-22 |
| 4.3.3 | Loading the Reset Configuration Words | 4-22 |
| 4.3.3.1 | Loading from Local Bus | 4-22 |
| 4.3.3.1.1 | Local Bus Controller Setting | 4-23 |
| 4.3.3.2 | Loading from I2C EEPROM | 4-23 |
| 4.3.3.2.1 | Using the Boot Sequencer Reset Configuration | 4-24 |
| 4.3.3.2.2 | EEPROM Calling Address | 4-24 |
| 4.3.3.2.3 | EEPROM Data Format in Reset Configuration Mode | 4-24 |
| 4.3.3.2.4 | Reset Configuration Load Fail | 4-27 |
| 4.3.3.3 | Default Reset Configuration Words | 4-27 |
| 4.3.3.3.1 | Examples for Hard-Coded Reset Configuration Words Usage | 4-28 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|---|-------------|
| 4.4 | Clocking | 4-29 |
| 4.4.1 | Clocking in PCI Host Mode..... | 4-30 |
| 4.4.1.1 | PCI Clock Outputs (PCI_CLK_OUT[0:2]) | 4-31 |
| 4.4.2 | Clocking In PCI Agent Mode | 4-31 |
| 4.4.3 | System Clock Domains..... | 4-31 |
| 4.4.4 | USB Clocking..... | 4-32 |
| 4.4.5 | Ethernet Clocking | 4-33 |
| 4.4.6 | Real-Time Clock (RTC)..... | 4-33 |
| 4.5 | Memory Map/Register Definitions | 4-33 |
| 4.5.1 | Reset Configuration Register Descriptions..... | 4-33 |
| 4.5.1.1 | Reset Configuration Word Low Register (RCWLR)..... | 4-34 |
| 4.5.1.2 | Reset Configuration Word High Register (RCWHR)..... | 4-34 |
| 4.5.1.3 | Reset Status Register (RSR) | 4-34 |
| 4.5.1.4 | Reset Mode Register (RMR) | 4-36 |
| 4.5.1.5 | Reset Protection Register (RPR) | 4-36 |
| 4.5.1.6 | Reset Control Register (RCR) | 4-37 |
| 4.5.1.7 | Reset Control Enable Register (RCER)..... | 4-38 |
| 4.5.2 | Clock Configuration Registers..... | 4-38 |
| 4.5.2.1 | System PLL Mode Register (SPMR) | 4-38 |
| 4.5.2.2 | Output Clock Control Register (OCCR)..... | 4-40 |
| 4.5.2.3 | System Clock Control Register (SCCR)..... | 4-41 |

Chapter 5 System Configuration

| | | |
|-----------|--|-----|
| 5.1 | Introduction..... | 5-1 |
| 5.2 | Local Memory Map Overview and Example | 5-1 |
| 5.2.1 | Address Translation and Mapping | 5-3 |
| 5.2.2 | Window into Configuration Space..... | 5-4 |
| 5.2.3 | Local Access Windows..... | 5-4 |
| 5.2.3.1 | Local Access Register Memory Map | 5-5 |
| 5.2.4 | Local Access Register Descriptions | 5-6 |
| 5.2.4.1 | Internal Memory Map Registers Base Address Register (IMMRBAR)..... | 5-6 |
| 5.2.4.1.1 | Updating IMMRBAR..... | 5-6 |
| 5.2.4.2 | Alternate Configuration Base Address Register (ALTCBAR)..... | 5-7 |
| 5.2.4.3 | LBC Local Access Window <i>n</i> Base Address Registers (LBLAWBAR0–LBLAWBAR3) | 5-8 |
| 5.2.4.3.1 | LBLAWBAR0[BASE_ADDR] Reset Value | 5-8 |
| 5.2.4.4 | LBC Local Access Window <i>n</i> Attributes Registers (LBLAWAR0–LBLAWAR3) | 5-9 |
| 5.2.4.4.1 | LBLAWAR0[EN] and LBLAWAR0[SIZE] Reset Value | 5-9 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 5.2.4.5 | PCI Local Access Window <i>n</i> Base Address Register (PCILAWBAR0–PCILAWBAR1) | 5-10 |
| 5.2.4.5.1 | PCILAWBAR0[BASE_ADDR] Reset Value | 5-10 |
| 5.2.4.6 | PCI Local Access Window <i>n</i> Attributes Registers (PCILAWAR0–PCILAWAR1) | 5-11 |
| 5.2.4.6.1 | PCILAWAR0[EN] and PCILAWAR0[SIZE] Reset Value | 5-11 |
| 5.2.4.7 | PCI Express 1 Local Access Window Base Address Register (PCIEXP1LAWBAR) | 5-12 |
| 5.2.4.8 | PCI Express 1 Local Access Window Attributes Registers (PCIEXP1LAWAR) . | 5-12 |
| 5.2.4.9 | PCI Express 2 Local Access Window Base Address Register (PCIEXP2LAWBAR).. 5-13 | 5-13 |
| 5.2.4.10 | PCI Express 2 Local Access Window Attributes Registers (PCIEXP2LAWAR) . | 5-14 |
| 5.2.4.11 | DDR Local Access Window <i>n</i> Base Address Registers (DDRLAWBAR0–DDRLAWBAR1) | 5-14 |
| 5.2.4.11.1 | DDRLAWBAR0[BASE_ADDR] Reset Value | 5-15 |
| 5.2.4.12 | DDR Local Access Window <i>n</i> Attributes Registers (DDRLAWAR0–DDRLAWAR1) 5-15 | 5-15 |
| 5.2.4.12.1 | DDRLAWAR0[EN] and DDRLAWAR0[SIZE] Reset Value | 5-16 |
| 5.2.5 | Precedence of Local Access Windows | 5-16 |
| 5.2.6 | Configuring Local Access Windows | 5-17 |
| 5.2.7 | Distinguishing Local Access Windows from Other Mapping Functions | 5-17 |
| 5.2.8 | Outbound Address Translation and Mapping Windows | 5-17 |
| 5.2.9 | Inbound Address Translation and Mapping Windows | 5-17 |
| 5.2.9.1 | PCI Inbound Windows | 5-18 |
| 5.2.10 | Internal Memory Map | 5-18 |
| 5.2.11 | Accessing Internal Memory from External Masters | 5-18 |
| 5.3 | System Configuration | 5-18 |
| 5.3.1 | System Configuration Register Memory Map | 5-19 |
| 5.3.2 | System Configuration Registers | 5-19 |
| 5.3.2.1 | System General Purpose Register Low (SGPRL) | 5-19 |
| 5.3.2.2 | System General Purpose Register High (SGPRH) | 5-20 |
| 5.3.2.3 | System Part and Revision ID Register (SPRIDR) | 5-20 |
| 5.3.2.3.1 | SPRIDR[PARTID] Coding | 5-21 |
| 5.3.2.4 | System Priority and Configuration Register (SPCR) | 5-21 |
| 5.3.2.5 | System I/O Configuration Register Low (SICRL) | 5-23 |
| 5.3.2.6 | System I/O Configuration Register High (SICRH) | 5-26 |
| 5.3.2.7 | Debug Configuration | 5-29 |
| 5.3.2.7.1 | DDR Debug Configuration | 5-29 |
| 5.3.2.7.2 | Local Bus Debug Configuration | 5-29 |
| 5.3.2.8 | DDR Control Driver Register (DDRCDR) | 5-29 |
| 5.3.2.9 | DDR Debug Status Register (DDRDSR) | 5-31 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 5.3.2.10 | PCI Express Control Registers (PECR1 and PECR2)..... | 5-31 |
| 5.4 | Software Watchdog Timer (WDT)..... | 5-33 |
| 5.4.1 | WDT Overview..... | 5-33 |
| 5.4.2 | WDT Features..... | 5-34 |
| 5.4.3 | WDT Modes of Operation..... | 5-34 |
| 5.4.4 | WDT Memory Map/Register Definition..... | 5-34 |
| 5.4.4.1 | System Watchdog Control Register (SWCRR)..... | 5-35 |
| 5.4.4.2 | System Watchdog Count Register (SWCNR)..... | 5-36 |
| 5.4.4.3 | System Watchdog Service Register (SWSRR)..... | 5-36 |
| 5.4.5 | Functional Description..... | 5-37 |
| 5.4.5.1 | Software Watchdog Timer Unit..... | 5-37 |
| 5.4.5.2 | Modes of Operation..... | 5-39 |
| 5.4.6 | Initialization/Application Information..... | 5-39 |
| 5.4.6.1 | WDT Programming Guidelines..... | 5-39 |
| 5.5 | Real Time Clock Module (RTC)..... | 5-40 |
| 5.5.1 | RTC Overview..... | 5-40 |
| 5.5.2 | RTC Features..... | 5-40 |
| 5.5.3 | RTC Modes of Operation..... | 5-41 |
| 5.5.4 | RTC External Signal Description..... | 5-41 |
| 5.5.5 | RTC Memory Map/Register Definition..... | 5-41 |
| 5.5.5.1 | Real Time Counter Control Register (RTCNR)..... | 5-42 |
| 5.5.5.2 | Real Time Counter Load Register (RTLDR)..... | 5-43 |
| 5.5.5.3 | Real Time Counter Prescale Register (RTPSR)..... | 5-43 |
| 5.5.5.4 | Real Time Counter Register (RTCTR)..... | 5-43 |
| 5.5.5.5 | Real Time Counter Event Register (RTEVR)..... | 5-44 |
| 5.5.5.6 | Real Time Counter Alarm Register (RTALR)..... | 5-45 |
| 5.5.6 | Functional Description..... | 5-45 |
| 5.5.6.1 | Real Time Counter Unit..... | 5-45 |
| 5.5.6.2 | RTC Operational Modes..... | 5-46 |
| 5.5.7 | RTC Programming Guidelines..... | 5-47 |
| 5.6 | Periodic Interval Timer (PIT)..... | 5-47 |
| 5.6.1 | PIT Overview..... | 5-47 |
| 5.6.2 | PIT Features..... | 5-48 |
| 5.6.3 | PIT Modes of Operation..... | 5-48 |
| 5.6.4 | PIT External Signal Description..... | 5-48 |
| 5.6.5 | PIT Memory Map/Register Definition..... | 5-48 |
| 5.6.5.1 | Periodic Interval Timer Control Register (PTCNR)..... | 5-49 |
| 5.6.5.2 | Periodic Interval Timer Load Register (PTLDR)..... | 5-50 |
| 5.6.5.3 | Periodic Interval Timer Prescale Register (PTPSR)..... | 5-50 |
| 5.6.5.4 | Periodic Interval Timer Counter Register (PTCTR)..... | 5-51 |
| 5.6.5.5 | Periodic Interval Timer Event Register (PTEVR)..... | 5-51 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 5.6.6 | Functional Description..... | 5-51 |
| 5.6.6.1 | Periodic Interval Timer Unit..... | 5-51 |
| 5.6.6.2 | PIT Operational Modes..... | 5-52 |
| 5.6.7 | PIT Programming Guidelines..... | 5-53 |
| 5.7 | General-Purpose Timers (GTM)..... | 5-53 |
| 5.7.1 | GTM Overview..... | 5-53 |
| 5.7.2 | GTM Features..... | 5-54 |
| 5.7.3 | GTM Modes of Operation..... | 5-55 |
| 5.7.3.1 | Cascaded Modes..... | 5-55 |
| 5.7.3.2 | Clock Source Modes..... | 5-55 |
| 5.7.3.3 | Reference Modes..... | 5-55 |
| 5.7.3.4 | Capture Modes..... | 5-56 |
| 5.7.4 | GTM External Signal Description..... | 5-56 |
| 5.7.5 | GTM Memory Map/Register Definition..... | 5-57 |
| 5.7.5.1 | Global Timers Configuration Registers (GTCFR _n)..... | 5-59 |
| 5.7.5.2 | Global Timers Mode Registers (GTMDR1–GTMDR4)..... | 5-62 |
| 5.7.5.3 | Global Timers Reference Registers (GTRFR1–GTRFR4)..... | 5-63 |
| 5.7.5.4 | Global Timers Capture Registers (GTCPR1–GTCPR4)..... | 5-63 |
| 5.7.5.5 | Global Timers Counter Registers (GTCNR1–GTCNR4)..... | 5-64 |
| 5.7.5.6 | Global Timers Event Registers (GTEVR1–GTEVR4)..... | 5-64 |
| 5.7.5.7 | Global Timers Prescale Registers (GTPSR1–GTPSR4)..... | 5-65 |
| 5.7.6 | Functional Description..... | 5-66 |
| 5.7.6.1 | General-Purpose Timer Units..... | 5-66 |
| 5.7.6.2 | Reference Modes..... | 5-66 |
| 5.7.6.3 | Capture Modes..... | 5-66 |
| 5.7.6.4 | Cascaded Modes..... | 5-67 |
| 5.7.7 | Initialization/Application Information..... | 5-69 |
| 5.7.7.1 | Programming Guidelines..... | 5-69 |
| 5.7.7.1.1 | GTM Registers..... | 5-69 |
| 5.8 | Power Management Control (PMC)..... | 5-69 |
| 5.8.1 | External Signal Description..... | 5-70 |
| 5.8.2 | PMC Memory Map/Register Definition..... | 5-71 |
| 5.8.2.1 | Power Management Controller Configuration Register (PMCCR)..... | 5-71 |
| 5.8.2.2 | Power Management Controller Event Register (PM CER)..... | 5-72 |
| 5.8.2.3 | Power Management Controller Mask Register (PM C MR)..... | 5-74 |
| 5.8.2.4 | Power Management Controller Configuration Register 1 (PM C CR1)..... | 5-75 |
| 5.8.2.5 | Power Management Controller Configuration Register 2 (PM C CR2)..... | 5-77 |
| 5.8.3 | Functional Description..... | 5-78 |
| 5.8.3.1 | Dynamic Power Management..... | 5-81 |
| 5.8.3.2 | Shutting Down Unused Blocks..... | 5-81 |
| 5.8.3.3 | Software-Controlled Power-Down States..... | 5-81 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 5.8.3.4 | Software-Controlled Power Supply Switching..... | 5-82 |
| 5.8.3.5 | Support of PCI Power Management Interface Specification..... | 5-82 |
| 5.8.3.5.1 | Entering Low Power States—Core-Only Mode..... | 5-85 |
| 5.8.3.5.2 | Entering Low Power States—Core and System Mode..... | 5-85 |
| 5.8.3.6 | Exiting Core and System Low Power States..... | 5-86 |
| 5.8.3.6.1 | Exiting Low Power States—Core-Only Mode..... | 5-86 |
| 5.8.3.6.2 | Exiting Low Power States—Core and System Mode..... | 5-86 |
| 5.8.3.7 | MPC8315E-Specific PMC Low Power States..... | 5-87 |
| 5.8.3.7.1 | Power State Transitions from an ACPI Perspective..... | 5-87 |
| 5.8.3.7.2 | MPC8315E Low Power Sequencing..... | 5-91 |
| 5.8.3.7.3 | PMC External Power Supply Control..... | 5-97 |
| 5.8.3.7.4 | Low-Power Considerations..... | 5-98 |
| 5.8.4 | Initialization/Application Information..... | 5-100 |
| 5.8.4.1 | Core Disable in Low-Power Mode..... | 5-100 |

Chapter 6 Arbiter and Bus Monitor

| | | |
|---------|---|------|
| 6.1 | Arbiter Overview..... | 6-1 |
| 6.1.1 | Coherent System Bus Overview..... | 6-1 |
| 6.2 | Arbiter Memory Map/Register Definition..... | 6-2 |
| 6.2.1 | Arbiter Configuration Register (ACR)..... | 6-3 |
| 6.2.2 | Arbiter Timers Register (ATR)..... | 6-4 |
| 6.2.3 | Arbiter Event Enable Register (AEER)..... | 6-5 |
| 6.2.4 | Arbiter Event Register (AER)..... | 6-6 |
| 6.2.5 | Arbiter Interrupt Definition Register (AIDR)..... | 6-7 |
| 6.2.6 | Arbiter Mask Register (AMR)..... | 6-8 |
| 6.2.7 | Arbiter Event Attributes Register (AEATR)..... | 6-9 |
| 6.2.8 | Arbiter Event Address Register (AEADR)..... | 6-11 |
| 6.2.9 | Arbiter Event Response Register (AERR)..... | 6-12 |
| 6.3 | Functional Description..... | 6-12 |
| 6.3.1 | Arbitration Policy..... | 6-12 |
| 6.3.1.1 | Address Bus Arbitration with <u>PRIORITY</u> [0:1]..... | 6-13 |
| 6.3.1.2 | Address Bus Arbitration with <u>REPEAT</u> | 6-14 |
| 6.3.1.3 | Address Bus Arbitration After <u>ARTRY</u> | 6-15 |
| 6.3.1.4 | Address Bus Parking..... | 6-15 |
| 6.3.1.5 | Data Bus Arbitration..... | 6-15 |
| 6.3.2 | Bus Error Detection..... | 6-15 |
| 6.3.2.1 | Address Time Out..... | 6-15 |
| 6.3.2.2 | Data Time Out..... | 6-16 |
| 6.3.2.3 | Transfer Error..... | 6-16 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|---|-------------|
| 6.3.2.4 | Address Only Transaction Type..... | 6-16 |
| 6.3.2.5 | Reserved Transaction Type..... | 6-17 |
| 6.3.2.6 | Illegal (eciwx/ecowx) Transaction Type..... | 6-17 |
| 6.4 | Initialization/Applications Information | 6-18 |
| 6.4.1 | Initialization Sequence..... | 6-18 |
| 6.4.2 | Error Handling Sequence..... | 6-18 |

Chapter 7 e300 Processor Core Overview

| | | |
|-----------|--|------|
| 7.1 | e300c3 Overview | 7-1 |
| 7.1.1 | e300c3 Features | 7-3 |
| 7.1.2 | Instruction Unit..... | 7-6 |
| 7.1.2.1 | Instruction Queue and Dispatch Unit | 7-6 |
| 7.1.2.2 | Branch Processing Unit (BPU)..... | 7-7 |
| 7.1.3 | Independent Execution Units..... | 7-7 |
| 7.1.3.1 | Integer Unit (IU)..... | 7-7 |
| 7.1.3.2 | Floating-Point Unit (FPU) | 7-7 |
| 7.1.3.3 | Load/Store Unit (LSU) | 7-8 |
| 7.1.3.4 | System Register Unit (SRU)..... | 7-8 |
| 7.1.4 | Completion Unit | 7-8 |
| 7.1.5 | Memory Subsystem Support..... | 7-8 |
| 7.1.5.1 | Memory Management Units (MMUs)..... | 7-9 |
| 7.1.5.2 | Cache Units..... | 7-10 |
| 7.1.6 | Bus Interface Unit (BIU) | 7-10 |
| 7.1.7 | System Support Functions | 7-11 |
| 7.1.7.1 | Power Management | 7-11 |
| 7.1.7.2 | Time Base/Decrementer | 7-11 |
| 7.1.7.3 | JTAG Test and Debug Interface..... | 7-12 |
| 7.1.7.4 | Clock Multiplier..... | 7-12 |
| 7.1.7.5 | Core Performance Monitor | 7-12 |
| 7.2 | e300 Processor and System Version Numbers..... | 7-13 |
| 7.3 | PowerPC Architecture Implementation | 7-13 |
| 7.4 | Implementation-Specific Information..... | 7-14 |
| 7.4.1 | Register Model..... | 7-14 |
| 7.4.1.1 | UISA Registers | 7-17 |
| 7.4.1.1.1 | General-Purpose Registers (GPRs) | 7-17 |
| 7.4.1.1.2 | Floating-Point Registers (FPRs)..... | 7-17 |
| 7.4.1.1.3 | Condition Register (CR)..... | 7-17 |
| 7.4.1.1.4 | Floating-Point Status and Control Register (FPSCR) | 7-17 |
| 7.4.1.1.5 | User-Level SPRs..... | 7-17 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 7.4.1.2 | VEA Registers | 7-18 |
| 7.4.1.3 | OEA Registers | 7-18 |
| 7.4.1.3.1 | Machine State Register (MSR)..... | 7-18 |
| 7.4.1.3.2 | Segment Registers (SRs) | 7-20 |
| 7.4.1.3.3 | Supervisor-Level SPRs | 7-20 |
| 7.4.2 | Instruction Set and Addressing Modes | 7-26 |
| 7.4.2.1 | PowerPC Instruction Set and Addressing Modes | 7-27 |
| 7.4.2.2 | Implementation-Specific Instruction Set | 7-28 |
| 7.4.3 | Cache Implementation | 7-29 |
| 7.4.3.1 | PowerPC Cache Characteristics | 7-29 |
| 7.4.3.2 | Implementation-Specific Cache Organization..... | 7-29 |
| 7.4.3.3 | Instruction and Data Cache Way-Locking..... | 7-31 |
| 7.4.4 | Interrupt Model | 7-31 |
| 7.4.4.1 | PowerPC Interrupt Model..... | 7-31 |
| 7.4.4.2 | Implementation-Specific Interrupt Model | 7-32 |
| 7.4.5 | Memory Management..... | 7-35 |
| 7.4.5.1 | PowerPC Memory Management..... | 7-35 |
| 7.4.5.2 | Implementation-Specific Memory Management..... | 7-35 |
| 7.4.6 | Instruction Timing | 7-36 |
| 7.4.7 | Core Interface | 7-37 |
| 7.4.7.1 | Memory Accesses..... | 7-38 |
| 7.4.7.2 | Signals..... | 7-38 |
| 7.4.8 | Debug Features | 7-39 |
| 7.4.8.1 | Breakpoint Signaling | 7-39 |
| 7.5 | Differences Between Cores..... | 7-40 |

Chapter 8 Integrated Programmable Interrupt Controller (IPIC)

| | | |
|-------|--|------|
| 8.1 | IPIC Introduction | 8-1 |
| 8.2 | IPIC Features | 8-4 |
| 8.3 | IPIC Modes of Operation | 8-4 |
| 8.3.1 | Core Enable Mode | 8-4 |
| 8.3.2 | Core Disable Mode | 8-5 |
| 8.4 | IPIC External Signal Description | 8-5 |
| 8.4.1 | IPIC External Signals Overview..... | 8-5 |
| 8.4.2 | IPIC Detailed Signal Descriptions..... | 8-5 |
| 8.5 | IPIC Memory Map/Register Definition..... | 8-6 |
| 8.5.1 | System Global Interrupt Configuration Register (SICFR) | 8-8 |
| 8.5.2 | System Global Interrupt Vector Register (SIVCR)..... | 8-9 |
| 8.5.3 | System Internal Interrupt Pending Registers (SIPNR_H and SIPNR_L)..... | 8-12 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 8.5.4 | System Internal Interrupt Group A Priority Register (SIPRR_A)..... | 8-15 |
| 8.5.5 | System Internal Interrupt Group B Priority Register (SIPRR_B)..... | 8-16 |
| 8.5.6 | System Internal Interrupt Group C Priority Register (SIPRR_C)..... | 8-16 |
| 8.5.7 | System Internal Interrupt Group D Priority Register (SIPRR_D)..... | 8-17 |
| 8.5.8 | System Internal Interrupt Mask Register (SIMSR_H and SIMSR_L)..... | 8-18 |
| 8.5.9 | System Internal Interrupt Control Register (SICNR)..... | 8-19 |
| 8.5.10 | System External Interrupt Pending Register (SEPNR)..... | 8-21 |
| 8.5.11 | System Mixed Interrupt Group A Priority Register (SMPRR_A)..... | 8-22 |
| 8.5.12 | System Mixed Interrupt Group B Priority Register (SMPRR_B)..... | 8-23 |
| 8.5.13 | System External Interrupt Mask Register (SEMSR)..... | 8-23 |
| 8.5.14 | System External Interrupt Control Register (SECNR)..... | 8-24 |
| 8.5.15 | System Error Status Register (SERSR)..... | 8-26 |
| 8.5.16 | System Error Mask Register (SERMR)..... | 8-27 |
| 8.5.17 | System Error Control Register (SERCR)..... | 8-27 |
| 8.5.18 | System External interrupt Polarity Control Register (SEPCR)..... | 8-28 |
| 8.5.19 | System Internal Interrupt Force Registers (SIFCR_H and SIFCR_L)..... | 8-29 |
| 8.5.20 | System External Interrupt Force Register (SEFCR)..... | 8-30 |
| 8.5.21 | System Error Force Register (SERFR)..... | 8-30 |
| 8.5.22 | System Critical Interrupt Vector Register (SCVCR)..... | 8-31 |
| 8.5.23 | System Management Interrupt Vector Register (SMVCR)..... | 8-31 |
| 8.6 | Functional Description..... | 8-32 |
| 8.6.1 | Interrupt Types..... | 8-32 |
| 8.6.2 | Interrupt Configuration..... | 8-33 |
| 8.6.3 | Internal Interrupts Group Relative Priority..... | 8-34 |
| 8.6.4 | Mixed Interrupts Group Relative Priority..... | 8-34 |
| 8.6.5 | Highest Priority Interrupt..... | 8-35 |
| 8.6.6 | Interrupt Source Priorities..... | 8-35 |
| 8.6.7 | Masking Interrupt Sources..... | 8-39 |
| 8.6.8 | Interrupt Vector Generation and Calculation..... | 8-39 |
| 8.6.9 | Machine Check Interrupts..... | 8-40 |
| 8.7 | Message Shared Interrupts..... | 8-40 |
| 8.7.1 | Memory Map/Register Definition..... | 8-40 |
| 8.7.2 | Message Shared Registers..... | 8-41 |
| 8.7.2.1 | Message Shared Interrupt Register (MSIRs)..... | 8-41 |
| 8.7.2.2 | Message Shared Interrupt Mask Register (MSIMR)..... | 8-41 |
| 8.7.2.3 | Message Shared Interrupt Status Register (MSISR)..... | 8-42 |
| 8.7.2.4 | Message Shared Interrupt Index Register (MSIIR)..... | 8-43 |

Chapter 9 DDR Memory Controller

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 9.1 | Introduction..... | 9-1 |
| 9.2 | Features..... | 9-2 |
| 9.2.1 | Modes of Operation..... | 9-3 |
| 9.3 | External Signal Descriptions..... | 9-3 |
| 9.3.1 | Signals Overview..... | 9-3 |
| 9.3.2 | Detailed Signal Descriptions..... | 9-5 |
| 9.3.2.1 | Memory Interface Signals..... | 9-5 |
| 9.3.2.2 | Clock Interface Signals..... | 9-7 |
| 9.3.2.3 | Debug Signals..... | 9-8 |
| 9.4 | Memory Map/Register Definition..... | 9-8 |
| 9.4.1 | Register Descriptions..... | 9-9 |
| 9.4.1.1 | Chip Select Memory Bounds (CS _n _BNDS)..... | 9-9 |
| 9.4.1.2 | Chip Select Configuration (CS _n _CONFIG)..... | 9-10 |
| 9.4.1.3 | DDR SDRAM Timing Configuration 3 (TIMING_CFG_3)..... | 9-12 |
| 9.4.1.4 | DDR SDRAM Timing Configuration 0 (TIMING_CFG_0)..... | 9-12 |
| 9.4.1.5 | DDR SDRAM Timing Configuration 1 (TIMING_CFG_1)..... | 9-14 |
| 9.4.1.6 | DDR SDRAM Timing Configuration 2 (TIMING_CFG_2)..... | 9-16 |
| 9.4.1.7 | DDR SDRAM Control Configuration (DDR_SDRAM_CFG)..... | 9-18 |
| 9.4.1.8 | DDR SDRAM Control Configuration 2 (DDR_SDRAM_CFG_2)..... | 9-21 |
| 9.4.1.9 | DDR SDRAM Mode Configuration (DDR_SDRAM_MODE)..... | 9-22 |
| 9.4.1.10 | DDR SDRAM Mode 2 Configuration (DDR_SDRAM_MODE_2)..... | 9-23 |
| 9.4.1.11 | DDR SDRAM Mode Control Register (DDR_SDRAM_MD_CNTL)..... | 9-24 |
| 9.4.1.12 | DDR SDRAM Interval Configuration (DDR_SDRAM_INTERVAL)..... | 9-27 |
| 9.4.1.13 | DDR SDRAM Data Initialization (DDR_DATA_INIT)..... | 9-27 |
| 9.4.1.14 | DDR SDRAM Clock Control (DDR_SDRAM_CLK_CNTL)..... | 9-28 |
| 9.4.1.15 | DDR Initialization Address (DDR_INIT_ADDR)..... | 9-28 |
| 9.4.1.16 | DDR IP Block Revision 1 (DDR_IP_REV1)..... | 9-29 |
| 9.4.1.17 | DDR IP Block Revision 2 (DDR_IP_REV2)..... | 9-29 |
| 9.4.1.18 | Memory Error Detect (ERR_DETECT)..... | 9-30 |
| 9.4.1.19 | Memory Error Disable (ERR_DISABLE)..... | 9-30 |
| 9.4.1.20 | Memory Error Interrupt Enable (ERR_INT_EN)..... | 9-31 |
| 9.5 | Functional Description..... | 9-31 |
| 9.5.1 | DDR SDRAM Interface Operation..... | 9-34 |
| 9.5.1.1 | Supported DDR SDRAM Organizations..... | 9-35 |
| 9.5.2 | DDR SDRAM Address Multiplexing..... | 9-36 |
| 9.5.3 | JEDEC Standard DDR SDRAM Interface Commands..... | 9-41 |
| 9.5.4 | DDR SDRAM Interface Timing..... | 9-42 |
| 9.5.4.1 | Clock Distribution..... | 9-46 |
| 9.5.5 | DDR SDRAM Mode-Set Command Timing..... | 9-46 |
| 9.5.6 | DDR SDRAM Registered DIMM Mode..... | 9-47 |
| 9.5.7 | DDR SDRAM Write Timing Adjustments..... | 9-47 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 9.5.8 | DDR SDRAM Refresh | 9-48 |
| 9.5.8.1 | DDR SDRAM Refresh Timing..... | 9-49 |
| 9.5.8.2 | DDR SDRAM Refresh and Power-Saving Modes | 9-49 |
| 9.5.8.2.1 | Self-Refresh in Sleep Mode..... | 9-51 |
| 9.5.9 | DDR Data Beat Ordering..... | 9-52 |
| 9.5.10 | Page Mode and Logical Bank Retention | 9-52 |
| 9.6 | Initialization/Application Information | 9-53 |
| 9.6.1 | Programming Differences between Memory Types | 9-54 |
| 9.6.2 | DDR SDRAM Initialization Sequence | 9-56 |
| 9.6.3 | Using Forced Self-Refresh Mode to Implement a Battery-Backed RAM System | 9-57 |
| 9.6.3.1 | Software Based Self-Refresh..... | 9-57 |
| 9.6.3.2 | Bypassing Re-initialization During Battery-Backed Operation | 9-57 |

Chapter 10 Enhanced Local Bus Controller

| | | |
|------------|--|-------|
| 10.1 | Introduction..... | 10-1 |
| 10.1.1 | Overview..... | 10-2 |
| 10.1.2 | Features..... | 10-2 |
| 10.1.3 | Modes of Operation | 10-3 |
| 10.1.3.1 | eLBC Bus Clock and Clock Ratios | 10-3 |
| 10.1.3.2 | Source ID Debug Mode | 10-4 |
| 10.2 | External Signal Descriptions | 10-4 |
| 10.3 | Memory Map/Register Definition | 10-8 |
| 10.3.1 | Register Descriptions..... | 10-9 |
| 10.3.1.1 | Base Registers (BR0–BR3) | 10-9 |
| 10.3.1.2 | Option Registers (OR0–OR3)..... | 10-11 |
| 10.3.1.2.1 | Address Mask | 10-11 |
| 10.3.1.2.2 | Option Registers (OR _n)—GPCM Mode | 10-13 |
| 10.3.1.2.3 | Option Registers (OR _n)—FCM Mode | 10-15 |
| 10.3.1.2.4 | Option Registers (OR _n)—UPM Mode | 10-18 |
| 10.3.1.3 | UPM Memory Address Register (MAR)..... | 10-19 |
| 10.3.1.4 | UPM Mode Registers (M _x MR) | 10-20 |
| 10.3.1.5 | Memory Refresh Timer Prescaler Register (MRTPR) | 10-22 |
| 10.3.1.6 | UPM/FCM Data Register (MDR) | 10-22 |
| 10.3.1.7 | Special Operation Initiation Register (LSOR)..... | 10-23 |
| 10.3.1.8 | UPM Refresh Timer (LURT)..... | 10-24 |
| 10.3.1.9 | Transfer Error Status Register (LTESR)..... | 10-25 |
| 10.3.1.10 | Transfer Error Check Disable Register (LTEDR)..... | 10-27 |
| 10.3.1.11 | Transfer Error Interrupt Enable Register (LTEIR) | 10-28 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 10.3.1.12 | Transfer Error Attributes Register (LTEATR) | 10-29 |
| 10.3.1.13 | Transfer Error Address Register (LTEAR) | 10-30 |
| 10.3.1.14 | Local Bus Configuration Register (LBCR) | 10-30 |
| 10.3.1.15 | Clock Ratio Register (LCRR) | 10-32 |
| 10.3.1.16 | Flash Mode Register (FMR) | 10-33 |
| 10.3.1.17 | Flash Instruction Register (FIR) | 10-34 |
| 10.3.1.18 | Flash Command Register (FCR) | 10-35 |
| 10.3.1.19 | Flash Block Address Register (FBAR) | 10-36 |
| 10.3.1.20 | Flash Page Address Register (FPAR) | 10-36 |
| 10.3.1.21 | Flash Byte Count Register (FBCR) | 10-38 |
| 10.4 | Functional Description | 10-38 |
| 10.4.1 | Basic Architecture | 10-40 |
| 10.4.1.1 | Address and Address Space Checking | 10-40 |
| 10.4.1.2 | External Address Latch Enable Signal (LALE) | 10-40 |
| 10.4.1.3 | Data Transfer Acknowledge (TA) | 10-41 |
| 10.4.1.4 | Data Buffer Control (LBCTL) | 10-42 |
| 10.4.1.5 | Bus Monitor | 10-42 |
| 10.4.2 | General-Purpose Chip-Select Machine (GPCM) | 10-43 |
| 10.4.2.1 | GPCM Read Signal Timing | 10-44 |
| 10.4.2.2 | GPCM Write Signal Timing | 10-46 |
| 10.4.2.3 | Chip-Select Assertion Timing | 10-47 |
| 10.4.2.3.1 | Programmable Wait State Configuration | 10-48 |
| 10.4.2.3.2 | Chip-Select and Write Enable Negation Timing | 10-48 |
| 10.4.2.3.3 | Relaxed Timing | 10-49 |
| 10.4.2.3.4 | Output Enable (LOE) Timing | 10-52 |
| 10.4.2.3.5 | Extended Hold Time on Read Accesses | 10-52 |
| 10.4.2.4 | External Access Termination (LGTA) | 10-53 |
| 10.4.2.5 | GPCM Boot Chip-Select Operation | 10-54 |
| 10.4.3 | Flash Control Machine (FCM) | 10-55 |
| 10.4.3.1 | FCM Buffer RAM | 10-57 |
| 10.4.3.1.1 | Buffer Layout and Page Mapping for Small-Page NAND Flash Devices | 10-58 |
| 10.4.3.1.2 | Buffer Layout and Page Mapping for Large-Page NAND Flash Devices | 10-59 |
| 10.4.3.1.3 | Error Correcting Codes and the Spare Region | 10-60 |
| 10.4.3.2 | Programming FCM | 10-61 |
| 10.4.3.2.1 | FCM Command Instructions | 10-62 |
| 10.4.3.2.2 | FCM No-Operation Instruction | 10-63 |
| 10.4.3.2.3 | FCM Address Instructions | 10-63 |
| 10.4.3.2.4 | FCM Data Read Instructions | 10-63 |
| 10.4.3.2.5 | FCM Data Write Instructions | 10-64 |
| 10.4.3.3 | FCM Signal Timing | 10-64 |
| 10.4.3.3.1 | FCM Chip-Select Timing | 10-64 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 10.4.3.3.2 | FCM Command, Address, and Write Data Timing | 10-65 |
| 10.4.3.3.3 | FCM Ready/Busy Timing..... | 10-66 |
| 10.4.3.3.4 | FCM Read Data Timing | 10-67 |
| 10.4.3.3.5 | FCM Extended Read Hold Timing..... | 10-68 |
| 10.4.3.4 | FCM Boot Chip-Select Operation | 10-68 |
| 10.4.3.4.1 | FCM Bank 0 Reset Initialization | 10-69 |
| 10.4.3.4.2 | Boot Block Loading into the FCM Buffer RAM..... | 10-69 |
| 10.4.4 | User-Programmable Machines (UPMs)..... | 10-71 |
| 10.4.4.1 | UPM Requests | 10-72 |
| 10.4.4.1.1 | Memory Access Requests..... | 10-73 |
| 10.4.4.1.2 | UPM Refresh Timer Requests | 10-73 |
| 10.4.4.1.3 | Software Requests—RUN Command | 10-74 |
| 10.4.4.1.4 | Exception Requests..... | 10-74 |
| 10.4.4.2 | Programming the UPMs | 10-74 |
| 10.4.4.2.1 | UPM Programming Example (Two Sequential Writes to the RAM Array).... | 10-75 |
| 10.4.4.2.2 | UPM Programming Example (Two Sequential Reads from the RAM Array) | 10-76 |
| 10.4.4.3 | UPM Signal Timing..... | 10-76 |
| 10.4.4.4 | RAM Array | 10-77 |
| 10.4.4.4.1 | RAM Words..... | 10-77 |
| 10.4.4.4.2 | Chip-Select Signal Timing (CST _n) | 10-81 |
| 10.4.4.4.3 | Byte Select Signal Timing (BST _n)..... | 10-82 |
| 10.4.4.4.4 | General-Purpose Signals (GnTn, GOn)..... | 10-83 |
| 10.4.4.4.5 | Loop Control (LOOP) | 10-83 |
| 10.4.4.4.6 | Repeat Execution of Current RAM Word (REDO)..... | 10-84 |
| 10.4.4.4.7 | Address Multiplexing (AMX) | 10-84 |
| 10.4.4.4.8 | Data Valid and Data Sample Control (UTA) | 10-86 |
| 10.4.4.4.9 | LGPL[0:5] Signal Negation (LAST)..... | 10-86 |
| 10.4.4.4.10 | Wait Mechanism (WAEN)..... | 10-86 |
| 10.4.4.5 | Extended Hold Time on Read Accesses | 10-87 |
| 10.5 | Initialization/Application Information..... | 10-88 |
| 10.5.1 | Interfacing to Peripherals in Different Address Modes | 10-88 |
| 10.5.1.1 | Multiplexed Address/Data Bus for 26-Bit Addressing..... | 10-88 |
| 10.5.1.2 | Peripheral Hierarchy on the Local Bus for High Bus Speeds | 10-88 |
| 10.5.1.3 | GPCM Timings..... | 10-89 |
| 10.5.2 | Bus Turnaround | 10-90 |
| 10.5.2.1 | Address Phase after Previous Read | 10-90 |
| 10.5.2.2 | Read Data Phase after Address Phase | 10-90 |
| 10.5.2.3 | Read-Modify-Write Cycle for Parity Protected Memory Banks | 10-91 |
| 10.5.2.4 | UPM Cycles with Additional Address Phases..... | 10-91 |
| 10.5.3 | Interface to Different Port-Size Devices..... | 10-91 |
| 10.5.4 | Command Sequence Examples for NAND Flash EEPROM..... | 10-92 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|---|-------------|
| 10.5.4.1 | NAND Flash Soft Reset Command Sequence Example | 10-93 |
| 10.5.4.2 | NAND Flash Read Status Command Sequence Example | 10-93 |
| 10.5.4.3 | NAND Flash Read Identification Command Sequence Example | 10-93 |
| 10.5.4.4 | NAND Flash Page Read Command Sequence Example | 10-94 |
| 10.5.4.5 | NAND Flash Block Erase Command Sequence Example | 10-95 |
| 10.5.4.6 | NAND Flash Program Command Sequence Example | 10-95 |
| 10.5.5 | Interfacing to Fast-Page Mode DRAM Using UPM | 10-96 |
| 10.5.6 | Interfacing to ZBT SRAM Using UPM..... | 10-101 |

Chapter 11 Sequencer

| | | |
|----------|---|------|
| 11.1 | Sequencer Overview | 11-1 |
| 11.1.1 | Sequencer Features | 11-2 |
| 11.2 | Sequencer External Signal Description | 11-2 |
| 11.3 | Sequencer Memory Map/Register Definition..... | 11-2 |
| 11.4 | Sequencer Register Descriptions | 11-3 |
| 11.4.1 | PCI Outbound Translation Address Registers (POTAR _n)..... | 11-3 |
| 11.4.2 | PCI Outbound Base Address Registers (POBAR _n) | 11-3 |
| 11.4.3 | PCI Outbound Comparison Mask Registers (POCMR _n) | 11-4 |
| 11.4.4 | Power Management Control Register (PMCR) | 11-5 |
| 11.4.5 | Discard Timer Control Register (DTCR) | 11-6 |
| 11.5 | Functional Description..... | 11-6 |
| 11.5.1 | Transaction Forwarding..... | 11-6 |
| 11.5.1.1 | Transactions from the Coherency System Bus (CSB) Port | 11-7 |
| 11.5.1.2 | Transactions from the PCI Port | 11-7 |
| 11.5.1.3 | Transactions from the DMA Port | 11-7 |
| 11.5.2 | PCI Outbound Address Translation | 11-7 |
| 11.5.3 | Transaction Ordering | 11-8 |

Chapter 12 DMA/Messaging Unit

| | | |
|--------|--|------|
| 12.1 | DMA Features..... | 12-1 |
| 12.2 | DMA External Signal Description..... | 12-2 |
| 12.2.1 | DMA Detailed Signal Descriptions | 12-2 |
| 12.3 | DMA Memory Map/Register Definition | 12-3 |
| 12.4 | DMA Register Descriptions..... | 12-4 |
| 12.4.1 | Outbound Message Interrupt Status Register (OMISR) | 12-4 |
| 12.4.2 | Outbound Message Interrupt Mask Register (OMIMR)..... | 12-5 |
| 12.4.3 | Inbound Message Registers (IMR0–IMR1) | 12-6 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|---|-------------|
| 12.4.4 | Outbound Message Registers (OMR0–OMR1)..... | 12-6 |
| 12.4.5 | Doorbell Registers | 12-6 |
| 12.4.5.1 | Outbound Doorbell Register (ODR)..... | 12-7 |
| 12.4.5.2 | Inbound Doorbell Register (IDR)..... | 12-7 |
| 12.4.6 | Inbound Message Interrupt Status Register (IMISR) | 12-8 |
| 12.4.7 | Inbound Message Interrupt Mask Register (IMIMR)..... | 12-9 |
| 12.4.8 | DMA Registers | 12-10 |
| 12.4.8.1 | DMA Mode Register (DMAMR _n) | 12-10 |
| 12.4.8.2 | DMA Status Register (DMASR _n) | 12-12 |
| 12.4.8.3 | DMA Current Descriptor Address Register (DMACDAR _n) | 12-13 |
| 12.4.8.4 | DMA Source Address Register (DMASAR _n)..... | 12-14 |
| 12.4.8.5 | DMA Destination Address Register (DMADAR _n)..... | 12-14 |
| 12.4.8.6 | DMA Byte Count Register (DMABCR _n) | 12-15 |
| 12.4.8.7 | DMA Next Descriptor Address Register (DMANDAR _n)..... | 12-15 |
| 12.4.8.8 | DMA General Status Register (DMAGSR)..... | 12-16 |
| 12.5 | Functional Description..... | 12-16 |
| 12.5.1 | Message Unit | 12-16 |
| 12.5.1.1 | Messaging Registers (IMR0–IMR1, OMR0–OMR1) | 12-16 |
| 12.5.1.2 | Doorbell Registers (IDR and ODR) | 12-17 |
| 12.5.2 | DMA Controller..... | 12-17 |
| 12.5.3 | DMA Operation | 12-17 |
| 12.5.3.1 | External Control..... | 12-18 |
| 12.5.3.2 | DMA Coherency..... | 12-19 |
| 12.5.3.3 | Halt and Error Conditions..... | 12-20 |
| 12.5.4 | DMA Segment Descriptors..... | 12-20 |
| 12.5.4.1 | Descriptor in Big-Endian Mode..... | 12-21 |
| 12.5.4.2 | Descriptor in Little-Endian Mode..... | 12-22 |
| 12.6 | Initialization/Application Information | 12-22 |
| 12.6.1 | Initialization Steps in Direct Mode | 12-22 |
| 12.6.2 | Initialization Steps in Chaining Mode | 12-22 |
| 12.6.3 | Initialization Steps in Direct Mode with External Control | 12-23 |
| 12.6.4 | Initialization Steps in Chaining Mode with External Control | 12-23 |

Chapter 13 PCI Bus Interface

| | | |
|----------|-------------------------------------|------|
| 13.1 | PCI Introduction | 13-2 |
| 13.1.1 | PCI Features..... | 13-3 |
| 13.1.2 | PCI Modes of Operation | 13-3 |
| 13.1.2.1 | Host/Agent Mode Configuration | 13-3 |
| 13.1.2.2 | PCI Arbiter Configuration | 13-4 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 13.2 | PCI External Signal Description..... | 13-4 |
| 13.3 | PCI Memory Map/Register Definitions..... | 13-11 |
| 13.3.1 | PCI Configuration Access Registers..... | 13-12 |
| 13.3.1.1 | PCI_CONFIG_ADDRESS..... | 13-13 |
| 13.3.1.2 | PCI_CONFIG_DATA..... | 13-14 |
| 13.3.1.3 | PCI Interrupt Acknowledge Register (PCI_INT_ACK)..... | 13-15 |
| 13.3.2 | PCI Memory-Mapped Control and Status Registers | 13-15 |
| 13.3.2.1 | PCI Error Status Register (PCI_ESR) | 13-15 |
| 13.3.2.2 | PCI Error Capture Disable Register (PCI_ECDR)..... | 13-16 |
| 13.3.2.3 | PCI Error Enable Register (PCI_EER)..... | 13-17 |
| 13.3.2.4 | PCI Error Attributes Capture Register (PCI_EATCR) | 13-18 |
| 13.3.2.5 | PCI Error Address Capture Register (PCI_EACR)..... | 13-19 |
| 13.3.2.6 | PCI Error Extended Address Capture Register (PCI_EEACR) | 13-20 |
| 13.3.2.7 | PCI Error Data Low Capture Register (PCI_EDLCR)..... | 13-20 |
| 13.3.2.8 | PCI General Control Register (PCI_GCR)..... | 13-20 |
| 13.3.2.9 | PCI Error Control Register (PCI_ECR) | 13-21 |
| 13.3.2.10 | PCI General Status Register (PCI_GSR)..... | 13-22 |
| 13.3.2.11 | PCI Inbound Translation Address Registers (PITAR _n)..... | 13-22 |
| 13.3.2.12 | PCI Inbound Base Address Registers (PIBAR _n)..... | 13-23 |
| 13.3.2.13 | PCI Inbound Extended Base Address Registers (PIEBAR _n)..... | 13-24 |
| 13.3.2.14 | PCI Inbound Window Attribute Registers (PIWAR _n)..... | 13-24 |
| 13.3.3 | PCI Configuration Space Registers | 13-25 |
| 13.3.3.1 | Vendor ID Configuration Register..... | 13-27 |
| 13.3.3.2 | Device ID Configuration Register..... | 13-27 |
| 13.3.3.3 | PCI Command Configuration Register..... | 13-28 |
| 13.3.3.4 | PCI Status Configuration Register..... | 13-29 |
| 13.3.3.5 | Revision ID Configuration Register..... | 13-30 |
| 13.3.3.6 | Standard Programming Interface Configuration Register | 13-30 |
| 13.3.3.7 | Subclass Code Configuration Register | 13-31 |
| 13.3.3.8 | Base Class Code Configuration Register..... | 13-31 |
| 13.3.3.9 | Cache Line Size Configuration Register | 13-32 |
| 13.3.3.10 | Latency Timer Configuration Register | 13-32 |
| 13.3.3.11 | Header Type Configuration Register | 13-33 |
| 13.3.3.12 | BIST Control Configuration Register..... | 13-33 |
| 13.3.3.13 | PIMMR Base Address Configuration Register | 13-33 |
| 13.3.3.14 | GPL Base Address Register 0..... | 13-34 |
| 13.3.3.15 | GPL Base Address Registers 1–2..... | 13-34 |
| 13.3.3.16 | GPL Extended Base Address Registers 1–2..... | 13-35 |
| 13.3.3.17 | Subsystem Vendor ID Configuration Register | 13-36 |
| 13.3.3.18 | Subsystem Device ID Configuration Register..... | 13-36 |
| 13.3.3.19 | Capabilities Pointer Configuration Register..... | 13-36 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 13.3.3.20 | Interrupt Line Configuration Register | 13-37 |
| 13.3.3.21 | Interrupt Pin Configuration Register | 13-37 |
| 13.3.3.22 | Minimum Grant Configuration Register | 13-37 |
| 13.3.3.23 | Maximum Latency Configuration Register | 13-38 |
| 13.3.3.24 | PCI Function Configuration Register | 13-38 |
| 13.3.3.25 | PCI Arbiter Control Register (PCIACR) | 13-39 |
| 13.3.3.26 | Hot Swap Register Block..... | 13-40 |
| 13.3.3.27 | PCI Power Management Register 0 (PCIPMR0) | 13-41 |
| 13.3.3.28 | PCI Power Management Register 1 (PCIPMR1) | 13-42 |
| 13.4 | Functional Description..... | 13-43 |
| 13.4.1 | PCI Bus Arbitration | 13-43 |
| 13.4.1.1 | Bus Parking..... | 13-44 |
| 13.4.1.2 | Arbitration Algorithm..... | 13-44 |
| 13.4.1.3 | Broken Master Lock-Out..... | 13-45 |
| 13.4.1.4 | Master Latency Timer..... | 13-45 |
| 13.4.2 | Bus Commands | 13-46 |
| 13.4.3 | PCI Protocol Fundamentals | 13-47 |
| 13.4.3.1 | Basic Transfer Control..... | 13-47 |
| 13.4.3.2 | Addressing | 13-47 |
| 13.4.3.3 | Device Selection | 13-48 |
| 13.4.3.4 | Byte Enable Signals..... | 13-48 |
| 13.4.3.5 | Bus Driving and Turnaround | 13-48 |
| 13.4.3.6 | Bus Transactions..... | 13-49 |
| 13.4.3.7 | Read and Write Transactions | 13-49 |
| 13.4.3.8 | Transaction Termination | 13-51 |
| 13.4.4 | Other Bus Operations..... | 13-53 |
| 13.4.4.1 | Fast Back-to-Back Transactions | 13-53 |
| 13.4.4.2 | Dual Address Cycles..... | 13-54 |
| 13.4.4.3 | Data Streaming | 13-54 |
| 13.4.4.4 | Host Mode Configuration Access..... | 13-54 |
| 13.4.4.5 | Agent Mode Configuration Access | 13-55 |
| 13.4.4.6 | Special Cycle Command..... | 13-55 |
| 13.4.4.7 | Interrupt Acknowledge | 13-56 |
| 13.4.5 | Error Functions | 13-57 |
| 13.4.5.1 | Parity..... | 13-57 |
| 13.4.5.2 | Error Reporting..... | 13-57 |
| 13.4.6 | PCI Inbound Address Translation..... | 13-59 |
| 13.4.7 | CompactPCI Hot Swap Specification Support | 13-60 |
| 13.4.8 | Byte Ordering | 13-60 |
| 13.4.8.1 | Byte Order for Configuration Transactions | 13-61 |
| 13.5 | Initialization/Application Information | 13-62 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 13.5.1 | Initialization Sequence for Host Mode | 13-62 |
| 13.5.2 | Initialization Sequence for Agent Mode | 13-62 |

Chapter 14 PCI Express Interface Controller

| | | |
|------------|---|-------|
| 14.1 | Introduction | 14-1 |
| 14.1.1 | MPC8315E/MPC8314E as a PCI Express Initiator | 14-3 |
| 14.1.2 | MPC8315E/MPC8314E as a PCI Express Target | 14-3 |
| 14.1.3 | Features | 14-4 |
| 14.1.4 | Modes of Operation | 14-4 |
| 14.1.4.1 | Root Complex/Endpoint Modes | 14-4 |
| 14.1.4.2 | Link Width | 14-4 |
| 14.1.4.3 | Reference Clock | 14-5 |
| 14.2 | External Signal Descriptions | 14-5 |
| 14.3 | Memory Map/Register Definitions | 14-5 |
| 14.3.1 | PCI Express Memory Map | 14-5 |
| 14.4 | PCI Express Core Configuration Header Registers | 14-14 |
| 14.4.1 | Common PCI-Compatible Configuration Header Registers | 14-15 |
| 14.4.1.1 | PCI Express Vendor ID Register | 14-15 |
| 14.4.1.2 | PCI Express Device ID Register | 14-16 |
| 14.4.1.3 | PCI Express Command Register | 14-16 |
| 14.4.1.4 | PCI Express Status Register | 14-17 |
| 14.4.1.5 | PCI Express Revision ID Register | 14-18 |
| 14.4.1.6 | PCI Express Class Code Register | 14-19 |
| 14.4.1.7 | PCI Express Cache Line Size Register | 14-19 |
| 14.4.1.8 | PCI Express Latency Timer Register | 14-20 |
| 14.4.1.9 | PCI Express Header Type Register | 14-20 |
| 14.4.1.10 | PCI Express BIST Register | 14-21 |
| 14.4.2 | Type 0 PCI-Compatible Configuration Header Registers | 14-21 |
| 14.4.2.1 | PCI Express Base Address Registers (EP Mode Only) | 14-22 |
| 14.4.2.1.1 | Base Address Registers 0 and 1 (BAR0/BAR1) | 14-22 |
| 14.4.2.1.2 | Base Address Registers 2 and 4 (BAR2/BAR4) | 14-23 |
| 14.4.2.1.3 | Base Address Registers 3 and 5 (BAR3/BAR5) | 14-23 |
| 14.4.2.2 | PCI Express Subsystem Vendor ID Register (EP Mode Only) | 14-24 |
| 14.4.2.3 | PCI Express Subsystem ID Register (EP Mode Only) | 14-24 |
| 14.4.2.4 | PCI Express Capabilities Pointer Register | 14-25 |
| 14.4.2.5 | PCI Express Interrupt Line Register (EP-Mode Only) | 14-25 |
| 14.4.2.6 | PCI Express Interrupt Pin Register | 14-26 |
| 14.4.2.7 | PCI Express Minimum Grant Register (EP Mode Only) | 14-26 |
| 14.4.2.8 | PCI Express Maximum Latency Register (EP Mode Only) | 14-26 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 14.4.3 | Type 1 PCI-Compatible Configuration Header Registers | 14-27 |
| 14.4.3.1 | PCI Express Primary Bus Number Register (RC Mode Only)..... | 14-27 |
| 14.4.3.2 | PCI Express Secondary Bus Number Register (RC Mode Only)..... | 14-28 |
| 14.4.3.3 | PCI Express Subordinate Bus Number Register (RC Mode Only) | 14-28 |
| 14.4.3.4 | PCI Express Secondary Latency Timer Register (RC Mode Only) | 14-29 |
| 14.4.3.5 | PCI Express I/O Base Register (RC Mode Only)..... | 14-29 |
| 14.4.3.6 | PCI Express I/O Limit Register (RC Mode Only)..... | 14-29 |
| 14.4.3.7 | PCI Express Secondary Status Register (RC Mode Only) | 14-30 |
| 14.4.3.8 | PCI Express Memory Base Register (RC Mode Only) | 14-31 |
| 14.4.3.9 | PCI Express Memory Limit Register (RC Mode Only) | 14-31 |
| 14.4.3.10 | PCI Express Prefetchable Memory Base Register (RC Mode Only) | 14-32 |
| 14.4.3.11 | PCI Express Prefetchable Memory Limit Register (RC Mode Only) | 14-32 |
| 14.4.3.12 | PCI Express Prefetchable Base Upper 32-Bit Register (RC Mode Only)..... | 14-33 |
| 14.4.3.13 | PCI Express Prefetchable Limit Upper 32-Bit Register (RC Mode Only)..... | 14-33 |
| 14.4.3.14 | PCI Express I/O Base Upper 16-Bit Register (RC Mode Only) | 14-33 |
| 14.4.3.15 | PCI Express I/O Limit Upper 16-Bit Register (RC Mode Only) | 14-34 |
| 14.4.3.16 | PCI Express Capabilities Pointer Register | 14-34 |
| 14.4.3.17 | PCI Express Interrupt Line Register..... | 14-35 |
| 14.4.3.18 | PCI Express Interrupt Pin Register..... | 14-35 |
| 14.4.3.19 | PCI Express Bridge Control Register (RC Mode Only)..... | 14-35 |
| 14.4.4 | PCI Compatible Device-Specific Configuration Space Registers..... | 14-37 |
| 14.4.4.1 | PCI Express Power Management Capability ID Register | 14-38 |
| 14.4.4.2 | PCI Express Power Management Next Capabilities Pointer Register..... | 14-38 |
| 14.4.4.3 | PCI Express Power Management Capabilities Register..... | 14-39 |
| 14.4.4.4 | PCI Express Power Management Status and Control Register | 14-39 |
| 14.4.4.5 | PCI Express Power Management Data Register | 14-40 |
| 14.4.4.6 | PCI Express Capability ID Register | 14-40 |
| 14.4.4.7 | PCI Express Next Capabilities Pointer Register..... | 14-41 |
| 14.4.4.8 | PCI Express Capabilities Register | 14-41 |
| 14.4.4.9 | PCI Express Device Capabilities Register..... | 14-42 |
| 14.4.4.10 | PCI Express Device Control Register..... | 14-43 |
| 14.4.4.11 | PCI Express Device Status Register | 14-44 |
| 14.4.4.12 | PCI Express Link Capabilities Register | 14-44 |
| 14.4.4.13 | PCI Express Link Control Register | 14-45 |
| 14.4.4.14 | PCI Express Link Status Register | 14-46 |
| 14.4.4.15 | PCI Express Slot Capabilities Register..... | 14-46 |
| 14.4.4.16 | PCI Express Slot Control Register | 14-47 |
| 14.4.4.17 | PCI Express Slot Status Register..... | 14-48 |
| 14.4.4.18 | PCI Express Root Control Register (RC Mode Only)..... | 14-48 |
| 14.4.4.19 | PCI Express Root Status Register (RC Mode Only) | 14-49 |
| 14.4.4.20 | PCI Express MSI Message Capability ID Register (EP Mode Only) | 14-49 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 14.4.4.21 | PCI Express MSI Message Control Register (EP Mode Only) | 14-50 |
| 14.4.4.22 | PCI Express MSI Message Address Register (EP Mode Only) | 14-50 |
| 14.4.4.23 | PCI Express MSI Message Upper Address Register (EP Mode Only) | 14-51 |
| 14.4.4.24 | PCI Express MSI Message Data Register (EP Mode Only)..... | 14-51 |
| 14.4.5 | PCI Express Extended Configuration Space | 14-52 |
| 14.4.5.1 | PCI Express Advanced Error Reporting Capability ID Register..... | 14-53 |
| 14.4.5.2 | PCI Express Uncorrectable Error Status Register | 14-53 |
| 14.4.5.3 | PCI Express Uncorrectable Error Mask Register | 14-54 |
| 14.4.5.4 | PCI Express Uncorrectable Error Severity Register..... | 14-55 |
| 14.4.5.5 | PCI Express Correctable Error Status Register | 14-56 |
| 14.4.5.6 | PCI Express Correctable Error Mask Register | 14-57 |
| 14.4.5.7 | PCI Express Advanced Error Capabilities and Control Register | 14-57 |
| 14.4.5.8 | PCI Express Header Log Register | 14-58 |
| 14.4.5.9 | PCI Express Root Error Command Register | 14-59 |
| 14.4.5.10 | PCI Express Root Error Status Register | 14-59 |
| 14.4.5.11 | PCI Express Error Source Identification Register | 14-60 |
| 14.4.6 | PCI Express Controller Internal Control and Status Registers (CSRs) | 14-61 |
| 14.4.6.1 | PCI Express LTSSM State Status Register (PEX_LTSSM_STAT)..... | 14-61 |
| 14.4.6.2 | PCI Express N_FTS Control Register (PEX_NFTS_CTRL)..... | 14-62 |
| 14.4.6.3 | PCI Express ACK Replay Timeout Register (PEX_ACKRPLY_TO) | 14-63 |
| 14.4.6.4 | PCI Express Controller Core Clock Ratio Register (PEX_GCLK_RATIO)..... | 14-64 |
| 14.4.6.5 | PCI Express Power Management Timer Register (PEX_PM_TIMER)..... | 14-65 |
| 14.4.6.6 | PCI Express PME Time-Out Register (PEX_PME_TIMEOUT) (EP Mode Only)..... | 14-66 |
| 14.4.6.7 | PCI Express ASPM Request Timer Register (PEX_ASPM_REQTMR)..... | 14-66 |
| 14.4.6.8 | PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE) | 14-67 |
| 14.4.6.9 | PCI Express Device Capabilities Update Register (PEX_DEVCAP_UPDATE) | 14-68 |
| 14.4.6.10 | PCI Express Link Capabilities Update Register (PEX_LINKCAP_UPDATE) | 14-69 |
| 14.4.6.11 | PCI Express Slot Capabilities Update Register (PEX_SLCAP_UPDATE) | 14-70 |
| 14.4.6.12 | PCI Express Configuration Ready Register | 14-71 |
| 14.4.7 | PCI Express BAR Configuration Registers (EP Mode) | 14-71 |
| 14.4.7.1 | PCI Express BAR Enable Register (PEX_BAR_ENABLE)..... | 14-71 |
| 14.4.7.2 | PCI Express BAR Size Low Configuration Register (PEX_BAR_SIZEL)..... | 14-72 |
| 14.4.7.3 | PCI Express BAR Select Configuration Register (PEX_BAR_SEL) | 14-73 |
| 14.4.7.4 | PCI Express BAR Prefetch Configuration Register (PEX_BAR_PF) | 14-73 |
| 14.4.8 | PCI Express Extended Status and Control Registers..... | 14-74 |
| 14.4.8.1 | PME_To_Ack Timeout Register (RC Mode Only) | 14-74 |
| 14.4.8.2 | PME_To_Ack Status Register (RC Mode Only)..... | 14-75 |
| 14.4.8.3 | Secondary Status Interrupt Mask Register (PEX_SS_INTR_MASK) (RC Mode Only)..... | 14-76 |
| 14.5 | PCI Express CSB Bridge | 14-76 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 14.5.1 | PCI Express CSB Bridge Configuration Space | 14-76 |
| 14.5.2 | Global Registers..... | 14-77 |
| 14.5.2.1 | PCI Express CSB Bridge Control Register (PEX_CSB_CTRL) | 14-77 |
| 14.5.2.2 | PCI Express DMA Descriptor Timer Register (PEX_DMA_DSTMR)..... | 14-78 |
| 14.5.2.3 | PCI Express CSB Bridge Status Register (PEX_CSB_STAT)..... | 14-78 |
| 14.5.3 | PCI Express Outbound PIO Registers | 14-79 |
| 14.5.3.1 | PCI Express Outbound PIO Control Register (PEX_CSB_OBCTRL)..... | 14-79 |
| 14.5.3.2 | PCI Express Outbound PIO Status Register (PEX_CSB_OBSTAT) | 14-80 |
| 14.5.4 | PCI Express Inbound PIO Registers..... | 14-81 |
| 14.5.4.1 | PCI Express Inbound PIO Control Register (PEX_CSB_IBCTRL) | 14-81 |
| 14.5.4.2 | PCI Express Inbound PIO Status Register (PEX_CSB_IBSTAT)..... | 14-81 |
| 14.5.5 | DMA Registers | 14-82 |
| 14.5.5.1 | PCI Express Write DMA Control Register (PEX_WDMA_CTRL) | 14-82 |
| 14.5.5.2 | PCI Express Write DMA First Address Register (PEX_WDMA_ADDR)..... | 14-83 |
| 14.5.5.3 | PCI Express Write DMA Status Register (PEX_WDMA_STAT)..... | 14-83 |
| 14.5.5.4 | PCI Express Read DMA Control Register (PEX_RDMA_CTRL)..... | 14-84 |
| 14.5.5.5 | PCI Express Read DMA First Address Register (PEX_RDMA_ADDR)..... | 14-85 |
| 14.5.5.6 | PCI Express Read DMA Status Register (PEX_RDMA_STAT) | 14-85 |
| 14.5.6 | Mailbox Registers | 14-86 |
| 14.5.6.1 | PCI Express Outbound Mailbox Control Register (PEX_OMBCR)..... | 14-86 |
| 14.5.6.2 | PCI Express Outbound Mailbox Data Register (PEX_OMBDR) | 14-87 |
| 14.5.6.3 | PCI Express Inbound Mailbox Control Register (PEX_IMBCR)..... | 14-87 |
| 14.5.6.4 | PCI Express Inbound Mailbox Data Register (PEX_IMBDR) | 14-88 |
| 14.5.7 | PCI Express Host Interrupt Registers | 14-88 |
| 14.5.7.1 | PCI Express Host Interrupt Enable Register (PEX_HIER)..... | 14-89 |
| 14.5.7.2 | PCI Express Host Interrupt Status Register (PEX_HISR)..... | 14-90 |
| 14.5.7.3 | PCI Express Host Outbound PIO Interrupt Vector Register (PEX_HOPIVR).... | 14-91 |
| 14.5.7.4 | PCI Express Host Inbound PIO Interrupt Vector Register (PEX_HIPIVR)..... | 14-91 |
| 14.5.7.5 | PCI Express Host Write DMA Interrupt Vector Register (PEX_HWDIVR) | 14-92 |
| 14.5.7.6 | PCI Express Host Read DMA Interrupt Vector Register (PEX_HRDIVR) | 14-92 |
| 14.5.7.7 | PCI Express Host Miscellaneous Interrupt Vector Register (PEX_HMIVR) | 14-93 |
| 14.5.8 | CSB System Interrupt Registers | 14-93 |
| 14.5.8.1 | CSB System PIO Interrupt Enable Register (PEX_CSPIER) | 14-93 |
| 14.5.8.2 | CSB System Write DMA Interrupt Enable Register (PEX_CSWDIER) | 14-94 |
| 14.5.8.3 | CSB System Read DMA Interrupt Enable Register (PEX_CSRDIER) | 14-95 |
| 14.5.8.4 | CSB System Miscellaneous Interrupt Enable Register (PEX_CSMIER) | 14-95 |
| 14.5.8.5 | CSB System PIO Interrupt Status Register (PEX_CSPISR) | 14-97 |
| 14.5.8.6 | CSB System Write DMA Interrupt Status Register (PEX_CSWDISR)..... | 14-97 |
| 14.5.8.7 | CSB System Read DMA Interrupt Status Register (PEX_CSRDISR) | 14-98 |
| 14.5.8.8 | CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR)..... | 14-99 |
| 14.5.9 | PCI Express Power Management Registers..... | 14-100 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 14.5.9.1 | PCI Express Power Management Control Register (PEX_PM_CTRL)..... | 14-100 |
| 14.5.9.2 | PCI Express Slot Control Misc Register | 14-101 |
| 14.5.10 | PCI Express Outbound Address Mapping Registers | 14-102 |
| 14.5.10.1 | PCI Express Outbound Window Attributes Register <i>n</i> (PEX_OWAR0–PEX_OWAR3)..... | 14-102 |
| 14.5.10.2 | PCI Express Outbound Window Base Address Register <i>n</i> (PEX_OWBAR0–PEX_OWBAR3)..... | 14-103 |
| 14.5.10.3 | PCI Express Outbound Window Translation Address Register Low <i>n</i> (PEX_OWTARL0–PEX_OWTARL3)..... | 14-103 |
| 14.5.10.4 | PCI Express Outbound Window Translation Address Register High <i>n</i> (PEX_OWTARH0–PEX_OWTARH3)..... | 14-104 |
| 14.5.11 | PCI Express EP Inbound Address Translation Registers | 14-105 |
| 14.5.11.1 | PCI Express EP Inbound Window Translation Address Register <i>n</i> (PEX_EPIWTAR0–PEX_EPIWTAR3) | 14-105 |
| 14.5.12 | PCI Express RC Inbound Address Mapping Registers | 14-106 |
| 14.5.12.1 | PCI Express RC Inbound Window Attributes Register <i>n</i> (PEX_RCIWAR0 –PEX_RCIWAR3) | 14-106 |
| 14.5.12.2 | PCI Express RC Inbound Window Translation Address Register <i>n</i> (PEX_RCIWTAR0–PEX_RCIWTAR3) | 14-107 |
| 14.5.12.3 | PCI Express RC Inbound Window Base Address Register Low <i>n</i> (PEX_RCIWBARL0–PEX_RCIWBARL3) | 14-108 |
| 14.5.12.4 | PCI Express RC Inbound Window Base Address Register High <i>n</i> (PEX_RCIWBARH0–PEX_RCIWBARH3) | 14-108 |
| 14.6 | Functional Description..... | 14-109 |
| 14.6.1 | Architecture | 14-110 |
| 14.6.1.1 | Address Translation Windows (ATMUs) | 14-110 |
| 14.6.1.2 | Byte Ordering | 14-111 |
| 14.6.1.2.1 | Byte Order for Configuration Transactions..... | 14-113 |
| 14.6.1.3 | Transaction Ordering Rule..... | 14-113 |
| 14.6.1.4 | Memory Space Addressing..... | 14-114 |
| 14.6.1.5 | I/O Space Addressing | 14-114 |
| 14.6.1.6 | Configuration Space Access | 14-114 |
| 14.6.1.6.1 | Outbound ATMU Configuration Transaction Generation (RC)..... | 14-114 |
| 14.6.1.6.2 | EP Configuration Register Access | 14-118 |
| 14.6.1.7 | Messages..... | 14-118 |
| 14.6.1.7.1 | Inbound Messages | 14-118 |
| 14.6.2 | Interrupts..... | 14-121 |
| 14.6.2.1 | EP Interrupt Generation..... | 14-121 |
| 14.6.2.1.1 | Hardware INTx Message Generation | 14-121 |
| 14.6.2.1.2 | Software INTx Message Generation | 14-121 |
| 14.6.2.1.3 | Hardware MSI Generation..... | 14-121 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|---|-------------|
| 14.6.2.1.4 | Software MSI Generation | 14-122 |
| 14.6.2.2 | RC Handling of INTx Message and MSI Interrupt | 14-122 |
| 14.6.2.2.1 | INTx Message Handling | 14-122 |
| 14.6.2.2.2 | MSI Handling | 14-122 |
| 14.6.2.3 | Initial Credit Advertisement | 14-123 |
| 14.6.3 | Mailbox | 14-123 |
| 14.6.3.1 | Outbound Mailbox | 14-123 |
| 14.6.3.2 | Inbound Mailbox | 14-124 |
| 14.6.4 | Power Management | 14-124 |
| 14.6.4.1 | L2/L3 Ready Link State | 14-125 |
| 14.6.5 | Hot Reset | 14-125 |
| 14.6.6 | Initialization Sequence | 14-126 |
| 14.7 | DMA Functional Operation | 14-127 |
| 14.7.1 | DMA Descriptor Format | 14-127 |
| 14.7.2 | Write DMA | 14-129 |
| 14.7.3 | Read DMA | 14-130 |
| 14.7.4 | Descriptor-Based DMA | 14-131 |
| 14.7.4.1 | Chain Descriptor | 14-131 |
| 14.7.4.2 | Block Descriptors | 14-132 |
| 14.7.4.3 | Descriptor Format | 14-132 |
| 14.7.4.4 | Software-Hardware Handshake | 14-133 |

Chapter 15 SATA Controller

| | | |
|----------|---|------|
| 15.1 | Overview | 15-1 |
| 15.2 | Command Operation | 15-2 |
| 15.2.1 | Command Issue | 15-2 |
| 15.2.2 | Command Service | 15-3 |
| 15.2.3 | Command Completion Interrupt Timing | 15-3 |
| 15.2.4 | DMA Context (Read Data) | 15-3 |
| 15.2.5 | DMA Context (Write Data) | 15-3 |
| 15.2.6 | DMAT Primitive Processing | 15-3 |
| 15.3 | Command Layer Overview | 15-3 |
| 15.3.1 | SATA Memory Map/Register Definition | 15-4 |
| 15.3.2 | Command Registers | 15-5 |
| 15.3.2.1 | Command Queue Register (CQR) | 15-5 |
| 15.3.2.2 | Command Active Register (CAR) | 15-6 |
| 15.3.2.3 | Command Completed Register (CCR) | 15-6 |
| 15.3.2.4 | Command Error Register (CER) | 15-7 |
| 15.3.2.5 | Device Error Register (DER) | 15-8 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 15.3.2.6 | Command Header Base Address Register (CHBA) | 15-8 |
| 15.3.2.7 | Host Status Register (HStatus) | 15-9 |
| 15.3.2.7.1 | Error Processing | 15-11 |
| 15.3.2.8 | Host Control Register (HControl) | 15-12 |
| 15.3.2.8.1 | Bringing the SATA Controller Online/Offline..... | 15-13 |
| 15.3.2.9 | Port Number Queue Register (CQPMP)..... | 15-13 |
| 15.3.2.10 | Signature Register (SIG)..... | 15-14 |
| 15.3.2.11 | Interrupt Coalescing Control Register (ICC)..... | 15-14 |
| 15.3.3 | SATA Superset Registers | 15-15 |
| 15.3.3.1 | SATA Interface Status Register (SStatus)..... | 15-15 |
| 15.3.3.2 | SATA Interface Error Register (SError) | 15-16 |
| 15.3.3.3 | SATA Interface Control Register (SControl)..... | 15-18 |
| 15.3.3.4 | SATA Interface Notification Register (SNotification)..... | 15-19 |
| 15.3.4 | Control Status Registers..... | 15-20 |
| 15.3.4.1 | Transport Layer Configuration Register (TransCfg) | 15-20 |
| 15.3.4.2 | Transport Layer Status Register (TransStatus) | 15-21 |
| 15.3.4.3 | Link Layer Configuration Register (LinkCfg) | 15-21 |
| 15.3.4.4 | Link Layer Configuration Register1 (LinkCfg1) | 15-22 |
| 15.3.4.5 | Link Layer Configuration Register2 (LinkCfg2) | 15-23 |
| 15.3.4.6 | Link Layer Status Register (LinkStatus) | 15-23 |
| 15.3.4.7 | Link Layer Status Register1 (LinkStatus1) | 15-24 |
| 15.3.4.8 | PHY Control Configuration Register1 (PhyCtrlCfg1) | 15-26 |
| 15.3.4.9 | Link Layer Command Status Register (CommandStatus)..... | 15-27 |
| 15.3.4.10 | PHY Control Configuration Register2 (PhyCtrlCfg2) | 15-29 |
| 15.3.5 | System Control Registers..... | 15-30 |
| 15.3.5.1 | System Priority Register (SYSPR) | 15-30 |
| 15.3.6 | Command Header | 15-31 |
| 15.3.7 | Command Descriptor | 15-34 |
| 15.3.7.1 | Command FIS Non-Queued Commands (CFIS)..... | 15-34 |
| 15.3.7.2 | Command FIS First Party DMA Commands NCQ | 15-35 |
| 15.3.7.3 | Status FIS (SFIS) | 15-35 |
| 15.3.7.4 | ATAPI Command (ACMD) | 15-35 |
| 15.3.7.5 | Physical Region Descriptor Table (PRDT)..... | 15-35 |
| 15.3.8 | Vendor-Specific BIST Operation..... | 15-36 |
| 15.4 | Transport Layer Architectural Overview | 15-38 |
| 15.5 | Link Layer Overview | 15-39 |
| 15.5.1 | Link Layer functionality | 15-40 |
| 15.5.1.1 | Link Idle State Machine..... | 15-40 |
| 15.5.1.2 | Transmit State Machine | 15-40 |
| 15.5.1.3 | Receive State Machine..... | 15-41 |
| 15.5.1.4 | Power Mode Change State Machine..... | 15-41 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 15.5.1.5 | Frame Content Scrambler and Descrambler | 15-42 |
| 15.5.1.6 | CRC Generator and Checker | 15-42 |
| 15.5.1.7 | 8B/10B Encode and Decode | 15-43 |
| 15.5.1.8 | CONT Primitive Processing | 15-43 |
| 15.5.1.9 | ALIGN Insertion | 15-43 |
| 15.5.1.10 | Debug Functionality | 15-43 |
| 15.5.1.11 | BIST Support | 15-44 |
| 15.6 | PHY Control Layer Overview | 15-44 |
| 15.7 | Initialization/Application Information | 15-44 |
| 15.7.1 | SATA Controller Initialization Steps | 15-44 |

Chapter 16 SerDes PHY

| | | |
|--------|---|-------|
| 16.1 | Introduction | 16-1 |
| 16.1.1 | Overview | 16-1 |
| 16.1.2 | Features | 16-2 |
| 16.1.3 | Modes of Operation | 16-2 |
| 16.1.4 | Clock | 16-2 |
| 16.2 | External Signals | 16-3 |
| 16.3 | Memory Map/Registers | 16-4 |
| 16.3.1 | SerDes Control Register 0 (SRDSCR0) | 16-5 |
| 16.3.2 | SerDes Control Register 1 (SRDSCR1) | 16-8 |
| 16.3.3 | SerDes Control Register 2 (SRDSCR2) | 16-9 |
| 16.3.4 | SerDes Control Register 3 (SRDSCR3) | 16-10 |
| 16.3.5 | SerDes Control Register 4 (SRDSCR4) | 16-11 |
| 16.3.6 | SerDesn Reset Control Register (SRDSRSTCTL) | 16-13 |
| 16.4 | Initialization Sequence and Reset | 16-13 |
| 16.5 | Power Management: Power Down | 16-14 |

Chapter 17 Universal Serial Bus Interface

| | | |
|--------|--------------------------|------|
| 17.1 | Introduction | 17-1 |
| 17.1.1 | Overview | 17-2 |
| 17.1.2 | Features | 17-2 |
| 17.1.3 | Modes of Operation | 17-3 |
| 17.2 | External Signals | 17-3 |
| 17.2.1 | UTMI Interface | 17-4 |
| 17.2.2 | ULPI Interface | 17-5 |
| 17.2.3 | PHY Clocks | 17-6 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 17.3 | Memory Map/Register Definitions | 17-6 |
| 17.3.1 | Capability Registers | 17-8 |
| 17.3.1.1 | Capability Registers Length (CAPLENGTH) | 17-9 |
| 17.3.1.2 | Host Controller Interface Version (HCIVERSION) | 17-9 |
| 17.3.1.3 | Host Controller Structural Parameters (HCSPARAMS) | 17-10 |
| 17.3.1.4 | Host Controller Capability Parameters (HCCPARAMS) | 17-10 |
| 17.3.1.5 | Device Controller Interface Version (DCIVERSION)—Non-EHCI | 17-11 |
| 17.3.1.6 | Device Controller Capability Parameters (DCCPARAMS)—Non-EHCI | 17-12 |
| 17.3.2 | Operational Registers | 17-12 |
| 17.3.2.1 | USB Command Register (USBCMD) | 17-13 |
| 17.3.2.2 | USB Status Register (USBSTS) | 17-15 |
| 17.3.2.3 | USB Interrupt Enable Register (USBINTR) | 17-17 |
| 17.3.2.4 | Frame Index Register (FRINDEX) | 17-19 |
| 17.3.2.5 | Control Data Structure Segment Register (CTRLDSSEGMENT) | 17-20 |
| 17.3.2.6 | Periodic Frame List Base Address Register (PERIODICLISTBASE) | 17-20 |
| 17.3.2.7 | Device Address Register (DEVICEADDR)—Non-EHCI | 17-21 |
| 17.3.2.8 | Current Asynchronous List Address Register (ASYNCLISTADDR) | 17-22 |
| 17.3.2.9 | Endpoint List Address Register (ENDPOINTLISTADDR)—Non-EHCI | 17-22 |
| 17.3.2.10 | Master Interface Data Burst Size Register (BURSTSIZE)—Non-EHCI | 17-23 |
| 17.3.2.11 | Transmit FIFO Tuning Controls Register (TXFILLTUNING)—Non-EHCI | 17-24 |
| 17.3.2.12 | ULPI Register Access (ULPI VIEWPORT) | 17-25 |
| 17.3.2.13 | Configure Flag Register (CONFIGFLAG) | 17-27 |
| 17.3.2.14 | Port Status and Control Register (PORTSC) | 17-27 |
| 17.3.2.15 | On-The-Go Status and Control (OTGSC)—Non-EHCI | 17-32 |
| 17.3.2.16 | USB Mode Register (USBMODE)—Non-EHCI | 17-35 |
| 17.3.2.17 | Endpoint Setup Status Register (ENDPTSETUPSTAT)—Non-EHCI | 17-36 |
| 17.3.2.18 | Endpoint Initialization Register (ENDPTPRIME)—Non-EHCI | 17-36 |
| 17.3.2.19 | Endpoint Flush Register (ENDPTFLUSH)—Non-EHCI | 17-37 |
| 17.3.2.20 | Endpoint Status Register (ENDPTSTATUS)—Non-EHCI | 17-38 |
| 17.3.2.21 | Endpoint Complete Register (ENDPTCOMPLETE)—Non-EHCI | 17-38 |
| 17.3.2.22 | Endpoint Control Register 0 (ENDPTCTRL0)—Non-EHCI | 17-39 |
| 17.3.2.23 | Endpoint Control Register <i>n</i> (ENDPTCTRL <i>n</i>)—Non-EHCI | 17-40 |
| 17.3.2.24 | SNOOP1 and SNOOP2—Non-EHCI | 17-42 |
| 17.3.2.25 | Age Count Threshold Register (AGE_CNT_THRESH)—Non-EHCI | 17-43 |
| 17.3.2.26 | Priority Control Register (PRI_CTRL)—Non-EHCI | 17-44 |
| 17.3.2.27 | System Interface Control Register (SI_CTRL)—Non-EHCI | 17-45 |
| 17.3.2.28 | USB General Purpose Register (CONTROL)—Non-EHCI | 17-46 |
| 17.4 | Functional Description | 17-48 |
| 17.4.1 | System Interface | 17-48 |
| 17.4.2 | DMA Engine | 17-48 |
| 17.4.3 | FIFO RAM Controller | 17-48 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 17.4.4 | PHY Interface | 17-49 |
| 17.5 | Host Data Structures | 17-49 |
| 17.5.1 | Periodic Frame List..... | 17-50 |
| 17.5.2 | Asynchronous List Queue Head Pointer..... | 17-51 |
| 17.5.3 | Isochronous (High-Speed) Transfer Descriptor (iTd)..... | 17-51 |
| 17.5.3.1 | Next Link Pointer | 17-52 |
| 17.5.3.2 | iTd Transaction Status and Control List | 17-53 |
| 17.5.3.3 | iTd Buffer Page Pointer List (Plus) | 17-53 |
| 17.5.4 | Split Transaction Isochronous Transfer Descriptor (siTD)..... | 17-55 |
| 17.5.4.1 | Next Link Pointer | 17-55 |
| 17.5.4.2 | siTD Endpoint Capabilities/Characteristics..... | 17-56 |
| 17.5.4.3 | siTD Transfer State | 17-57 |
| 17.5.4.4 | siTD Buffer Pointer List (Plus)..... | 17-58 |
| 17.5.4.5 | siTD Back Link Pointer | 17-58 |
| 17.5.5 | Queue Element Transfer Descriptor (qTD) | 17-59 |
| 17.5.5.1 | Next qTD Pointer..... | 17-60 |
| 17.5.5.2 | Alternate Next qTD Pointer | 17-60 |
| 17.5.5.3 | qTD Token | 17-60 |
| 17.5.5.4 | qTD Buffer Page Pointer List | 17-64 |
| 17.5.6 | Queue Head..... | 17-65 |
| 17.5.6.1 | Queue Head Horizontal Link Pointer | 17-65 |
| 17.5.6.2 | Endpoint Capabilities/Characteristics..... | 17-66 |
| 17.5.6.3 | Transfer Overlay | 17-68 |
| 17.5.7 | Periodic Frame Span Traversal Node (FSTN)..... | 17-69 |
| 17.5.7.1 | FSTN Normal Path Pointer..... | 17-70 |
| 17.5.7.2 | FSTN Back Path Link Pointer | 17-70 |
| 17.6 | Host Operations | 17-71 |
| 17.6.1 | Host Controller Initialization | 17-71 |
| 17.6.2 | Power Port..... | 17-72 |
| 17.6.3 | Reporting Over-Current | 17-72 |
| 17.6.4 | Suspend/Resume..... | 17-73 |
| 17.6.4.1 | Port Suspend/Resume | 17-73 |
| 17.6.5 | Schedule Traversal Rules..... | 17-75 |
| 17.6.6 | Periodic Schedule Frame Boundaries vs. Bus Frame Boundaries..... | 17-76 |
| 17.6.7 | Periodic Schedule | 17-79 |
| 17.6.8 | Managing Isochronous Transfers Using iTDs | 17-80 |
| 17.6.8.1 | Host Controller Operational Model for iTDs | 17-81 |
| 17.6.8.2 | Software Operational Model for iTDs | 17-82 |
| 17.6.8.2.1 | Periodic Scheduling Threshold..... | 17-84 |
| 17.6.9 | Asynchronous Schedule..... | 17-85 |
| 17.6.9.1 | Adding Queue Heads to Asynchronous Schedule | 17-86 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 17.6.9.2 | Removing Queue Heads from Asynchronous Schedule..... | 17-86 |
| 17.6.9.3 | Empty Asynchronous Schedule Detection | 17-88 |
| 17.6.9.4 | Asynchronous Schedule Traversal: Start Event..... | 17-89 |
| 17.6.9.5 | Reclamation Status Bit (USBSTS Register)..... | 17-89 |
| 17.6.10 | Managing Control/Bulk/Interrupt Transfers via Queue Heads..... | 17-89 |
| 17.6.10.1 | Buffer Pointer List Use for Data Streaming with qTDs | 17-90 |
| 17.6.10.2 | Adding Interrupt Queue Heads to the Periodic Schedule..... | 17-92 |
| 17.6.10.3 | Managing Transfer Complete Interrupts from Queue Heads | 17-92 |
| 17.6.11 | Ping Control..... | 17-93 |
| 17.6.12 | Split Transactions..... | 17-94 |
| 17.6.12.1 | Split Transactions for Asynchronous Transfers..... | 17-94 |
| 17.6.12.1.1 | Asynchronous—Do-Start-Split..... | 17-95 |
| 17.6.12.1.2 | Asynchronous—Do-Complete-Split | 17-95 |
| 17.6.12.2 | Split Transaction Interrupt | 17-96 |
| 17.6.12.2.1 | Split Transaction Scheduling Mechanisms for Interrupt | 17-96 |
| 17.6.12.2.2 | Host Controller Operational Model for FSTNs..... | 17-99 |
| 17.6.12.2.3 | Software Operational Model for FSTNs | 17-101 |
| 17.6.12.2.4 | Tracking Split Transaction Progress for Interrupt Transfers | 17-102 |
| 17.6.12.2.5 | Split Transaction Execution State Machine for Interrupt | 17-102 |
| 17.6.12.2.6 | Periodic Interrupt—Do-Start-Split | 17-103 |
| 17.6.12.2.7 | Periodic Interrupt—Do-Complete-Split | 17-104 |
| 17.6.12.2.8 | Managing the QH[FrameTag] Field | 17-107 |
| 17.6.12.2.9 | Rebalancing the Periodic Schedule | 17-108 |
| 17.6.12.3 | Split Transaction Isochronous | 17-108 |
| 17.6.12.3.1 | Split Transaction Scheduling Mechanisms for Isochronous | 17-109 |
| 17.6.12.3.2 | Tracking Split Transaction Progress for Isochronous Transfers..... | 17-113 |
| 17.6.12.3.3 | Split Transaction Execution State Machine for Isochronous..... | 17-114 |
| 17.6.12.3.4 | Periodic Isochronous—Do-Start-Split..... | 17-115 |
| 17.6.12.3.5 | Periodic Isochronous—Do Complete Split | 17-117 |
| 17.6.12.3.6 | Complete-Split for Scheduling Boundary Cases 2a, 2b | 17-120 |
| 17.6.12.3.7 | Split Transaction for Isochronous—Processing Example | 17-121 |
| 17.6.13 | Port Test Modes | 17-122 |
| 17.6.14 | Interrupts..... | 17-123 |
| 17.6.14.1 | Transfer/Transaction Based Interrupts..... | 17-124 |
| 17.6.14.1.1 | Transaction Error | 17-124 |
| 17.6.14.1.2 | Serial Bus Babble | 17-124 |
| 17.6.14.1.3 | Data Buffer Error | 17-125 |
| 17.6.14.1.4 | USB Interrupt (Interrupt on Completion (IOC)) | 17-126 |
| 17.6.14.1.5 | Short Packet..... | 17-126 |
| 17.6.14.2 | Host Controller Event Interrupts | 17-126 |
| 17.6.14.2.1 | Port Change Events | 17-126 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 17.6.14.2.2 | Frame List Rollover | 17-126 |
| 17.6.14.2.3 | Interrupt on Async Advance | 17-126 |
| 17.6.14.2.4 | Host System Error | 17-127 |
| 17.7 | Device Data Structures | 17-127 |
| 17.7.1 | Endpoint Queue Head | 17-128 |
| 17.7.1.1 | Endpoint Capabilities/Characteristics | 17-129 |
| 17.7.1.2 | Transfer Overlay | 17-130 |
| 17.7.1.3 | Current dTD Pointer | 17-130 |
| 17.7.1.4 | Setup Buffer | 17-130 |
| 17.7.2 | Endpoint Transfer Descriptor (dTD) | 17-131 |
| 17.8 | Device Operational Model | 17-133 |
| 17.8.1 | Device Controller Initialization | 17-133 |
| 17.8.2 | Port State and Control | 17-135 |
| 17.8.2.1 | Bus Reset | 17-137 |
| 17.8.2.2 | Suspend/Resume | 17-138 |
| 17.8.2.2.1 | Suspend Description | 17-138 |
| 17.8.2.2.2 | Suspend Operational Model | 17-138 |
| 17.8.2.2.3 | Resume | 17-138 |
| 17.8.3 | Managing Endpoints | 17-138 |
| 17.8.3.1 | Endpoint Initialization | 17-139 |
| 17.8.3.1.1 | Stalling | 17-140 |
| 17.8.3.2 | Data Toggle | 17-140 |
| 17.8.3.2.1 | Data Toggle Reset | 17-140 |
| 17.8.3.2.2 | Data Toggle Inhibit | 17-141 |
| 17.8.3.3 | Device For Packet Transfers | 17-141 |
| 17.8.3.3.1 | Priming Transmit Endpoints | 17-141 |
| 17.8.3.3.2 | Priming Receive Endpoints | 17-142 |
| 17.8.3.4 | Interrupt/Bulk Endpoint Operational Model | 17-142 |
| 17.8.3.4.1 | Interrupt/Bulk Endpoint Bus Response Matrix | 17-144 |
| 17.8.3.5 | Control Endpoint Operation Model | 17-144 |
| 17.8.3.5.1 | Setup Phase | 17-144 |
| 17.8.3.5.2 | Data Phase | 17-145 |
| 17.8.3.5.3 | Status Phase | 17-145 |
| 17.8.3.5.4 | Control Endpoint Bus Response Matrix | 17-146 |
| 17.8.3.6 | Isochronous Endpoint Operational Model | 17-147 |
| 17.8.3.6.1 | Isochronous Pipe Synchronization | 17-148 |
| 17.8.3.6.2 | Isochronous Endpoint Bus Response Matrix | 17-149 |
| 17.8.4 | Managing Queue Heads | 17-149 |
| 17.8.4.1 | Queue Head Initialization | 17-150 |
| 17.8.4.2 | Operational Model for Setup Transfers | 17-150 |
| 17.8.5 | Managing Transfers with Transfer Descriptors | 17-151 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 17.8.5.1 | Software Link Pointers | 17-151 |
| 17.8.5.2 | Building a Transfer Descriptor | 17-151 |
| 17.8.5.3 | Executing a Transfer Descriptor | 17-152 |
| 17.8.5.4 | Transfer Completion | 17-152 |
| 17.8.5.5 | Flushing/Depriming an Endpoint | 17-153 |
| 17.8.5.6 | Device Error Matrix | 17-154 |
| 17.8.6 | Servicing Interrupts | 17-154 |
| 17.8.6.1 | High-Frequency Interrupts | 17-154 |
| 17.8.6.2 | Low-Frequency Interrupts | 17-155 |
| 17.8.6.3 | Error Interrupts | 17-155 |
| 17.9 | Deviations from the EHCI Specifications | 17-155 |
| 17.9.1 | Embedded Transaction Translator Function | 17-156 |
| 17.9.1.1 | Capability Registers | 17-156 |
| 17.9.1.2 | Operational Registers | 17-156 |
| 17.9.1.3 | Discovery | 17-156 |
| 17.9.1.4 | Data Structures | 17-157 |
| 17.9.1.5 | Operational Model | 17-157 |
| 17.9.1.5.1 | Microframe Pipeline | 17-158 |
| 17.9.1.5.2 | Split State Machines | 17-158 |
| 17.9.1.5.3 | Asynchronous Transaction Scheduling and Buffer Management | 17-159 |
| 17.9.1.5.4 | Periodic Transaction Scheduling and Buffer Management | 17-159 |
| 17.9.1.5.5 | Multiple Transaction Translators | 17-159 |
| 17.9.2 | Device Operation | 17-159 |
| 17.9.3 | Non-Zero Fields the Register File | 17-160 |
| 17.9.4 | SOF Interrupt | 17-160 |
| 17.9.5 | Embedded Design | 17-160 |
| 17.9.5.1 | Frame Adjust Register | 17-160 |
| 17.9.6 | Miscellaneous Variations from EHCI | 17-160 |
| 17.9.6.1 | Discovery | 17-160 |
| 17.9.6.1.1 | Port Reset | 17-160 |
| 17.9.6.1.2 | Port Speed Detection | 17-161 |
| 17.10 | Timing Diagrams | 17-162 |

Chapter 18 Security Engine (SEC) 3.3

| | | |
|--------|--------------------------------------|------|
| 18.1 | SEC 3.3 Architecture Overview | 18-2 |
| 18.1.1 | Descriptor Overview | 18-4 |
| 18.1.2 | Polychannel Overview | 18-5 |
| 18.1.3 | Controller Overview | 18-6 |
| 18.1.4 | Execution Units (EUs) Overview | 18-6 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 18.1.4.1 | Public Key Execution Unit (PKEU) | 18-7 |
| 18.1.4.1.1 | Elliptic Curve Operations | 18-7 |
| 18.1.4.1.2 | Modular Exponentiation Operations | 18-7 |
| 18.1.4.2 | Data Encryption Standard Execution Unit (DEU)..... | 18-8 |
| 18.1.4.3 | Advanced Encryption Standard Execution Unit (AESU)..... | 18-8 |
| 18.1.4.4 | Message Digest Execution Unit (MDEU) | 18-9 |
| 18.1.4.5 | Cyclical Redundancy Check Unit (CRCU) | 18-9 |
| 18.1.4.6 | Random Number Generator (RNGU)..... | 18-10 |
| 18.2 | Configuration of Internal Memory Space | 18-10 |
| 18.3 | Descriptors | 18-16 |
| 18.3.1 | Descriptor Structure | 18-16 |
| 18.3.2 | Descriptor Format: Header Dword | 18-17 |
| 18.3.2.1 | Selecting Execution Units—EU_SEL0 and EU_SEL1 | 18-19 |
| 18.3.2.2 | Selecting Descriptor Type—DESC_TYPE | 18-20 |
| 18.3.3 | Descriptor Format: Pointer Dwords..... | 18-22 |
| 18.3.4 | Link Table Format | 18-23 |
| 18.3.4.1 | Example of Link Table Operation | 18-26 |
| 18.4 | Descriptor Types | 18-27 |
| 18.5 | Polychannel..... | 18-29 |
| 18.5.1 | Channel Operation | 18-29 |
| 18.5.1.1 | Channel Descriptor Processing..... | 18-29 |
| 18.5.1.2 | Channel Arbitration | 18-30 |
| 18.5.1.3 | Channel Host Notification | 18-31 |
| 18.5.2 | Channel Interrupts..... | 18-31 |
| 18.5.2.1 | Channel Done Interrupt | 18-31 |
| 18.5.2.2 | Channel Error Interrupt..... | 18-32 |
| 18.5.3 | Polychannel Registers..... | 18-32 |
| 18.5.3.1 | Traffic Counters | 18-32 |
| 18.5.3.1.1 | Fetch FIFO Enqueue Counter..... | 18-32 |
| 18.5.3.1.2 | Descriptor Finished Counter..... | 18-33 |
| 18.5.3.1.3 | Data Bytes In Counter | 18-33 |
| 18.5.3.1.4 | Data Bytes Out Counter..... | 18-33 |
| 18.5.4 | Channel Registers | 18-33 |
| 18.5.4.1 | Channel Configuration Register (CCR)..... | 18-34 |
| 18.5.4.2 | Channel Status Register (CSR)..... | 18-36 |
| 18.5.4.3 | Current Descriptor Pointer Register (CDPR) | 18-38 |
| 18.5.4.4 | Fetch FIFO Enqueue Register (FFER) | 18-39 |
| 18.5.5 | Channel Buffers and Tables | 18-40 |
| 18.5.5.1 | Descriptor Buffer (DB)..... | 18-40 |
| 18.5.5.2 | Scatter and Gather Link Tables (SLT, GLT) | 18-40 |
| 18.6 | Controller | 18-41 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 18.6.1 | Bus Transfers | 18-41 |
| 18.6.1.1 | Host-Controlled Access | 18-41 |
| 18.6.1.2 | Channel-Controlled Access | 18-42 |
| 18.6.1.2.1 | Channel Controlled Read—Detailed Description | 18-42 |
| 18.6.1.2.2 | System Bus Master Write—Detailed Description | 18-42 |
| 18.6.2 | Arbitration Algorithms | 18-43 |
| 18.6.2.1 | Round-Robin Arbitration | 18-43 |
| 18.6.2.2 | Weighted Priority Arbitration | 18-43 |
| 18.6.3 | Controller Interrupts | 18-43 |
| 18.6.3.1 | Controller Interrupt Conditions and Interrupt Generation | 18-43 |
| 18.6.3.2 | Blocking of Interrupts | 18-44 |
| 18.6.3.3 | Interrupt Handling | 18-44 |
| 18.6.4 | Controller Registers | 18-44 |
| 18.6.4.1 | EU Assignment Status Register (EUASR) | 18-44 |
| 18.6.4.2 | Interrupt Enable Register (IER) | 18-45 |
| 18.6.4.3 | Interrupt Status Register (ISR) | 18-47 |
| 18.6.4.4 | Interrupt Clear Register (ICR) | 18-48 |
| 18.6.4.5 | ID Register | 18-49 |
| 18.6.4.6 | IP Block Revision Register | 18-50 |
| 18.6.4.7 | Master Control Register (MCR) | 18-51 |
| 18.7 | Power Saving Mode | 18-52 |
| 18.8 | Execution Units | 18-52 |
| 18.8.1 | Advanced Encryption Standard Execution Unit (AESU) | 18-53 |
| 18.8.1.1 | ICV Checking in AESU | 18-53 |
| 18.8.1.2 | AESU Mode Register | 18-54 |
| 18.8.1.3 | AESU Key Size Register | 18-57 |
| 18.8.1.4 | AESU Data Size Register | 18-57 |
| 18.8.1.5 | AESU Reset Control Register | 18-58 |
| 18.8.1.6 | AESU Status Register | 18-58 |
| 18.8.1.7 | AESU Interrupt Status Register | 18-59 |
| 18.8.1.8 | AESU Interrupt Mask Register | 18-61 |
| 18.8.1.9 | ICV Size Register | 18-63 |
| 18.8.1.10 | AESU End of Message Register | 18-63 |
| 18.8.1.11 | AESU Context Registers | 18-64 |
| 18.8.1.11.1 | Context for CBC, CBC-RBP, OFB, and CFB128 Cipher Modes | 18-66 |
| 18.8.1.11.2 | Context for Counter Cipher Mode | 18-66 |
| 18.8.1.11.3 | Context for SRT Cipher Mode | 18-66 |
| 18.8.1.11.4 | Context for CCM Cipher Mode | 18-67 |
| 18.8.1.11.5 | Context and Operation for GCM Cipher Mode | 18-69 |
| 18.8.1.11.6 | Context and Operation for XCBC-MAC Cipher Mode | 18-76 |
| 18.8.1.11.7 | Context and Operation for CMAC Cipher Mode | 18-78 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 18.8.1.11.8 | AESU Key Registers | 18-79 |
| 18.8.1.11.9 | AESU FIFOs..... | 18-79 |
| 18.8.2 | Cyclical Redundancy Check Unit (CRCU) | 18-79 |
| 18.8.2.1 | ICV Checking | 18-80 |
| 18.8.2.2 | CRCU Mode Register..... | 18-80 |
| 18.8.2.3 | CRCU Key Size Register..... | 18-81 |
| 18.8.2.4 | CRCU Data Size Register..... | 18-81 |
| 18.8.2.5 | CRCU Reset Control Register | 18-82 |
| 18.8.2.6 | CRCU Control Register..... | 18-82 |
| 18.8.2.7 | CRCU Status Register | 18-83 |
| 18.8.2.8 | CRCU Interrupt Status Register | 18-84 |
| 18.8.2.9 | CRCU Interrupt Mask Register | 18-86 |
| 18.8.2.10 | CRCU ICV Size Register | 18-87 |
| 18.8.2.11 | CRCU End of Message Register | 18-87 |
| 18.8.2.12 | CRCU Received ICV Register | 18-87 |
| 18.8.2.13 | CRCU Context Register | 18-87 |
| 18.8.2.14 | CRCU Key Register | 18-88 |
| 18.8.2.15 | CRCU FIFO..... | 18-89 |
| 18.8.3 | Data Encryption Standard Execution Unit (DEU)..... | 18-89 |
| 18.8.3.1 | DEU Mode Register | 18-89 |
| 18.8.3.2 | DEU Key Size Register | 18-90 |
| 18.8.3.3 | DEU Data Size Register | 18-90 |
| 18.8.3.4 | DEU Reset Control Register..... | 18-91 |
| 18.8.3.5 | DEU Status Register | 18-92 |
| 18.8.3.6 | DEU Interrupt Status Register | 18-92 |
| 18.8.3.7 | DEU Interrupt Mask Register..... | 18-94 |
| 18.8.3.8 | DEU End_of_message Register | 18-96 |
| 18.8.3.9 | DEU IV Register..... | 18-96 |
| 18.8.3.10 | DEU Key Registers..... | 18-96 |
| 18.8.3.11 | DEU FIFOs..... | 18-97 |
| 18.8.4 | Message Digest Execution Unit (MDEU) | 18-97 |
| 18.8.4.1 | ICV Checking | 18-97 |
| 18.8.4.2 | MDEU Mode Register..... | 18-98 |
| 18.8.4.3 | Recommended Settings for MDEU Mode Register | 18-101 |
| 18.8.4.4 | MDEU Key Size Register..... | 18-102 |
| 18.8.4.5 | MDEU Data Size Register..... | 18-102 |
| 18.8.4.6 | MDEU Reset Control Register | 18-102 |
| 18.8.4.7 | MDEU Status Register | 18-103 |
| 18.8.4.8 | MDEU Interrupt Status Register..... | 18-104 |
| 18.8.4.9 | MDEU Interrupt Mask Register | 18-106 |
| 18.8.4.10 | MDEU ICV Size Register | 18-107 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 18.8.4.11 | MDEU End_of_message Register | 18-107 |
| 18.8.4.12 | MDEU Context Registers | 18-108 |
| 18.8.4.13 | MDEU Key Registers | 18-111 |
| 18.8.4.14 | MDEU FIFOs | 18-111 |
| 18.8.5 | Public Key Execution Units (PKEU)..... | 18-111 |
| 18.8.5.1 | PKEU Mode Register | 18-111 |
| 18.8.5.2 | PKEU Key Size Register | 18-112 |
| 18.8.5.3 | PKEU AB Size Register | 18-113 |
| 18.8.5.4 | PKEU Data Size Register | 18-114 |
| 18.8.5.5 | PKEU Reset Control Register | 18-114 |
| 18.8.5.6 | PKEU Status Register | 18-115 |
| 18.8.5.7 | PKEU Interrupt Status Register..... | 18-116 |
| 18.8.5.8 | PKEU Interrupt Mask Register..... | 18-117 |
| 18.8.5.9 | PKEU End_Of_Message Register | 18-118 |
| 18.8.5.10 | PKEU Parameter Memories | 18-119 |
| 18.8.5.10.1 | PKEU Parameter Memory A | 18-119 |
| 18.8.5.10.2 | PKEU Parameter Memory B | 18-119 |
| 18.8.5.10.3 | PKEU Parameter Memory E | 18-119 |
| 18.8.5.10.4 | PKEU Parameter Memory N..... | 18-119 |
| 18.8.6 | Random Number Generator (RNGU)..... | 18-119 |
| 18.8.6.1 | RNGU Mode Register | 18-120 |
| 18.8.6.2 | RNGU Data Size Register | 18-120 |
| 18.8.6.3 | RNGU Reset Control Register..... | 18-121 |
| 18.8.6.4 | RNGU Status Register | 18-121 |
| 18.8.6.5 | RNGU Interrupt Status Register..... | 18-122 |
| 18.8.6.6 | RNGU Interrupt Mask Register..... | 18-123 |
| 18.8.6.7 | RNGU End_Of_Message Register..... | 18-124 |
| 18.8.6.8 | RNGU ENTROPY | 18-125 |
| 18.8.6.9 | RNGU FIFO | 18-125 |

Chapter 19 Enhanced Three-Speed Ethernet Controllers

| | | |
|--------|--------------------------------------|-------|
| 19.1 | Overview..... | 19-1 |
| 19.2 | Features | 19-2 |
| 19.3 | Modes of Operation | 19-4 |
| 19.4 | External Signals Description | 19-6 |
| 19.4.1 | Detailed Signal Descriptions | 19-8 |
| 19.5 | Memory Map/Register Definition | 19-11 |
| 19.5.1 | Top-Level Module Memory Map | 19-11 |
| 19.5.2 | Detailed Memory Map..... | 19-12 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|---|-------------|
| 19.5.3 | Memory-Mapped Register Descriptions..... | 19-22 |
| 19.5.3.1 | eTSEC General Control and Status Registers..... | 19-22 |
| 19.5.3.1.1 | Controller ID Register (TSEC_ID)..... | 19-22 |
| 19.5.3.1.2 | Controller ID Register (TSEC_ID2)..... | 19-23 |
| 19.5.3.1.3 | Interrupt Event Register (IEVENT) | 19-24 |
| 19.5.3.1.4 | Interrupt Mask Register (IMASK) | 19-28 |
| 19.5.3.1.5 | Error Disabled Register (EDIS)..... | 19-30 |
| 19.5.3.1.6 | Ethernet Control Register (ECNTRL)..... | 19-31 |
| 19.5.3.1.7 | Pause Time Value Register (PTV) | 19-33 |
| 19.5.3.1.8 | DMA Control Register (DMACTRL) | 19-34 |
| 19.5.3.1.9 | TBI Physical Address Register (TBIPA) | 19-35 |
| 19.5.3.2 | eTSEC Transmit Control and Status Registers..... | 19-36 |
| 19.5.3.2.1 | Transmit Control Register (TCTRL) | 19-36 |
| 19.5.3.2.2 | Transmit Status Register (TSTAT)..... | 19-38 |
| 19.5.3.2.3 | Default VLAN Control Word Register (DFVLAN)..... | 19-42 |
| 19.5.3.2.4 | Transmit Interrupt Coalescing Register (TXIC)..... | 19-43 |
| 19.5.3.2.5 | Transmit Queue Control Register (TQUEUE) | 19-44 |
| 19.5.3.2.6 | TxBD Ring 0–3 Weighting Register (TR03WT)..... | 19-44 |
| 19.5.3.2.7 | TxBD Ring 4–7 Weighting Register (TR47WT)..... | 19-45 |
| 19.5.3.2.8 | Transmit Buffer Descriptor Pointers 0–7 (TBPTR0–TBPTR7) | 19-46 |
| 19.5.3.2.9 | Transmit Descriptor Base Address Registers (TBASE0–TBASE7) | 19-47 |
| 19.5.3.2.10 | Transmit Time Stamp Identification Register (TMR_TXTS1–2_ID)..... | 19-47 |
| 19.5.3.2.11 | Transmit Time Stamp Register (TMR_TXTS1–2_H/L) | 19-48 |
| 19.5.3.3 | eTSEC Receive Control and Status Registers | 19-48 |
| 19.5.3.3.1 | Receive Control Register (RCTRL) | 19-48 |
| 19.5.3.3.2 | Receive Status Register (RSTAT)..... | 19-50 |
| 19.5.3.3.3 | Receive Interrupt Coalescing Register (RXIC) | 19-52 |
| 19.5.3.3.4 | Receive Queue Control Register (RQUEUE) | 19-53 |
| 19.5.3.3.5 | Receive Bit Field Extract Control Register (RBIFX)..... | 19-54 |
| 19.5.3.3.6 | Receive Queue Filer Table Address Register (RQFAR) | 19-56 |
| 19.5.3.3.7 | Receive Queue Filer Table Control Register (RQFCR) | 19-56 |
| 19.5.3.3.8 | Receive Queue Filer Table Property Register (RQFPR) | 19-57 |
| 19.5.3.3.9 | Maximum Receive Buffer Length Register (MRBLR)..... | 19-61 |
| 19.5.3.3.10 | Receive Buffer Descriptor Pointers 0–7 (RBPTR0–RBPTR7) | 19-61 |
| 19.5.3.3.11 | Receive Descriptor Base Address Registers (RBASE0–RBASE7) | 19-62 |
| 19.5.3.3.12 | Receive Stamp Register (TMR_RXTS_H/L)..... | 19-62 |
| 19.5.3.4 | MAC Functionality..... | 19-63 |
| 19.5.3.4.1 | Configuring the MAC | 19-63 |
| 19.5.3.4.2 | Controlling CSMA/CD..... | 19-63 |
| 19.5.3.4.3 | Handling Packet Collisions | 19-64 |
| 19.5.3.4.4 | Controlling Packet Flow | 19-64 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 19.5.3.4.5 | Controlling PHY Links..... | 19-65 |
| 19.5.3.5 | MAC Registers | 19-66 |
| 19.5.3.5.1 | MAC Configuration 1 Register (MACCFG1)..... | 19-66 |
| 19.5.3.5.2 | MAC Configuration 2 Register (MACCFG2)..... | 19-67 |
| 19.5.3.5.3 | Inter-Packet Gap/Inter-Frame Gap Register (IPGIFG) | 19-69 |
| 19.5.3.5.4 | Half-Duplex Register (HAFDUP) | 19-70 |
| 19.5.3.5.5 | Maximum Frame Length Register (MAXFRM) | 19-71 |
| 19.5.3.5.6 | MII Management Configuration Register (MIIMCFG) | 19-71 |
| 19.5.3.5.7 | MII Management Command Register (MIIMCOM)..... | 19-72 |
| 19.5.3.5.8 | MII Management Address Register (MIIMADD)..... | 19-73 |
| 19.5.3.5.9 | MII Management Control Register (MIIMCON)..... | 19-74 |
| 19.5.3.5.10 | MII Management Status Register (MIIMSTAT) | 19-74 |
| 19.5.3.5.11 | MII Management Indicator Register (MIIMIND)..... | 19-75 |
| 19.5.3.5.12 | Interface Status Register (IFSTAT)..... | 19-75 |
| 19.5.3.5.13 | MAC Station Address Part 1 Register (MACSTNADDR1) | 19-76 |
| 19.5.3.5.14 | MAC Station Address Part 2 Register (MACSTNADDR2) | 19-77 |
| 19.5.3.5.15 | MAC Exact Match Address 1–15 Part 1 Registers (MAC01ADDR1–MAC15ADDR1)..... | 19-77 |
| 19.5.3.5.16 | MAC Exact Match Address 1–15 Part 2 Registers (MAC01ADDR2–MAC15ADDR2)..... | 19-78 |
| 19.5.3.6 | MIB Registers..... | 19-78 |
| 19.5.3.6.1 | Transmit and Receive 64-Byte Frame Counter (TR64) | 19-79 |
| 19.5.3.6.2 | Transmit and Receive 65- to 127-Byte Frame Counter (TR127) | 19-80 |
| 19.5.3.6.3 | Transmit and Receive 128- to 255-Byte Frame Counter (TR255) | 19-80 |
| 19.5.3.6.4 | Transmit and Receive 256- to 511-Byte Frame Counter (TR511) | 19-81 |
| 19.5.3.6.5 | Transmit and Receive 512- to 1023-Byte Frame Counter (TR1K) | 19-81 |
| 19.5.3.6.6 | Transmit and Receive 1024- to 1518-Byte Frame Counter (TRMAX)..... | 19-82 |
| 19.5.3.6.7 | Transmit and Receive 1519- to 1522-Byte VLAN Frame Counter (TRMGV)..... | 19-82 |
| 19.5.3.6.8 | Receive Byte Counter (RBYT)..... | 19-83 |
| 19.5.3.6.9 | Receive Packet Counter (RPKT)..... | 19-83 |
| 19.5.3.6.10 | Receive FCS Error Counter (RFCS) | 19-83 |
| 19.5.3.6.11 | Receive Multicast Packet Counter (RMCA) | 19-84 |
| 19.5.3.6.12 | Receive Broadcast Packet Counter (RBCA) | 19-84 |
| 19.5.3.6.13 | Receive Control Frame Packet Counter (RXCF) | 19-85 |
| 19.5.3.6.14 | Receive Pause Frame Packet Counter (RXPF)..... | 19-85 |
| 19.5.3.6.15 | Receive Unknown Opcode Packet Counter (RXUO)..... | 19-86 |
| 19.5.3.6.16 | Receive Alignment Error Counter (RALN) | 19-86 |
| 19.5.3.6.17 | Receive Frame Length Error Counter (RFLR)..... | 19-87 |
| 19.5.3.6.18 | Receive Code Error Counter (RCDE) | 19-87 |
| 19.5.3.6.19 | Receive Carrier Sense Error Counter (RCSE)..... | 19-88 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 19.5.3.6.20 | Receive Undersize Packet Counter (RUND)..... | 19-88 |
| 19.5.3.6.21 | Receive Oversize Packet Counter (ROVR)..... | 19-89 |
| 19.5.3.6.22 | Receive Fragments Counter (RFRG) | 19-89 |
| 19.5.3.6.23 | Receive Jabber Counter (RJBR)..... | 19-90 |
| 19.5.3.6.24 | Receive Dropped Packet Counter (RDRP)..... | 19-90 |
| 19.5.3.6.25 | Transmit Byte Counter (TBYT) | 19-91 |
| 19.5.3.6.26 | Transmit Packet Counter (TPKT)..... | 19-91 |
| 19.5.3.6.27 | Transmit Multicast Packet Counter (TMCA)..... | 19-92 |
| 19.5.3.6.28 | Transmit Broadcast Packet Counter (TBCA)..... | 19-92 |
| 19.5.3.6.29 | Transmit Pause Control Frame Counter (TXPF)..... | 19-93 |
| 19.5.3.6.30 | Transmit Deferral Packet Counter (TDFR) | 19-93 |
| 19.5.3.6.31 | Transmit Excessive Deferral Packet Counter (TEDF) | 19-94 |
| 19.5.3.6.32 | Transmit Single Collision Packet Counter (TSCL) | 19-94 |
| 19.5.3.6.33 | Transmit Multiple Collision Packet Counter (TMCL) | 19-95 |
| 19.5.3.6.34 | Transmit Late Collision Packet Counter (TLCL)..... | 19-95 |
| 19.5.3.6.35 | Transmit Excessive Collision Packet Counter (TXCL)..... | 19-96 |
| 19.5.3.6.36 | Transmit Total Collision Counter (TNCL) | 19-96 |
| 19.5.3.6.37 | Transmit Drop Frame Counter (TDRP)..... | 19-97 |
| 19.5.3.6.38 | Transmit Jabber Frame Counter (TJBR) | 19-97 |
| 19.5.3.6.39 | Transmit FCS Error Counter (TFCS) | 19-98 |
| 19.5.3.6.40 | Transmit Control Frame Counter (TXCF)..... | 19-98 |
| 19.5.3.6.41 | Transmit Oversize Frame Counter (TOVR)..... | 19-99 |
| 19.5.3.6.42 | Transmit Undersize Frame Counter (TUND)..... | 19-99 |
| 19.5.3.6.43 | Transmit Fragment Counter (TFRG)..... | 19-100 |
| 19.5.3.6.44 | Carry Register 1 (CAR1)..... | 19-100 |
| 19.5.3.6.45 | Carry Register 2 (CAR2)..... | 19-102 |
| 19.5.3.6.46 | Carry Mask Register 1 (CAM1)..... | 19-103 |
| 19.5.3.6.47 | Carry Mask Register 2 (CAM2)..... | 19-104 |
| 19.5.3.6.48 | Receive Filer Rejected Packet Counter (RREJ) | 19-105 |
| 19.5.3.7 | Hash Function Registers..... | 19-106 |
| 19.5.3.7.1 | Individual/Group Address Registers 0–7 (IGADDR n) | 19-106 |
| 19.5.3.7.2 | Group Address Registers 0–7 (GADDR n) | 19-106 |
| 19.5.3.8 | DMA Attribute Registers..... | 19-107 |
| 19.5.3.8.1 | Attribute Register (ATTR)..... | 19-107 |
| 19.5.3.9 | Lossless Flow Control Configuration Registers..... | 19-108 |
| 19.5.3.9.1 | Receive Queue Parameters 0–7 (RQPRM0–PQPRM7)..... | 19-108 |
| 19.5.3.9.2 | Receive Free Buffer Descriptor Pointer Registers 0–7 (RFBPTR0–RFBPTR7) | 19-109 |
| 19.5.3.10 | IEEE 1588-Compatible Timestamping Registers..... | 19-109 |
| 19.5.3.10.1 | Timer Control Register (TMR_CTRL) | 19-110 |
| 19.5.3.10.2 | Timer Event Register (TMR_TEVENT) | 19-111 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 19.5.3.10.3 | Timer Event Mask Register (TMR_TEMASK) | 19-113 |
| 19.5.3.10.4 | Timer PTP Packet Event Register (TMR_PEVENT) | 19-113 |
| 19.5.3.10.5 | Timer Event Mask Register (TMR_PEMASK) | 19-114 |
| 19.5.3.10.6 | Timer Status Register (TMR_STAT) | 19-115 |
| 19.5.3.10.7 | Timer Counter Register (TMR_CNT_H/L) | 19-115 |
| 19.5.3.10.8 | Timer Drift Compensation Addend Register (TMR_ADD) | 19-116 |
| 19.5.3.10.9 | Timer Accumulator Register (TMR_ACC) | 19-117 |
| 19.5.3.10.10 | Timer Prescale Register (TMR_PRSC) | 19-117 |
| 19.5.3.10.11 | Timer Offset Register (TMROFF_H/L) | 19-118 |
| 19.5.3.10.12 | Alarm Time Comparator Register (TMR_ALARM1-2_H/L) | 19-118 |
| 19.5.3.10.13 | Timer Fixed Interval Period Register (TMR_FIPER1-3) | 19-119 |
| 19.5.3.10.14 | External Trigger Stamp Register (TMR_ETTS1-2_H/L) | 19-120 |
| 19.5.4 | Ten-Bit Interface (TBI) | 19-120 |
| 19.5.4.1 | TBI Transmit Process | 19-121 |
| 19.5.4.1.1 | Packet Encapsulation | 19-121 |
| 19.5.4.1.2 | 8B10B Encoding | 19-121 |
| 19.5.4.1.3 | Preamble Shortening | 19-121 |
| 19.5.4.2 | TBI Receive Process | 19-121 |
| 19.5.4.2.1 | Synchronization | 19-121 |
| 19.5.4.2.2 | Auto-Negotiation for 1000BASE-X | 19-122 |
| 19.5.4.3 | TBI MII Set Register Descriptions | 19-122 |
| 19.5.4.3.1 | Control Register (CR) | 19-123 |
| 19.5.4.3.2 | Status Register (SR) | 19-124 |
| 19.5.4.3.3 | AN Advertisement Register (ANA) | 19-125 |
| 19.5.4.3.4 | AN Link Partner Base Page Ability Register (ANLPBPA) | 19-127 |
| 19.5.4.3.5 | AN Expansion Register (ANEX) | 19-128 |
| 19.5.4.3.6 | AN Next Page Transmit Register (ANNPT) | 19-128 |
| 19.5.4.3.7 | AN Link Partner Ability Next Page Register (ANLPANP) | 19-129 |
| 19.5.4.3.8 | Extended Status Register (EXST) | 19-130 |
| 19.5.4.3.9 | Jitter Diagnostics Register (JD) | 19-131 |
| 19.5.4.3.10 | TBI Control Register (TBICON) | 19-132 |
| 19.6 | Functional Description | 19-133 |
| 19.6.1 | Connecting to Physical Interfaces on Ethernet | 19-133 |
| 19.6.1.1 | Media-Independent Interface (MII) | 19-134 |
| 19.6.1.2 | Reduced Media-Independent Interface (RMII) | 19-134 |
| 19.6.1.3 | Reduced Gigabit Media-Independent Interface (RGMII) | 19-135 |
| 19.6.1.4 | Reduced Ten-Bit Interface (RTBI) | 19-136 |
| 19.6.1.5 | Serial Gigabit Media-Independent Interface (SGMII) | 19-137 |
| 19.6.1.6 | Ethernet Physical Interfaces Signal Summary | 19-138 |
| 19.6.2 | Gigabit Ethernet Controller Channel Operation | 19-140 |
| 19.6.2.1 | Initialization Sequence | 19-141 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 19.6.2.1.1 | Hardware Controlled Initialization | 19-141 |
| 19.6.2.1.2 | User Initialization | 19-141 |
| 19.6.2.2 | Soft Reset and Reconfiguring Procedure | 19-142 |
| 19.6.2.3 | Gigabit Ethernet Frame Transmission | 19-143 |
| 19.6.2.4 | Gigabit Ethernet Frame Reception | 19-144 |
| 19.6.2.5 | Ethernet Preamble Customization | 19-145 |
| 19.6.2.5.1 | User-Defined Preamble Transmission | 19-145 |
| 19.6.2.5.2 | User-Visible Preamble Reception | 19-146 |
| 19.6.2.6 | RMON Support | 19-147 |
| 19.6.2.7 | Frame Recognition | 19-147 |
| 19.6.2.7.1 | Destination Address Recognition and Frame Filtering | 19-148 |
| 19.6.2.7.2 | Hash Table Algorithm | 19-149 |
| 19.6.2.8 | Magic Packet Mode | 19-151 |
| 19.6.2.9 | Flow Control | 19-151 |
| 19.6.2.10 | Interrupt Handling | 19-152 |
| 19.6.2.10.1 | Interrupt Coalescing | 19-153 |
| 19.6.2.10.2 | Interrupt Coalescing By Frame Count Threshold | 19-153 |
| 19.6.2.10.3 | Interrupt Coalescing By Timer Threshold | 19-154 |
| 19.6.2.11 | Inter-Frame Gap Time | 19-155 |
| 19.6.2.12 | Internal and External Loop Back | 19-155 |
| 19.6.2.13 | Error-Handling Procedure | 19-155 |
| 19.6.3 | TCP/IP Off-Load | 19-157 |
| 19.6.3.1 | Frame Control Blocks | 19-158 |
| 19.6.3.2 | Transmit Path Off-Load and Tx PTP Packet Parsing | 19-159 |
| 19.6.3.3 | Receive Path Off-Load | 19-160 |
| 19.6.4 | Quality of Service (QoS) Provision | 19-162 |
| 19.6.4.1 | Receive Parser | 19-162 |
| 19.6.4.2 | Receive Queue Filer | 19-164 |
| 19.6.4.2.1 | Filing Rules | 19-165 |
| 19.6.4.2.2 | Comparing Properties with Bit Masks | 19-166 |
| 19.6.4.2.3 | Special-Case Rules | 19-167 |
| 19.6.4.2.4 | Filer Interrupt Events | 19-167 |
| 19.6.4.2.5 | Setting Up the Receive Queue Filer Table | 19-168 |
| 19.6.4.2.6 | Filer Example—802.1p Priority Filing | 19-168 |
| 19.6.4.2.7 | Filer Example—IP Diff-Serv Code Points Filing | 19-169 |
| 19.6.4.2.8 | Filer Example—TCP and UDP Port Filing | 19-169 |
| 19.6.4.2.9 | Filer Example—Interrupt from Deep Sleep | 19-170 |
| 19.6.4.3 | Transmission Scheduling | 19-172 |
| 19.6.4.3.1 | Priority-Based Queuing (PBQ) | 19-172 |
| 19.6.4.3.2 | Modified Weighted Round-Robin Queuing (MWRR) | 19-173 |
| 19.6.5 | Lossless Flow Control | 19-174 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 19.6.5.1 | Back Pressure Determination through Free Buffers | 19-174 |
| 19.6.5.2 | Software Use of Hardware-Initiated Back Pressure | 19-176 |
| 19.6.5.2.1 | Initialization | 19-176 |
| 19.6.5.2.2 | Operation | 19-176 |
| 19.6.6 | Hardware Assist for IEEE 1588-Compatible Timestamping..... | 19-176 |
| 19.6.6.1 | Features..... | 19-177 |
| 19.6.6.2 | Timer Logic Overview..... | 19-178 |
| 19.6.6.3 | Time-Stamp Insertion on the Received Packets | 19-178 |
| 19.6.6.3.1 | Timestamp Point | 19-178 |
| 19.6.6.4 | PTP Packet Parsing | 19-179 |
| 19.6.6.4.1 | General Purpose Filter Rule..... | 19-180 |
| 19.6.6.5 | Time-Stamp Insertion on Transmit Packets..... | 19-180 |
| 19.6.6.5.1 | Interrupts..... | 19-180 |
| 19.6.6.5.2 | Error Condition..... | 19-181 |
| 19.6.6.6 | Tx PTP Packet Parsing..... | 19-182 |
| 19.6.7 | Buffer Descriptors..... | 19-183 |
| 19.6.7.1 | Data Buffer Descriptors | 19-183 |
| 19.6.7.2 | Transmit Data Buffer Descriptors (TxBD)..... | 19-185 |
| 19.6.7.3 | Receive Buffer Descriptors (RxBd)..... | 19-188 |
| 19.7 | Initialization/Application Information | 19-190 |
| 19.7.1 | Interface Mode Configuration | 19-190 |
| 19.7.1.1 | MII Interface Mode..... | 19-191 |
| 19.7.1.2 | RGMII Interface Mode | 19-194 |
| 19.7.1.3 | RMII Interface Mode | 19-198 |
| 19.7.1.4 | RTBI Interface Mode..... | 19-202 |
| 19.7.1.5 | SGMII Interface Support | 19-205 |

Chapter 20 I²C Interface

| | | |
|----------|---|------|
| 20.1 | I ² C Introduction | 20-1 |
| 20.1.1 | I ² C Features | 20-2 |
| 20.1.2 | I ² C Modes of Operation..... | 20-2 |
| 20.2 | I ² C External Signal Descriptions | 20-3 |
| 20.2.1 | I ² C Signal Overview..... | 20-3 |
| 20.2.2 | I ² C Detailed Signal Descriptions..... | 20-3 |
| 20.3 | I ² C Memory Map/Register Definition..... | 20-4 |
| 20.3.1 | I ² C Register Descriptions | 20-4 |
| 20.3.1.1 | I ² C Address Register (I2CADR) | 20-4 |
| 20.3.1.2 | I ² C Frequency Divider Register (I2CFDR)..... | 20-5 |
| 20.3.1.3 | I ² C Control Register (I2CCR) | 20-6 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 20.3.1.4 | I ² C Status Register (I2CSR) | 20-7 |
| 20.3.1.5 | I ² C Data Register (I2CDR)..... | 20-9 |
| 20.3.1.6 | Digital Filter Sampling Rate Register (I2CDFSRR) | 20-9 |
| 20.4 | Functional Description..... | 20-9 |
| 20.4.1 | Transaction Protocol | 20-10 |
| 20.4.1.1 | START Condition | 20-10 |
| 20.4.1.2 | Slave Address Transmission | 20-10 |
| 20.4.1.3 | Repeated START Condition | 20-11 |
| 20.4.1.4 | STOP Condition..... | 20-11 |
| 20.4.1.5 | Protocol Implementation Details | 20-12 |
| 20.4.1.5.1 | Transaction Monitoring—Implementation Details..... | 20-12 |
| 20.4.1.5.2 | Control Transfer—Implementation Details | 20-12 |
| 20.4.1.6 | Address Compare—Implementation Details | 20-13 |
| 20.4.2 | Arbitration Procedure | 20-13 |
| 20.4.2.1 | Arbitration Control | 20-13 |
| 20.4.3 | Handshaking | 20-14 |
| 20.4.4 | Clock Control..... | 20-14 |
| 20.4.4.1 | Clock Synchronization..... | 20-14 |
| 20.4.4.2 | Input Synchronization and Digital Filter | 20-14 |
| 20.4.4.2.1 | Input Signal Synchronization | 20-15 |
| 20.4.4.2.2 | Filtering of SCL and SDA Lines | 20-15 |
| 20.4.4.3 | Clock Stretching | 20-15 |
| 20.4.5 | Boot Sequencer Mode..... | 20-15 |
| 20.4.5.1 | Using the Boot Sequencer for Reset Configuration | 20-16 |
| 20.4.5.2 | EEPROM Calling Address | 20-16 |
| 20.4.5.3 | EEPROM Data Format | 20-16 |
| 20.4.5.4 | Boot Sequencer Done Indication | 20-19 |
| 20.5 | Initialization/Application Information | 20-19 |
| 20.5.1 | Interrupt Service Routine Flowchart..... | 20-19 |
| 20.5.2 | Initialization Sequence..... | 20-20 |
| 20.5.3 | Generation of START | 20-21 |
| 20.5.4 | Post-Transfer Software Response | 20-21 |
| 20.5.5 | Generation of STOP..... | 20-22 |
| 20.5.6 | Generation of Repeated START | 20-22 |
| 20.5.7 | Generation of SCL When SDA is Negated | 20-22 |
| 20.5.8 | Slave Mode Interrupt Service Routine..... | 20-22 |
| 20.5.8.1 | Slave Transmitter and Received Acknowledge | 20-22 |
| 20.5.8.2 | Loss of Arbitration and Forcing of Slave Mode | 20-23 |

Chapter 21 DUART

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 21.1 | DUART Overview | 21-1 |
| 21.1.1 | DUART Features | 21-2 |
| 21.1.2 | DUART Modes of Operation..... | 21-3 |
| 21.2 | DUART External Signal Descriptions | 21-3 |
| 21.2.1 | DUART Signal Overview | 21-3 |
| 21.2.2 | DUART Detailed Signal Descriptions..... | 21-3 |
| 21.3 | DUART Memory Map/Register Definition | 21-4 |
| 21.3.1 | DUART Register Descriptions | 21-6 |
| 21.3.1.1 | Receiver Buffer Registers (URBR1 and URBR2)..... | 21-6 |
| 21.3.1.2 | Transmitter Holding Registers (UTHR1 and UTHR2)..... | 21-6 |
| 21.3.1.3 | Divisor Most and Least Significant Byte Registers (UDMB and UDLB) | 21-7 |
| 21.3.1.4 | Interrupt Enable Registers (UIER1 and UIER2) | 21-8 |
| 21.3.1.5 | Interrupt ID Registers (UIIR1 and UIIR2) | 21-9 |
| 21.3.1.6 | FIFO Control Registers (UFCR1 and UFCR2) | 21-10 |
| 21.3.1.7 | Alternate Function Registers (UAFR1 and UAFR2)..... | 21-11 |
| 21.3.1.8 | Line Control Registers (ULCR1 and ULCR2) | 21-12 |
| 21.3.1.9 | MODEM Control Registers (UMCR1 and UMCR2)..... | 21-14 |
| 21.3.1.10 | Line Status Registers (ULSR1 and ULSR2) | 21-15 |
| 21.3.1.11 | MODEM Status Registers (UMSR1 and UMSR2) | 21-16 |
| 21.3.1.12 | Scratch Registers (USCR1 and USCR2) | 21-17 |
| 21.3.1.13 | DMA Status Registers (UDSR1 and UDSR2)..... | 21-17 |
| 21.4 | Functional Description..... | 21-18 |
| 21.4.1 | Serial Interface | 21-19 |
| 21.4.1.1 | START Bit | 21-19 |
| 21.4.1.2 | Data Transfer | 21-20 |
| 21.4.1.3 | Parity Bit..... | 21-20 |
| 21.4.1.4 | STOP Bit..... | 21-20 |
| 21.4.2 | Baud-Rate Generator Logic | 21-20 |
| 21.4.3 | Local Loopback Mode | 21-21 |
| 21.4.4 | Errors | 21-21 |
| 21.4.4.1 | Framing Error | 21-21 |
| 21.4.4.2 | Parity Error | 21-21 |
| 21.4.4.3 | Overrun Error..... | 21-21 |
| 21.4.5 | FIFO Mode | 21-21 |
| 21.4.5.1 | FIFO Interrupts | 21-22 |
| 21.4.5.2 | DMA Mode Select..... | 21-22 |
| 21.4.5.3 | Interrupt Control Logic..... | 21-22 |
| 21.5 | DUART Initialization/Application Information | 21-23 |

Chapter 22 Serial Peripheral Interface

Contents

| Paragraph Number | Title | Page Number |
|------------------|--|-------------|
| 22.1 | Overview..... | 22-1 |
| 22.2 | Introduction..... | 22-2 |
| 22.2.1 | Features..... | 22-2 |
| 22.2.2 | SPI Transmission and Reception Process..... | 22-3 |
| 22.2.3 | Modes of Operation..... | 22-3 |
| 22.2.3.1 | SPI as a Master Device..... | 22-3 |
| 22.2.3.2 | SPI as a Slave Device..... | 22-4 |
| 22.2.3.3 | SPI in Multiple-Master Operation..... | 22-5 |
| 22.3 | External Signal Descriptions..... | 22-6 |
| 22.3.1 | Overview..... | 22-7 |
| 22.3.2 | Detailed Signal Descriptions..... | 22-7 |
| 22.4 | Memory Map/Register Definition..... | 22-8 |
| 22.4.1 | Register Descriptions..... | 22-9 |
| 22.4.1.1 | SPI Mode Register (SPMODE)..... | 22-9 |
| 22.4.1.2 | SPI Event Register (SPIE)..... | 22-11 |
| 22.4.1.3 | SPI Mask Register (SPIM)..... | 22-12 |
| 22.4.1.4 | SPI Command Register (SPCOM)..... | 22-14 |
| 22.4.1.5 | SPI Transmit Data Hold Register (SPITD)..... | 22-14 |
| 22.4.1.6 | SPI Receive Data Hold Register (SPIRD)..... | 22-15 |
| 22.4.1.6.1 | Reverse Mode SPMODE[REV] Examples..... | 22-15 |
| 22.5 | Initialization/Application Information..... | 22-16 |
| 22.5.1 | SPI Master Programming Example..... | 22-16 |
| 22.5.2 | SPI Slave Programming Example..... | 22-16 |

Chapter 23 JTAG/Testing Support

| | | |
|--------|--|------|
| 23.1 | JTAG Overview..... | 23-1 |
| 23.2 | JTAG Signals..... | 23-1 |
| 23.2.1 | JTAG External Signal Descriptions..... | 23-2 |
| 23.3 | JTAG Registers and Scan Chains..... | 23-3 |

Chapter 24 General Purpose I/O (GPIO)

| | | |
|--------|-------------------------------------|------|
| 24.1 | Introduction..... | 24-1 |
| 24.1.1 | Overview..... | 24-1 |
| 24.1.2 | Features..... | 24-1 |
| 24.2 | External Signal Description..... | 24-2 |
| 24.2.1 | Signals Overview..... | 24-2 |
| 24.3 | Memory Map/Register Definition..... | 24-2 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|--|----------------|
| 24.3.1 | GPIO Direction Register (GPDIR)..... | 24-3 |
| 24.3.2 | GPIO Open Drain Register (GPODR)..... | 24-3 |
| 24.3.3 | GPIO Data Register (GPDAT)..... | 24-4 |
| 24.3.4 | GPIO Interrupt Event Register (GPIER) | 24-4 |
| 24.3.5 | GPIO Interrupt Mask Register (GPIMR)..... | 24-4 |
| 24.3.6 | GPIO Interrupt Control Register (GPICR)..... | 24-5 |

Chapter 25 Time Division Multiplexing (TDM) Interface

| | | |
|------------|--|-------|
| 25.1 | Introduction..... | 25-1 |
| 25.1.1 | Features..... | 25-1 |
| 25.2 | Signal Descriptions..... | 25-2 |
| 25.2.1 | Overview..... | 25-2 |
| 25.2.2 | External Signals Descriptions..... | 25-2 |
| 25.3 | TDM Overview..... | 25-4 |
| 25.3.1 | TDM Basics..... | 25-5 |
| 25.3.2 | Common Signals for the TDM Modules..... | 25-6 |
| 25.4 | TDM Programming Model..... | 25-7 |
| 25.4.1 | SB Interface..... | 25-7 |
| 25.4.2 | SB Memory Map..... | 25-7 |
| 25.4.2.1 | Configuration Registers..... | 25-8 |
| 25.4.2.1.1 | TDM General Interface Register (TDMGIR)..... | 25-8 |
| 25.4.2.1.2 | TDM Receive Interface Register (TDMRIR)..... | 25-9 |
| 25.4.2.1.3 | TDM Transmit Interface Register (TDMTIR)..... | 25-12 |
| 25.4.2.1.4 | TDM Receive Frame Parameters (TDMRFP)..... | 25-14 |
| 25.4.2.1.5 | TDM Transmit Frame Parameters (TDMTFP)..... | 25-16 |
| 25.4.2.2 | Control Registers..... | 25-17 |
| 25.4.2.2.1 | TDM Receive Channel Enable <i>n</i> (TDMRCEN _{<i>n</i>})..... | 25-17 |
| 25.4.2.2.2 | TDM Transmit Channel Enable <i>n</i> (TDMTCEN _{<i>n</i>})..... | 25-17 |
| 25.4.2.2.3 | TDM Transmit Channel Mask <i>n</i> (TDMTCMA _{<i>n</i>})..... | 25-18 |
| 25.4.2.2.4 | TDM Receive Control Register (TDMRCR)..... | 25-19 |
| 25.4.2.2.5 | TDM Transmit Control Register (TDMTCR)..... | 25-19 |
| 25.4.2.2.6 | TDM Receive Interrupt Enable Register (TDMRIER)..... | 25-20 |
| 25.4.2.2.7 | TDM Transmit Interrupt Enable Register (TDMTIER)..... | 25-21 |
| 25.4.2.3 | Status Registers..... | 25-22 |
| 25.4.2.3.1 | TDM Receive Event Register (TDMRER)..... | 25-22 |
| 25.4.2.3.2 | TDM Transmit Event Register (TDMTER)..... | 25-24 |
| 25.4.2.3.3 | TDM Receive Status Register (TDMRSR)..... | 25-26 |
| 25.4.2.3.4 | TDM Transmit Status Register (TDMTSR)..... | 25-26 |
| 25.4.3 | AHB Interface..... | 25-27 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|---|-------------|
| 25.4.4 | AHB Memory Map..... | 25-28 |
| 25.4.4.1 | Data Registers..... | 25-28 |
| 25.4.4.1.1 | TDM Receive Data Register (TDMRDREG) | 25-28 |
| 25.4.4.1.2 | TDM Transmit Data Register (TDMTDREG) | 25-29 |
| 25.5 | Clocks and Reset..... | 25-29 |
| 25.5.1 | TDM Clock and Frame Sync Generation | 25-30 |
| 25.5.2 | Reset..... | 25-31 |
| 25.6 | TDM Configurations..... | 25-31 |
| 25.6.1 | Typical Configurations | 25-32 |
| 25.7 | TDM Detailed Operation | 25-33 |
| 25.7.1 | Serial Interface..... | 25-33 |
| 25.7.1.1 | Sync Out Configuration..... | 25-33 |
| 25.7.1.2 | Sync In Configuration..... | 25-34 |
| 25.7.1.3 | Serial Interface Synchronization..... | 25-37 |
| 25.7.1.4 | Reverse Data Order..... | 25-38 |
| 25.7.2 | Receiver and Transmitter Independent or Shared Operation..... | 25-39 |
| 25.7.3 | TDM Multichannel (Network) Mode | 25-40 |
| 25.7.3.1 | Operation Using Tx Channel Mask Register..... | 25-41 |
| 25.7.3.2 | Operation Using Rx Channel Enable Register | 25-42 |
| 25.7.4 | Data Structures..... | 25-43 |
| 25.7.5 | FIFO Configuration | 25-44 |
| 25.7.6 | DMA Configuration..... | 25-45 |
| 25.7.6.1 | Setting up the TDM for Correct Operation with the DMA | 25-45 |
| 25.8 | Software Programming..... | 25-46 |
| 25.8.1 | Software Programming Sequence..... | 25-46 |
| 25.8.1.1 | Initialization—Shared Operation Starting on the Same Frame | 25-46 |
| 25.8.1.2 | Initialization—Non-Shared Operation..... | 25-47 |
| 25.8.1.3 | Dynamic Channel Configuration while TDM Operating—Shared | 25-48 |
| 25.8.1.4 | Dynamic Channel Configuration while TDM Operating—Non-Shared..... | 25-48 |
| 25.8.1.5 | Configuring the TDM for I2S Operation..... | 25-49 |
| 25.8.1.6 | TDM Power Down | 25-49 |
| 25.8.1.7 | Synchronization Errors | 25-50 |
| 25.8.2 | Interrupts..... | 25-50 |
| 25.8.2.1 | Receiver Normal and Receiver Error Interrupts..... | 25-50 |
| 25.8.2.2 | Transmit Normal and Transmit Error TDMRDREG | 25-50 |
| 25.9 | TDM Clock Control..... | 25-50 |
| 25.9.1 | Features..... | 25-50 |
| 25.9.2 | Theory of Operation..... | 25-51 |
| 25.9.3 | TDM Clock Memory Map..... | 25-52 |
| 25.9.3.1 | TDMCLK_DIV_VAL_RX Register..... | 25-52 |
| 25.9.3.2 | TDMCLK_DIV_VAL_TX Register..... | 25-53 |

Contents

| Paragraph Number | Title | Page Number |
|---------------------|---|----------------|
| 25.10 | DMA Controller (DMAC) | 25-53 |
| 25.10.1 | Overview | 25-54 |
| 25.10.1.1 | Features | 25-55 |
| 25.10.2 | DMAC Memory Map/Register Definition | 25-55 |
| 25.10.2.1 | DMA Control Register (DMACR) | 25-57 |
| 25.10.3 | DMA Error Status (DMAES) | 25-59 |
| 25.10.3.1 | DMA Enable Request Register (DMAERQ) | 25-62 |
| 25.10.3.2 | DMA Enable Error Interrupt Register (DMAEEI) | 25-63 |
| 25.10.3.3 | DMA Set Enable Request (DMASERQ) | 25-64 |
| 25.10.3.4 | DMA Clear Enable Request (DMACERQ) | 25-64 |
| 25.10.3.5 | DMA Set Enable Error Interrupt (DMASEEI) | 25-65 |
| 25.10.3.6 | DMA Clear Enable Error Interrupt (DMACEEI) | 25-65 |
| 25.10.3.7 | DMA Clear Interrupt Request (DMACINT) | 25-66 |
| 25.10.3.8 | DMA Clear Error (DMACERR) | 25-67 |
| 25.10.3.9 | DMA Set START Bit (DMASSRT) | 25-67 |
| 25.10.3.10 | DMA Clear DONE Status (DMACDNE) | 25-68 |
| 25.10.3.11 | DMA Interrupt Request Register (DMAINT) | 25-68 |
| 25.10.3.12 | DMA Error Register (DMAERR) | 25-69 |
| 25.10.3.13 | DMA Hardware Request Status Register (DMAHRS) | 25-70 |
| 25.10.3.14 | DMA General Purpose Output Register (DMAGPOR) | 25-71 |
| 25.10.3.15 | DMA Channel <i>n</i> Priority (DCHPRI _{<i>n</i>} , <i>n</i> = 0, ..., {15,31,63} | 25-71 |
| 25.10.3.16 | Transfer Control Descriptor (TCD) | 25-72 |
| 25.10.4 | Functional Description | 25-80 |
| 25.10.4.1 | DMA Microarchitecture | 25-80 |
| 25.10.4.2 | DMA Basic Data Flow | 25-81 |
| 25.10.5 | Initialization/Application Information | 25-84 |
| 25.10.5.1 | DMA Initialization | 25-84 |
| 25.10.5.2 | DMA Programming Errors | 25-85 |
| 25.10.6 | DMA Transfer | 25-86 |
| 25.10.6.1 | Single Request | 25-86 |
| 25.10.6.2 | Multiple Requests | 25-87 |
| 25.10.7 | TCD Status | 25-88 |
| 25.10.7.1 | Minor Loop Complete | 25-88 |
| 25.10.7.2 | Active Channel TCD Reads | 25-89 |
| 25.10.8 | Hardware Request Release Timing | 25-89 |

Appendix A

Complete List of Configuration, Control, and Status Registers

| | | |
|-----|--------------------------------------|-----|
| A.1 | Local Access Windows | A-1 |
| A.2 | System Configuration Registers | A-3 |

Contents

| Paragraph Number | Title | Page Number |
|------------------|---|-------------|
| A.3 | Watchdog Timer (WDT) | A-3 |
| A.4 | Real Time Clock (RTC) | A-4 |
| A.5 | Periodic Interval Timer (PIT) | A-4 |
| A.6 | General Purpose (Global) Timers (GTMs) | A-4 |
| A.7 | Integrated Programmable Interrupt Controller (IPIC) | A-5 |
| A.8 | System Arbiter | A-6 |
| A.9 | Reset Configuration | A-7 |
| A.10 | Clock Configuration | A-7 |
| A.11 | Power Management Controller (PMC) | A-8 |
| A.12 | General Purpose I/O (GPIO) | A-8 |
| A.13 | DDR Memory Controller | A-9 |
| A.14 | I ² C Controller | A-10 |
| A.15 | DUART | A-10 |
| A.16 | Enhanced Local Bus Controller (eLBC) | A-11 |
| A.17 | Serial Peripheral Interface (SPI) | A-13 |
| A.18 | DMA Controller | A-13 |
| A.19 | PCI Configuration Access | A-15 |
| A.20 | I/O Sequencer (IOS) | A-15 |
| A.21 | PCI Controller | A-16 |
| A.22 | PCI Express Controller | A-17 |
| A.23 | Time Division Multiplexing (TDM) Interface | A-25 |
| A.24 | Serial ATA (SATA) Controller | A-27 |
| A.25 | USB DR Controller | A-28 |
| A.26 | Enhanced Three-Speed Ethernet Controllers (eTSECs) | A-29 |
| A.27 | TDM DMA Controller (DMAC) | A-40 |
| A.28 | Security Engine Controller (SEC) | A-41 |
| A.29 | SerDes PHY | A-45 |

Appendix B Revision History

| | | |
|-----|---|-------|
| B.1 | Changes From Revision 1 to Revision 2 | B-1 |
| B.2 | Changes From Revision 0 to Revision 1 | B-110 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 1-1 | MPC8315E Block Diagram | 1-2 |
| 1-2 | MPC8315E Integrated e300c3 Core Block Diagram..... | 1-12 |
| 1-3 | Integrated Security Engine Functional Blocks..... | 1-13 |
| 1-4 | USB Controllers Port Configuration..... | 1-17 |
| 1-5 | MPC8315E as a Media Server | 1-23 |
| 1-6 | MPC8314E Serving as a Low-End Voice Gateway Application | 1-24 |
| 1-7 | MPC8315E as a WLAN Access Point..... | 1-25 |
| 2-1 | MPC8315E Signal Groupings (1 of 2)..... | 2-3 |
| 2-2 | MPC8315E Signal Groupings (2 of 2)..... | 2-4 |
| 4-1 | Power-On Reset Flow | 4-8 |
| 4-2 | Hard Reset Flow..... | 4-9 |
| 4-3 | Reset Configuration Word Low Register (RCWLR)..... | 4-13 |
| 4-4 | Reset Configuration Word High Register (RCWHR)..... | 4-15 |
| 4-5 | EEPROM Data Format for Reset Configuration Words Preload Command | 4-25 |
| 4-6 | EEPROM Contents | 4-26 |
| 4-7 | Clock Subsystem Block Diagram | 4-30 |
| 4-8 | Reset Status Register (RSR)..... | 4-34 |
| 4-9 | Reset Mode Register (RMR)..... | 4-36 |
| 4-10 | Reset Protection Register (RPR)..... | 4-36 |
| 4-11 | Reset Control Register (RCR)..... | 4-37 |
| 4-12 | Reset Control Enable Register (RCER)..... | 4-38 |
| 4-13 | System PLL Mode Register | 4-39 |
| 4-14 | Output Clock Control Register (OCCR)..... | 4-40 |
| 4-15 | System Clock Control Register (SCCR)..... | 4-41 |
| 5-1 | Local Memory Map Example | 5-2 |
| 5-2 | Internal Memory Map Registers' Base Address Register (IMMRBAR)..... | 5-7 |
| 5-3 | Alternate Configuration Base Address Register (ALTCBAR) | 5-8 |
| 5-4 | LBC Local Access Window <i>n</i> Base Address Registers (LBLAWBAR0–LBLAWBAR3) ... | 5-8 |
| 5-5 | LBC Local Access Window <i>n</i> Attributes Registers (LBLAWAR0–LBLAWAR3) | 5-9 |
| 5-6 | PCI Local Access Window <i>n</i> Base Address Registers (PCILAWBAR0–PCILAWBAR1) . | 5-10 |
| 5-7 | PCI Local Access Window <i>n</i> Attributes Registers (PCILAWAR0–PCILAWAR1)..... | 5-11 |
| 5-8 | PCI Express 1 Local Access Window Base Address Register (PCIEXP1LAWBAR) | 5-12 |
| 5-9 | PCI Express 1 Local Access Window Attributes Register (PCIEXP1LAWAR) | 5-12 |
| 5-10 | PCI Express 2 Local Access Window Base Address Register (PCIEXP2LAWBAR) | 5-13 |
| 5-11 | PCI Express 2 Local Access Window Attributes Register (PCIEXP2LAWAR) | 5-14 |
| 5-12 | DDR Local Access Window <i>n</i> Base Address Registers (DDRLAWBAR0–DDRLAWBAR1)... | 5-14 |
| 5-13 | DDR Local Access Window <i>n</i> Attributes Registers (DDRLAWAR0–DDRLAWAR1) | 5-15 |
| 5-14 | System General Purpose Register Low (SGPRL)..... | 5-19 |
| 5-15 | System General Purpose Register High (SGPRH) | 5-20 |
| 5-16 | System Part and Revision ID Register (SPRIDR) | 5-20 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 5-17 | System Priority Configuration Register (SPCR) | 5-22 |
| 5-18 | System I/O Configuration Register Low (SICRL) | 5-24 |
| 5-19 | System I/O Configuration Register High (SICRH) | 5-27 |
| 5-20 | DDR Control Driver Register (DDRCDR)..... | 5-30 |
| 5-21 | DDR Debug Status Register (DDRDSR)..... | 5-31 |
| 5-22 | PCI Express Controller Registers (PECR1 and PECR2)..... | 5-32 |
| 5-23 | Software Watchdog Timer High-Level Block Diagram | 5-33 |
| 5-24 | System Watchdog Control Register (SWCRR)..... | 5-35 |
| 5-25 | System Watchdog Count Register (SWCNR)..... | 5-36 |
| 5-26 | System Watchdog Service Register (SWSRR) | 5-37 |
| 5-27 | Software Watchdog Timer Service State Diagram..... | 5-38 |
| 5-28 | Software Watchdog Timer Functional Block Diagram..... | 5-38 |
| 5-29 | Real Time Clock Module High Level Block Diagram | 5-40 |
| 5-30 | Real Time Counter Control Register (RTCNR)..... | 5-42 |
| 5-31 | Real Time Counter Load Register (RTLDR) | 5-43 |
| 5-32 | Real Time Counter Prescale Register (RTPSR)..... | 5-43 |
| 5-33 | Real Time Counter Register (RTCTR)..... | 5-44 |
| 5-34 | Real Time Counter Event Register (RTEVR) | 5-44 |
| 5-35 | Real Time Counter Alarm Register (RTALR) | 5-45 |
| 5-36 | Real Time Clock Module Functional Block Diagram | 5-46 |
| 5-37 | Periodic Interval Timer High Level Block Diagram..... | 5-47 |
| 5-38 | Periodic Interval Timer Control Register (PTCNR) | 5-49 |
| 5-39 | Periodic Interval Timer Load Register (PTLDR)..... | 5-50 |
| 5-40 | Periodic Interval Timer Prescale Register (PTPSR) | 5-50 |
| 5-41 | Periodic Interval Timer Counter Register (PTCTR) | 5-51 |
| 5-42 | Periodic Interval Timer Event Register (PTEVR)..... | 5-51 |
| 5-43 | Periodic Interval Timer Functional Block Diagram..... | 5-52 |
| 5-44 | Global Timers Block Diagram | 5-54 |
| 5-45 | Global Timers Configuration Register 1 (GTCFR1)..... | 5-59 |
| 5-46 | Global Timers Configuration Register 2 (GTCFR2)..... | 5-60 |
| 5-47 | Global Timers Mode Registers (GTMDR1–GTMDR4)..... | 5-62 |
| 5-48 | Global Timers Reference Registers (GTRFR1–GTRFR4)..... | 5-63 |
| 5-49 | Global Timers Capture Registers (GTCPR1–GTCPR4) | 5-63 |
| 5-50 | Global Timers Counter Registers (GTCNR1—GTCNR4)..... | 5-64 |
| 5-51 | Global Timers Event Registers (GTEVR1—GTEVR4)..... | 5-64 |
| 5-52 | Global Timers Prescale Registers (GTPSR1–GTPSR4)..... | 5-65 |
| 5-53 | Timers Non-Cascaded Mode Block Diagram | 5-67 |
| 5-54 | Timer Pair-Cascaded Mode Block Diagram | 5-68 |
| 5-55 | Timers Super-Cascaded Mode Block Diagram..... | 5-68 |
| 5-56 | Power Management Controller Configuration Register | 5-71 |
| 5-57 | Power Management Controller Event Register | 5-72 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 5-58 | Power Management Controller Mask Register | 5-74 |
| 5-59 | Power Management Controller Configuration Register 1 | 5-75 |
| 5-60 | Power Management Controller Configuration Register 2 | 5-77 |
| 5-61 | Power Segmentation in Deep Sleep Mode..... | 5-82 |
| 5-62 | Power State Transitions Supported | 5-91 |
| 5-63 | PMC Power-Down Sequence to D3Warm | 5-93 |
| 5-64 | PMC Power-Down Sequence to D3Warm (continued)..... | 5-94 |
| 5-65 | PMC Wake-Up Sequence from D3Warm | 5-96 |
| 5-66 | PMC Wake-Up Sequence from D3Warm (continued)..... | 5-97 |
| 5-67 | Example VDD Control of Device as Host, Using Optional PMC_PWR_OK Signal | 5-98 |
| 6-1 | Arbiter Configuration Register (ACR) | 6-3 |
| 6-2 | Arbiter Timers Register (ATR) | 6-5 |
| 6-3 | Arbiter Event Enable Register (AEER) | 6-5 |
| 6-4 | Arbiter Event Register (AER)..... | 6-6 |
| 6-5 | Arbiter Interrupt Definition Register (AIDR)..... | 6-7 |
| 6-6 | Arbiter Mask Register (AMR) | 6-8 |
| 6-7 | Arbiter Event Attributes Register (AEATR)..... | 6-9 |
| 6-8 | Arbiter Event Address Register (AEADR)..... | 6-11 |
| 6-9 | Arbiter Event Response Register (AERR)..... | 6-12 |
| 6-10 | Address Bus Arbitration..... | 6-13 |
| 6-11 | An Example of Priority-Based Arbitration Algorithm | 6-14 |
| 7-1 | e300c3 Core Block Diagram..... | 7-2 |
| 7-2 | e300 Programming Model—Registers..... | 7-16 |
| 7-3 | Machine State Register (MSR) | 7-18 |
| 7-4 | e300c3 Data Cache Organization..... | 7-30 |
| 7-5 | Core Interface..... | 7-38 |
| 8-1 | Interrupt Sources Block Diagram | 8-3 |
| 8-2 | System Global Interrupt Configuration Register (SICFR) | 8-8 |
| 8-3 | System Global Interrupt Vector Register (SIVCR)..... | 8-9 |
| 8-4 | System Internal Interrupt Pending Register (SIPNR_H) | 8-12 |
| 8-5 | System Internal Interrupt Pending Register (SIPNR_L)..... | 8-13 |
| 8-6 | System Internal Interrupt Group A Priority Register (SIPRR_A) | 8-15 |
| 8-7 | System Internal Interrupt Group B Priority Register (SIPRR_B)..... | 8-16 |
| 8-8 | System Internal Interrupt Group C Priority Register (SIPRR_C)..... | 8-17 |
| 8-9 | System Internal Interrupt Group D Priority Register (SIPRR_D) | 8-17 |
| 8-10 | System Internal Interrupt Mask Register (SIMSR_H)..... | 8-18 |
| 8-11 | System Internal Interrupt Mask Register (SIMSR_L) | 8-19 |
| 8-12 | System Internal Interrupt Control Register (SICNR) | 8-20 |
| 8-13 | System External Interrupt Pending Register (SEPNR)..... | 8-21 |
| 8-14 | System Mixed Interrupt Group A Priority Register (SMPRR_A)..... | 8-22 |
| 8-15 | System Mixed Interrupt Group B Priority Register (SMPRR_B) | 8-23 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 8-16 | System External Interrupt Mask Register (SEMSR) | 8-24 |
| 8-17 | System External Interrupt Control Register (SECNR) | 8-25 |
| 8-18 | System Error Status Register (SERSR)..... | 8-26 |
| 8-19 | System Error Mask Register (SERMR)..... | 8-27 |
| 8-20 | System Error Control Register (SERCR)..... | 8-27 |
| 8-21 | System External Interrupt Polarity Control Register (SEPCR) | 8-28 |
| 8-22 | System Internal Interrupt Force Register (SIFCR_H) | 8-29 |
| 8-23 | System Internal Interrupt Force Register (SIFCR_L)..... | 8-29 |
| 8-24 | System External Interrupt Force Register (SEFCR) | 8-30 |
| 8-25 | System Error Status Register (SERFR)..... | 8-30 |
| 8-26 | System Critical Interrupt Vector Register (SCVCR) | 8-31 |
| 8-27 | System Management Interrupt Vector Register (SMVCR)..... | 8-32 |
| 8-28 | Interrupt Structure | 8-33 |
| 8-29 | DDR Interrupt Request Masking | 8-39 |
| 8-30 | Message Shared Interrupt Register (MSIRs) | 8-41 |
| 8-31 | Message Shared Interrupt Mask Register (MSIMR) | 8-41 |
| 8-32 | Message Shared Interrupt Status Register (MSISR) | 8-42 |
| 8-33 | Message Shared Interrupt Index Register (MSIIR) | 8-43 |
| 9-1 | DDR Memory Controller Simplified Block Diagram..... | 9-2 |
| 9-2 | Chip Select Bounds Registers (CS _n _BNDS)..... | 9-10 |
| 9-3 | Chip Select Configuration Register (CS _n _CONFIG)..... | 9-10 |
| 9-4 | DDR SDRAM Timing Configuration 3 (TIMING_CFG_3) | 9-12 |
| 9-5 | DDR SDRAM Timing Configuration 0 (TIMING_CFG_0) | 9-12 |
| 9-6 | DDR SDRAM Timing Configuration 1 (TIMING_CFG_1) | 9-14 |
| 9-7 | DDR SDRAM Timing Configuration 2 Register (TIMING_CFG_2)..... | 9-16 |
| 9-8 | DDR SDRAM Control Configuration Register (DDR_SDRAM_CFG)..... | 9-18 |
| 9-9 | DDR SDRAM Control Configuration Register 2 (DDR_SDRAM_CFG_2)..... | 9-21 |
| 9-10 | DDR SDRAM Mode Configuration Register (DDR_SDRAM_MODE)..... | 9-22 |
| 9-11 | DDR SDRAM Mode 2 Configuration Register (DDR_SDRAM_MODE_2)..... | 9-23 |
| 9-12 | DDR SDRAM Mode Control Register (DDR_SDRAM_MD_CNTL)..... | 9-24 |
| 9-13 | DDR SDRAM Interval Configuration Register (DDR_SDRAM_INTERVAL) | 9-27 |
| 9-14 | DDR SDRAM Data Initialization Configuration Register (DDR_DATA_INIT)..... | 9-27 |
| 9-15 | DDR SDRAM Clock Control Configuration Register (DDR_SDRAM_CLK_CNTL)..... | 9-28 |
| 9-16 | DDR Initialization Address Configuration Register (DDR_INIT_ADDR) | 9-28 |
| 9-17 | DDR IP Block Revision 1 (DDR_IP_REV1) | 9-29 |
| 9-18 | DDR IP Block Revision 2 (DDR_IP_REV2) | 9-29 |
| 9-19 | Memory Error Detect Register (ERR_DETECT)..... | 9-30 |
| 9-20 | Memory Error Disable Register (ERR_DISABLE)..... | 9-30 |
| 9-21 | Memory Error Interrupt Enable Register (ERR_INT_EN)..... | 9-31 |
| 9-22 | DDR Memory Controller Block Diagram | 9-32 |
| 9-23 | Typical Dual Data Rate SDRAM Internal Organization..... | 9-33 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 9-24 | Typical DDR SDRAM Interface Signals | 9-33 |
| 9-25 | Example 64-Mbyte DDR SDRAM Configuration..... | 9-34 |
| 9-26 | DDR SDRAM Burst Read Timing—ACTTORW = 3, MCAS Latency = 2 | 9-44 |
| 9-27 | DDR SDRAM Single-Beat (Double Word) Write Timing—ACTTOR | 9-45 |
| 9-28 | DDR SDRAM Single-Beat (Double Word) Write Timing—ACTTORW = 3..... | 9-45 |
| 9-29 | DDR SDRAM Clock Distribution Example for x8 DDR SDRAMs | 9-46 |
| 9-30 | DDR SDRAM Mode-Set Command Timing | 9-46 |
| 9-31 | Registered DDR SDRAM DIMM Burst Write Timing | 9-47 |
| 9-32 | Write Timing Adjustments Example for Write Latency = 1 | 9-48 |
| 9-33 | DDR SDRAM Bank Staggered Auto Refresh Timing..... | 9-49 |
| 9-34 | DDR SDRAM Power-Down Mode | 9-50 |
| 9-35 | DDR SDRAM Self-Refresh Entry Timing | 9-51 |
| 9-36 | DDR SDRAM Self-Refresh Exit Timing | 9-51 |
| 10-1 | Enhanced Local Bus Controller Block Diagram..... | 10-1 |
| 10-2 | Base Registers (BR _n) | 10-10 |
| 10-3 | Option Registers (OR _n) in GPCM Mode..... | 10-13 |
| 10-4 | Option Registers (OR _n) in FCM Mode..... | 10-15 |
| 10-5 | Option Registers (OR _n) in UPM Mode | 10-18 |
| 10-6 | UPM Memory Address Register (MAR) | 10-19 |
| 10-7 | UPM Mode Registers (MxMR)..... | 10-20 |
| 10-8 | Memory Refresh Timer Prescaler Register (MRTPR)..... | 10-22 |
| 10-9 | UPM Data Register in UPM Mode (MDR) | 10-23 |
| 10-10 | FCM Data Register in FCM Mode (MDR)..... | 10-23 |
| 10-11 | Special Operation Initiation Register (LSOR)..... | 10-24 |
| 10-12 | UPM Refresh Timer (LURT) | 10-24 |
| 10-13 | Transfer Error Status Register (LTESR) | 10-25 |
| 10-14 | Transfer Error Check Disable Register (LTEDR)..... | 10-27 |
| 10-15 | Transfer Error Interrupt Enable Register (LTEIR)..... | 10-28 |
| 10-16 | Transfer Error Attributes Register (LTEATR)..... | 10-29 |
| 10-17 | Transfer Error Address Register (LTEAR) | 10-30 |
| 10-18 | Local Bus Configuration Register..... | 10-30 |
| 10-19 | Clock Ratio Register (LCRR) | 10-32 |
| 10-20 | Flash Mode Register | 10-33 |
| 10-21 | Flash Instruction Register | 10-35 |
| 10-22 | Flash Command Register | 10-35 |
| 10-23 | Flash Block Address Register | 10-36 |
| 10-24 | Flash Page Address Register, Small Page Device (ORx[PGS] = 0) | 10-36 |
| 10-25 | Flash Page Address Register, Large Page Device (ORx[PGS] = 1) | 10-37 |
| 10-26 | Flash Byte Count Register | 10-38 |
| 10-27 | Basic Operation of Memory Controllers in the eLBC | 10-39 |
| 10-28 | Example of 8-Bit GPCM Writing 32 Bytes to Address 0x5420 (LCRR[PBYP] = 0)..... | 10-41 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 10-29 | Basic eLBC Bus Cycle with LALE, TA, and $\overline{\text{LCSn}}$ | 10-42 |
| 10-30 | Enhanced Local Bus to GPCM Device Interface..... | 10-43 |
| 10-31 | GPCM Basic Read Timing (XACS = 0, ACS = 1x, TRLX = 0, CLKDIV = 4,8)..... | 10-44 |
| 10-32 | GPCM General Read Timing Parameters | 10-44 |
| 10-33 | GPCM General Write Timing Parameters | 10-46 |
| 10-34 | GPCM Basic Write Timing (XACS = 0, ACS = 00, CSNT = 1, SCY = 1, TRLX = 0, CLKDIV = 4, 8)..... | 10-48 |
| 10-35 | GPCM Relaxed Timing Back-to-Back Reads (XACS = 0, ACS = 1x, SCY = 1, CSNT = 0, TRLX = 1, EHTR = 0, CLKDIV = 4, 8) | 10-50 |
| 10-36 | GPCM Relaxed Timing Back-to-Back Writes (XACS = 0, ACS = 1x, SCY = 0, CSNT = 0, TRLX = 1, CLKDIV = 4, 8)..... | 10-50 |
| 10-37 | GPCM Relaxed Timing Write (XACS = 0, ACS = 10, SCY = 0, CSNT = 1, TRLX = 1, CLKDIV = 4, 8)..... | 10-51 |
| 10-38 | GPCM Relaxed Timing Write (XACS = 0, ACS = 00, SCY = 1, CSNT = 1, TRLX = 1, CLKDIV = 4, 8)..... | 10-51 |
| 10-39 | GPCM Read Followed by Read (TRLX = 0, EHTR = 0, Fastest Timing)..... | 10-52 |
| 10-40 | GPCM Read Followed by Write (TRLX = 0, EHTR = 1, One-Cycle Extended Hold Time on Reads) | 10-53 |
| 10-41 | External Termination of GPCM Access (PLL Bypass Mode)..... | 10-54 |
| 10-42 | Local Bus to 8-Bit FCM Device Interface | 10-56 |
| 10-43 | FCM Basic Page Read Timing (PGS = 1, CSCT = 0, CST = 0, CHT = 1, RST = 1, SCY = 0, TRLX = 0, EHTR = 1) | 10-57 |
| 10-44 | FCM Buffer RAM Memory Map for Small-Page (512-byte page) NAND Flash Devices | 10-59 |
| 10-45 | FCM Buffer RAM Memory Map for Large-Page (2-Kbyte page) NAND Flash Devices . | 10-60 |
| 10-46 | FCM ECC Calculation | 10-60 |
| 10-47 | ECC Placement in NAND Flash Spare Regions in Relation to FMR[ECCM] | 10-61 |
| 10-48 | FCM Instruction Sequencer Mechanism..... | 10-62 |
| 10-49 | Timing of FCM Command/Address and Write Data Cycles (for TRLX = 0, CHT = 0, CST = 1, SCY = 1, CLKDIV = 4*N)..... | 10-65 |
| 10-50 | Example of FCM Command and Address Timing with Minimum Delay Parameters (for TRLX = 0, CHT = 0, CST = 0, SCY = 0, CLKDIV = 4*N)..... | 10-66 |
| 10-51 | Example of FCM Command and Address Timing with Relaxed Parameters (for TRLX = 1, CHT = 0, CST = 1, SCY = 2, CLKDIV = 4*N)..... | 10-66 |
| 10-52 | FCM Delay Prior to Sampling LFRB State | 10-67 |
| 10-53 | FCM Read Data Timing (for TRLX = 0, RST = 0, SCY = 1, CLKDIV = 4*N) | 10-67 |
| 10-54 | FCM Read Data Timing with Extended Hold Time (for TRLX = 0, EHTR = 1, RST = 1, SCY = 1, CLKDIV = 4*N) | 10-68 |
| 10-55 | FCM Buffer RAM Memory Map During Boot Loading | 10-70 |
| 10-56 | User-Programmable Machine Functional Block Diagram..... | 10-71 |
| 10-57 | RAM Array Indexing | 10-72 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 10-58 | Memory Refresh Timer Request Block Diagram | 10-73 |
| 10-59 | UPM Clock Scheme for LCRR[CLKDIV] = 2..... | 10-77 |
| 10-60 | UPM Clock Scheme for LCRR[CLKDIV] = 4 or 8 | 10-77 |
| 10-61 | RAM Array and Signal Generation | 10-77 |
| 10-62 | RAM Word Fields | 10-78 |
| 10-63 | LCSn Signal Selection | 10-82 |
| 10-64 | LBS Signal Selection | 10-83 |
| 10-65 | UPM Read Access Data Sampling..... | 10-86 |
| 10-66 | Effect of LUPWAIT Signal..... | 10-87 |
| 10-67 | Multiplexed Address/Data Bus for 26-Bit Addressing | 10-88 |
| 10-68 | Local Bus Peripheral Hierarchy for High Bus Speeds..... | 10-89 |
| 10-69 | GPCM Address Timings | 10-89 |
| 10-70 | GPCM Data Timings..... | 10-90 |
| 10-71 | Interface to Different Port-Size Devices | 10-91 |
| 10-72 | Single-Beat Read Access to FPM DRAM | 10-97 |
| 10-73 | Single-Beat Write Access to FPM DRAM | 10-98 |
| 10-74 | Burst Read Access to FPM DRAM Using LOOP (Two Beats Shown)..... | 10-99 |
| 10-75 | Refresh Cycle (CBR) to FPM DRAM | 10-100 |
| 10-76 | Exception Cycle | 10-101 |
| 10-77 | Interface to ZBT SRAM | 10-102 |
| 11-1 | I/O Sequencer Block Diagram | 11-1 |
| 11-2 | PCI Outbound Translation Address Registers (POTAR _n)..... | 11-3 |
| 11-3 | PCI Outbound Base Address Registers (POBAR _n)..... | 11-3 |
| 11-4 | PCI Outbound Comparison Mask Registers (POCMR _n) | 11-4 |
| 11-5 | Power Management Control Register (PMCR) | 11-5 |
| 11-6 | Discard Timer Control Register (DTCR)..... | 11-6 |
| 11-7 | Outbound PCI Memory Address Translation | 11-8 |
| 12-1 | DMA/Messaging Unit Block Diagram | 12-1 |
| 12-2 | Outbound Message Interrupt Status Register (OMISR) | 12-4 |
| 12-3 | Outbound Message Interrupt Mask Register (OMIMR)..... | 12-5 |
| 12-4 | Inbound Message Registers (IMR0, IMR1)..... | 12-6 |
| 12-5 | Outbound Message Registers (OMR0–OMR1)..... | 12-6 |
| 12-6 | Outbound Doorbell Register (ODR) | 12-7 |
| 12-7 | Inbound Doorbell Register (IDR) | 12-7 |
| 12-8 | Inbound Message Interrupt Status Register (IMISR)..... | 12-8 |
| 12-9 | Inbound Message Interrupt Mask Register (IMIMR) | 12-9 |
| 12-10 | DMA Mode Register (DMAMR _n) | 12-10 |
| 12-11 | DMA Status Register (DMASR _n) | 12-12 |
| 12-12 | DMA Current Descriptor Address Register (DMACDAR _n)..... | 12-13 |
| 12-13 | DMA Source Address Register (DMASAR _n)..... | 12-14 |
| 12-14 | DMA Destination Address Register (DMADAR _n)..... | 12-14 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 12-15 | DMA Byte Count Register (DMABCR _n)..... | 12-15 |
| 12-16 | DMA Next Descriptor Address Register (DMANDAR _n)..... | 12-15 |
| 12-17 | DMA General Status Register (DMAGSR)..... | 12-16 |
| 12-18 | DMA Controller Block Diagram | 12-17 |
| 12-19 | DMA Chain of Segment Descriptors | 12-21 |
| 13-1 | PCI Controller Block Diagram | 13-2 |
| 13-2 | PCI Interface External Signals | 13-5 |
| 13-3 | PCI_CONFIG_ADDRESS Register | 13-13 |
| 13-4 | PCI_CONFIG_DATA | 13-15 |
| 13-5 | PCI Error Status Register (PCI_ESR)..... | 13-15 |
| 13-6 | PCI Error Capture Disable Register (PCI_ECDR) | 13-16 |
| 13-7 | PCI Error Enable Register (PCI_EER) | 13-17 |
| 13-8 | PCI Error Attributes Capture Register (PCI_EATCR) | 13-18 |
| 13-9 | PCI Error Address Capture Register (PCI_EACR) | 13-19 |
| 13-10 | PCI Error Extended Address Capture Register (PCI_EEACR)..... | 13-20 |
| 13-11 | PCI Error Data Low Capture Register (PCI_EDLCR) | 13-20 |
| 13-12 | PCI General Control Register (PCI_GCR) | 13-21 |
| 13-13 | PCI Error Control Register (PCI_ECR)..... | 13-21 |
| 13-14 | PCI General Status Register (PCI_GSR)..... | 13-22 |
| 13-15 | PCI Inbound Translation Address Registers (PITAR _n) | 13-23 |
| 13-16 | PCI Inbound Base Address Registers (PIBAR _n)..... | 13-23 |
| 13-17 | PCI Inbound Extended Base Address Registers (PIEBAR _n) | 13-24 |
| 13-18 | PCI Inbound Window Attribute Registers (PIWAR _n) | 13-24 |
| 13-19 | Vendor ID Configuration Register | 13-27 |
| 13-20 | Device ID Configuration Register | 13-27 |
| 13-21 | PCI Command Configuration Register | 13-28 |
| 13-22 | PCI Status Configuration Register | 13-29 |
| 13-23 | Revision ID Configuration Register | 13-30 |
| 13-24 | Standard Programming Interface Configuration Register..... | 13-30 |
| 13-25 | Subclass Code Configuration Register | 13-31 |
| 13-26 | Base Class Code Configuration Register | 13-31 |
| 13-27 | Cache Line Size Configuration Register..... | 13-32 |
| 13-28 | Latency Timer Configuration Register | 13-32 |
| 13-29 | Header Type Configuration Register | 13-33 |
| 13-30 | BIST Control Configuration Register | 13-33 |
| 13-31 | PIMMR Base Address Configuration Register..... | 13-33 |
| 13-32 | GPL Base Address Register 0..... | 13-34 |
| 13-33 | GPL Base Address Registers 1–2 | 13-35 |
| 13-34 | GPL Extended Base Address Registers 1–2 | 13-35 |
| 13-35 | Subsystem Vendor ID Configuration Register..... | 13-36 |
| 13-36 | Subsystem Device ID Configuration Register | 13-36 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 13-37 | Capabilities Pointer Configuration Register | 13-37 |
| 13-38 | Interrupt Line Configuration Register..... | 13-37 |
| 13-39 | Interrupt Pin Register | 13-37 |
| 13-40 | Minimum Grant Configuration Register..... | 13-38 |
| 13-41 | Maximum Latency Configuration Register | 13-38 |
| 13-42 | PCI Function Configuration Register | 13-38 |
| 13-43 | PCI Arbiter Control Register (PCIACR) | 13-39 |
| 13-44 | Hot Swap Register Block..... | 13-40 |
| 13-45 | PCI Power Management Register 0 (PCIPMR0)..... | 13-41 |
| 13-46 | PCI Power Management Register 1 (PCIPMR1)..... | 13-42 |
| 13-47 | PCI Arbitration Example | 13-45 |
| 13-48 | Single Beat Read Example..... | 13-49 |
| 13-49 | Burst Read Example..... | 13-50 |
| 13-50 | Single Beat Write Example | 13-50 |
| 13-51 | Burst Write Example | 13-51 |
| 13-52 | Target-Initiated Terminations | 13-52 |
| 13-53 | PCI Parity Operation | 13-58 |
| 13-54 | Inbound PCI Memory Address Translation | 13-59 |
| 13-55 | Address Invariant Byte Ordering—4 bytes Outbound..... | 13-60 |
| 13-56 | Address Invariant Byte Ordering—4 bytes Inbound | 13-61 |
| 13-57 | Address Invariant Byte Ordering—8 bytes Outbound..... | 13-61 |
| 13-58 | Address Invariant Byte Ordering—2 bytes Inbound | 13-61 |
| 13-59 | CFG_DATA Byte Ordering..... | 13-62 |
| 14-1 | PCI Express Controller Block Diagram..... | 14-2 |
| 14-2 | PCI Express PCI-Compatible Configuration Header Common Registers..... | 14-15 |
| 14-3 | PCI Express Vendor ID Register..... | 14-15 |
| 14-4 | PCI Express Device ID Register | 14-16 |
| 14-5 | PCI Express Command Register..... | 14-16 |
| 14-6 | PCI Express Status Register..... | 14-17 |
| 14-7 | PCI Express Revision ID Register | 14-18 |
| 14-8 | PCI Express Class Code Register | 14-19 |
| 14-9 | PCI Express Bus Cache Line Size Register | 14-19 |
| 14-10 | PCI Express Latency Timer Register | 14-20 |
| 14-11 | PCI Express Header Type Register | 14-20 |
| 14-12 | PCI Express PCI-Compatible Configuration Header—Type 0..... | 14-21 |
| 14-13 | 32-Bit Base Address Registers (BAR0/BAR1) | 14-22 |
| 14-14 | 64-Bit Low Memory Base Address Register (BAR2) | 14-23 |
| 14-15 | 64-Bit High Memory Base Address Registers 3 and 5 (BAR3/BAR5)..... | 14-23 |
| 14-16 | PCI Express Subsystem Vendor ID Register | 14-24 |
| 14-17 | PCI Express Subsystem ID Register | 14-24 |
| 14-18 | PCI Express Capabilities Pointer Register..... | 14-25 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 14-19 | PCI Express Interrupt Line Register | 14-25 |
| 14-20 | PCI Express Interrupt Pin Register | 14-26 |
| 14-21 | PCI Express Minimum Grant Register (MAX_GNT) | 14-26 |
| 14-22 | PCI Express Maximum Latency Register (MAX_LAT)..... | 14-26 |
| 14-23 | PCI Express PCI-Compatible Configuration Header—Type 1 | 14-27 |
| 14-24 | PCI Express Primary Bus Number Register | 14-27 |
| 14-25 | PCI Express Secondary Bus Number Register | 14-28 |
| 14-26 | PCI Express Subordinate Bus Number Register | 14-28 |
| 14-27 | PCI Express I/O Base Register | 14-29 |
| 14-28 | PCI Express I/O Limit Register | 14-29 |
| 14-29 | PCI Express Secondary Status Register | 14-30 |
| 14-30 | PCI Express Memory Base Register | 14-31 |
| 14-31 | PCI Express Memory Limit Register | 14-31 |
| 14-32 | PCI Express Prefetchable Memory Base Register | 14-32 |
| 14-33 | PCI Express Prefetchable Memory Limit Register | 14-32 |
| 14-34 | PCI Express Prefetchable Base Upper 32-Bit Register | 14-33 |
| 14-35 | PCI Express Prefetchable Limit Upper 32-Bit Register | 14-33 |
| 14-36 | PCI Express I/O Base Upper 16-Bit Register | 14-33 |
| 14-37 | PCI Express I/O Limit Upper 16-Bit Register | 14-34 |
| 14-38 | PCI Express Capabilities Pointer Register | 14-34 |
| 14-39 | PCI Express Interrupt Line Register | 14-35 |
| 14-40 | PCI Express Interrupt Pin Register | 14-35 |
| 14-41 | PCI Express Bridge Control Register | 14-35 |
| 14-42 | PCI-Compatible Device-Specific Configuration Space | 14-37 |
| 14-43 | PCI Express Power Management Capability ID Register | 14-38 |
| 14-44 | PCI Express Power Management Next Capabilities Pointer | 14-38 |
| 14-45 | PCI Express Power Management Capabilities Register | 14-39 |
| 14-46 | PCI Express Power Management Status and Control Register..... | 14-39 |
| 14-47 | PCI Express Power Management Data Register | 14-40 |
| 14-48 | PCI Express Capability ID Register..... | 14-40 |
| 14-49 | PCI Express Next Capabilities Pointer | 14-41 |
| 14-50 | PCI Express Capabilities Register | 14-41 |
| 14-51 | PCI Express Device Capabilities Register | 14-42 |
| 14-52 | PCI Express Device Control Register | 14-43 |
| 14-53 | PCI Express Device Status Register | 14-44 |
| 14-54 | PCI Express Link Capabilities Register..... | 14-44 |
| 14-55 | PCI Express Link Control Register..... | 14-45 |
| 14-56 | PCI Express Link Status Register | 14-46 |
| 14-57 | PCI Express Slot Capabilities Register | 14-46 |
| 14-58 | PCI Express Slot Control Register | 14-47 |
| 14-59 | PCI Express Slot Status Register | 14-48 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 14-60 | PCI Express Root Control Register | 14-48 |
| 14-61 | PCI Express Root Status Register | 14-49 |
| 14-62 | PCI Express Capability ID Register..... | 14-49 |
| 14-63 | PCI Express MSI Message Control Register | 14-50 |
| 14-64 | PCI Express MSI Message Address Register | 14-50 |
| 14-65 | PCI Express MSI Message Upper Address Register | 14-51 |
| 14-66 | PCI Express MSI Message Data Register | 14-51 |
| 14-67 | PCI Express Extended Configuration Space..... | 14-52 |
| 14-68 | PCI Express Advanced Error Reporting Capability ID Register..... | 14-53 |
| 14-69 | PCI Express Uncorrectable Error Status Register..... | 14-53 |
| 14-70 | PCI Express Uncorrectable Error Mask Register | 14-54 |
| 14-71 | PCI Express Uncorrectable Error Severity Register | 14-55 |
| 14-72 | PCI Express Correctable Error Status Register..... | 14-56 |
| 14-73 | PCI Express Correctable Error Mask Register | 14-57 |
| 14-74 | PCI Express Advanced Error Capabilities and Control Register..... | 14-57 |
| 14-75 | PCI Express Header Log Register | 14-58 |
| 14-76 | PCI Express Root Error Command Register..... | 14-59 |
| 14-77 | PCI Express Root Error Command Register..... | 14-59 |
| 14-78 | PCI Express Error Source Identification Register | 14-60 |
| 14-79 | PCI Express LTSSM State Status Register (PEX_LTSSM_STAT) | 14-61 |
| 14-80 | PCI Express N_FTS Control Register | 14-63 |
| 14-81 | PCI Express ACK Replay Timeout Register | 14-64 |
| 14-82 | PCI Express Core Clock Ratio Register (PEX_GCLK_RATIO)..... | 14-65 |
| 14-83 | PCI Express Power Management Timer Register (PEX_PM_TIMER) | 14-65 |
| 14-84 | PCI Express PME Time-Out Register (PEX_PME_TIMEOUT) | 14-66 |
| 14-85 | PCI Express ASPM Request Timer Register | 14-67 |
| 14-86 | PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE)..... | 14-67 |
| 14-87 | PCI Express Device Capabilities Update Register | 14-68 |
| 14-88 | PCI Express Link Capabilities Update Register | 14-69 |
| 14-89 | PCI Express Slot Capabilities Update Register | 14-70 |
| 14-90 | PCI Express Configuration Ready Register (PEX_CFG_READY) | 14-71 |
| 14-91 | PCI Express BAR Enable Register (PEX_BAR_ENABLE)..... | 14-72 |
| 14-92 | PCI Express BAR Size Low Configuration Register (PEX_BAR_SIZEL) | 14-72 |
| 14-93 | PCI Express BAR Select Configuration Register (PEX_BAR_SEL) | 14-73 |
| 14-94 | PCI Express BAR Prefetch Configuration Register (PEX_BAR_PF)..... | 14-74 |
| 14-95 | PCI Express PME_To_Ack Timeout Register (PEX_PME_TO_ACK_TOR)..... | 14-74 |
| 14-96 | PME_To_Ack Status Register (PEX_PME_TO_ACK_SR)..... | 14-75 |
| 14-97 | PCI Express PCI Interrupt Mask Register (PEX_SS_INTR_MASK)..... | 14-76 |
| 14-98 | PCI Express CSB Bridge Control Register (PEX_CSB_CTRL)..... | 14-77 |
| 14-99 | PCI Express DMA Descriptor Timer Register (PEX_DMA_DSTMR) | 14-78 |
| 14-100 | PCI Express CSB Bridge Status Register (PEX_CSB_STAT) | 14-78 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 14-101 | PCI Express Outbound PIO Control Register (PEX_CSB_OBCTRL) | 14-79 |
| 14-102 | PCI Express Outbound PIO Status Register (PEX_CSB_OBSTAT)..... | 14-80 |
| 14-103 | PCI Express Inbound PIO Control Register (PEX_CSB_IBCTRL) | 14-81 |
| 14-104 | PCI Express Inbound PIO Status Register (PEX_CSB_IBSTAT)..... | 14-81 |
| 14-105 | PCI Express Write DMA Control Register (PEX_WDMA_CTRL) | 14-82 |
| 14-106 | PCI Express Write DMA First Address Register (PEX_WDMA_ADDR)..... | 14-83 |
| 14-107 | PCI Express Write DMA Status Register (PEX_WDMA_STAT)..... | 14-83 |
| 14-108 | PCI Express Read DMA Control Register (PEX_RDMA_CTRL) | 14-84 |
| 14-109 | PCI Express Read DMA First Address Register (PEX_RDMA_ADDR)..... | 14-85 |
| 14-110 | PCI Express Read DMA Status Register (PEX_RDMA_STAT)..... | 14-85 |
| 14-111 | PCI Express Outbound Mailbox Control Register (PEX_OMBCR) | 14-86 |
| 14-112 | MPCI Express Outbound Mailbox Data Register (PEX_OMBDR)..... | 14-87 |
| 14-113 | PCI Express Inbound Mailbox Control Register (PEX_IMBCR) | 14-87 |
| 14-114 | PCI Express Inbound Mailbox Data Register (PEX_IMBDR)..... | 14-88 |
| 14-115 | PCI Express Host Interrupt Enable Register (PEX_HIER) | 14-89 |
| 14-116 | PCI Express Host Interrupt Status Register (PEX_HISR)..... | 14-90 |
| 14-117 | PCI Express Host Outbound PIO Interrupt Vector Register (PEX_HOPIVR)..... | 14-91 |
| 14-118 | PCI Express Host Inbound PIO Interrupt Vector Register (PEX_HIPIVR) | 14-91 |
| 14-119 | PCI Express Host Write DMA Interrupt Vector Register (PEX_HWDIVR)..... | 14-92 |
| 14-120 | PCI Express Host Read DMA Interrupt Vector Register (PEX_HRDIVR) | 14-92 |
| 14-121 | PCI Express Host Miscellaneous Interrupt Vector Register (PEX_HMIVR)..... | 14-93 |
| 14-122 | CSB System PIO Interrupt Enable Register (PEX_CSPIER)..... | 14-93 |
| 14-123 | CSB System Write DMA Interrupt Enable Register (PEX_CSWDIER) | 14-94 |
| 14-124 | CSB System Read DMA Interrupt Enable Register (PEX_CSRDIER) | 14-95 |
| 14-125 | CSB System Miscellaneous Interrupt Enable Register (PEX_CSMIER)..... | 14-95 |
| 14-126 | CSB System PIO Interrupt Status Register (PEX_CSPISR) | 14-97 |
| 14-127 | CSB System Write DMA Interrupt Status Register (PEX_CSWDISR) | 14-97 |
| 14-128 | CSB System Read DMA Interrupt Status Register (PEX_CSRDISR)..... | 14-98 |
| 14-129 | CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR) | 14-99 |
| 14-130 | PCI Express PM Control Register (PEX_PM_CTRL) | 14-100 |
| 14-131 | PCI Express Slot Control Misc Register | 14-101 |
| 14-132 | PCI Express Outbound Window Attributes Register <i>n</i> (PEX_OWAR0–PEX_OWAR3). | 14-102 |
| 14-133 | PCI Express Outbound Window Base Address Register <i>n</i> (PEX_OWBAR0–PEX_OWBAR3) | 14-103 |
| 14-134 | PCI Express Outbound Window Translation Address Register Low <i>n</i> (PEX_OWTARL0–PEX_OWTARL3)..... | 14-104 |
| 14-135 | PCI Express Outbound Window Translation Address Register High <i>n</i> (PEX_OWTARH0–PEX_OWTARH3)..... | 14-104 |
| 14-136 | PCI Express EP Inbound Window Translation Address Register <i>n</i> (PEX_EPIWTAR0–PEX_EPIWTAR3) | 14-105 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 14-137 | PCI Express RC Inbound Window Attributes Register <i>n</i> (PEX_RCIWAR0–PEX_RCIWAR3). 14-106 | |
| 14-138 | PCI Express RC Inbound Window Translation Address Register <i>n</i> (PEX_RCIWTAR0–PEX_RCIWTAR3)..... | 14-107 |
| 14-139 | PCI Express RC Inbound Window Base Address Register Low <i>n</i> (PEX_RCIWBARL0–PEX_RCIWBARL3)..... | 14-108 |
| 14-140 | CI Express RC Inbound Window Base Address Register High <i>n</i> (PEX_RCIWBARH0–PEX_RCIWBARH3)..... | 14-108 |
| 14-141 | Requestor/Completer Relationship | 14-109 |
| 14-142 | PCI Express High-Level Layering | 14-109 |
| 14-143 | PCI Express Packet Flow | 14-110 |
| 14-144 | Address Invariant Byte Ordering—4 bytes Outbound..... | 14-112 |
| 14-145 | Address Invariant Byte Ordering—4 bytes Inbound | 14-112 |
| 14-146 | Address Invariant Byte Ordering—8 bytes Outbound..... | 14-112 |
| 14-147 | Address Invariant Byte Ordering—2 bytes Inbound | 14-113 |
| 14-148 | PEX_CONFIG_DATA Byte Ordering | 14-113 |
| 14-149 | Example on How to Generate WAKE#..... | 14-125 |
| 14-150 | DMA Descriptor Format..... | 14-128 |
| 14-151 | <i>n</i> -Way Chain Descriptor Organization in Host Memory | 14-131 |
| 14-152 | Block Descriptor Organization in Host Memory | 14-132 |
| 15-1 | SATA Block Diagram | 15-2 |
| 15-2 | Command Queue Register (CQR) | 15-6 |
| 15-3 | Command Active Register (CAR) | 15-6 |
| 15-4 | Command Completed Register (CCR) | 15-7 |
| 15-5 | Command Error Register (CER)..... | 15-7 |
| 15-6 | Device Error Register (DER)..... | 15-8 |
| 15-7 | Command Header Base Address Register (CHBA) | 15-8 |
| 15-8 | Host Status Register (HStatus)..... | 15-9 |
| 15-9 | Host Control (HControl) Register..... | 15-12 |
| 15-10 | Port Number Queue Register (CQPMP)..... | 15-13 |
| 15-11 | Signature Register (SIG)..... | 15-14 |
| 15-12 | Interrupt Coalescing Control Register (ICC)..... | 15-14 |
| 15-13 | SATA Interface Status Register (SStatus) | 15-15 |
| 15-14 | SATA Interface Error Register (SError)..... | 15-16 |
| 15-15 | SATA Interface Control Register (SControl) | 15-18 |
| 15-16 | SATA Interface Notification Register (SNotification)..... | 15-20 |
| 15-17 | Transport Layer Configuration Register (TransCfg)..... | 15-20 |
| 15-18 | Transport Layer Status Register (TransStatus)..... | 15-21 |
| 15-19 | Link Layer Configuration Register (LinkCfg)..... | 15-21 |
| 15-20 | Link Layer Configuration Register1 (LinkCfg1)..... | 15-22 |
| 15-21 | Link Layer Configuration Register1 (LinkCfg1)..... | 15-23 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 15-22 | Link Layer Status Register (LinkStatus)..... | 15-23 |
| 15-23 | Link Layer Status Register1 (LinkStatus1)..... | 15-24 |
| 15-24 | PHY Control Configuration Register1 (PhyCtrlCfg1)..... | 15-26 |
| 15-25 | Link Layer Command Status Register (CommandStatus)..... | 15-28 |
| 15-26 | PHY Control Configuration Register2 (PhyCtrlCfg2)..... | 15-29 |
| 15-27 | System Priority Register (SYSPR) | 15-30 |
| 15-28 | Command Header | 15-32 |
| 15-29 | Command Descriptor | 15-34 |
| 15-30 | Register Host-to-Device..... | 15-34 |
| 15-31 | Register Host-to-Device First Party DMA Commands NCQ | 15-35 |
| 15-32 | Register Device-to-Host..... | 15-35 |
| 15-33 | Vendor-Specific BIST Operation | 15-38 |
| 15-34 | PRD Entry | 15-46 |
| 16-1 | SerDes PHY Block Diagram..... | 16-1 |
| 16-2 | Modes of Operation | 16-2 |
| 16-3 | SerDes Control Register 0 (SRDSCR0)..... | 16-5 |
| 16-4 | SerDes Control Register 1 (SRDSCR1)..... | 16-8 |
| 16-5 | SerDes Control Register 2 (SRDSCR2)..... | 16-9 |
| 16-6 | SerDes Control Register 3 (SRDSCR3)..... | 16-10 |
| 16-7 | SerDes Control Register 4 (SRDSCR4)..... | 16-12 |
| 16-8 | SerDesn Reset Control Register (SRDSRSTCTL) | 16-13 |
| 17-1 | USB Interface Block Diagram | 17-2 |
| 17-2 | Capability Registers Length (CAPLENGTH)..... | 17-9 |
| 17-3 | Host Controller Interface Version (HCIVERSION) | 17-9 |
| 17-4 | Host Controller Structural Parameters (HCCPARAMS)..... | 17-10 |
| 17-5 | Host Control Capability Parameters (HCCPARAMS) | 17-10 |
| 17-6 | Device Interface Version (DCIVERSION) | 17-11 |
| 17-7 | Device Control Capability Parameters (DCCPARAMS)..... | 17-12 |
| 17-8 | USB Command Register (USBCMD) | 17-13 |
| 17-9 | USB Status Register (USBSTS)..... | 17-15 |
| 17-10 | USB Interrupt Enable (USBINTR)..... | 17-18 |
| 17-11 | USB Frame Index (FRINDEX)..... | 17-19 |
| 17-12 | Periodic Frame List Base Address (PERIODICLISTBASE) | 17-21 |
| 17-13 | Device Address (DEVICEADDR)..... | 17-21 |
| 17-14 | Current Asynchronous List Address (ASYNCLISTADDR) | 17-22 |
| 17-15 | Endpoint List Address (ENDPOINTLISTADDR)..... | 17-23 |
| 17-16 | Master Interface Data Burst Size (BURSTSIZE) | 17-23 |
| 17-17 | Transmit FIFO Tuning Controls (TXFILLTUNING) | 17-24 |
| 17-18 | ULPI Register Access (ULPI VIEWPORT) | 17-26 |
| 17-19 | Configure Flag Register (CONFIGFLAG) | 17-27 |
| 17-20 | Port Status and Control (PORTSC)..... | 17-28 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 17-21 | OTG Status Control (OTGSC)..... | 17-33 |
| 17-22 | USB Mode (USBMODE) | 17-35 |
| 17-23 | Endpoint Setup Status (ENDPTSETUPSTAT) | 17-36 |
| 17-24 | Endpoint Initialization (ENDPTPRIME)..... | 17-36 |
| 17-25 | Endpoint Flush (ENDPTFLUSH) | 17-37 |
| 17-26 | Endpoint Status (ENDPTSTATUS)..... | 17-38 |
| 17-27 | Endpoint Complete (ENDPTCOMPLETE)..... | 17-39 |
| 17-28 | Endpoint Control 0 (ENDPTCTRL0) | 17-39 |
| 17-29 | Endpoint Control 1 to 5 (ENDPTCTRL _n)..... | 17-40 |
| 17-30 | Snoop 1 and Snoop 2 (SNOOP _n)..... | 17-42 |
| 17-31 | Age Count Threshold (AGE_CNT_THRESH)..... | 17-44 |
| 17-32 | Priority Control (PRI_CTRL) | 17-45 |
| 17-33 | System Interface Control Register (SI_CTRL)..... | 17-45 |
| 17-34 | USB General-Purpose Register (CONTROL) | 17-46 |
| 17-35 | Periodic Schedule Organization..... | 17-50 |
| 17-36 | Frame List Link Pointer Format..... | 17-50 |
| 17-37 | Asynchronous Schedule Organization | 17-51 |
| 17-38 | Isochronous Transaction Descriptor (iT _D) | 17-52 |
| 17-39 | Split-Transaction Isochronous Transaction Descriptor (siT _D) | 17-55 |
| 17-40 | Queue Element Transfer Descriptor (qT _D)..... | 17-59 |
| 17-41 | Queue Head Layout | 17-65 |
| 17-42 | Frame Span Traversal Node Structure | 17-69 |
| 17-43 | Derivation of Pointer into Frame List Array..... | 17-76 |
| 17-44 | General Format of Asynchronous Schedule List | 17-76 |
| 17-45 | Frame Boundary Relationship Between HS Bus and FS/LS Bus..... | 17-77 |
| 17-46 | Relationship of Periodic Schedule Frame Boundaries to Bus Frame Boundaries | 17-78 |
| 17-47 | Example Periodic Schedule | 17-80 |
| 17-48 | Example Association of iT _D s to Client Request Buffer | 17-83 |
| 17-49 | Generic Queue Head Unlink Scenario | 17-88 |
| 17-50 | Asynchronous Schedule List with Annotation to Mark Head of List..... | 17-89 |
| 17-51 | Example Mapping of qT _D Buffer Pointers to Buffer Pages | 17-91 |
| 17-52 | Host Controller Asynchronous Schedule Split-Transaction State Machine | 17-94 |
| 17-53 | Split Transaction, Interrupt Scheduling Boundary Conditions | 17-97 |
| 17-54 | General Structure of EHCI Periodic Schedule Utilizing Interrupt Spreading | 17-98 |
| 17-55 | Example Host Controller Traversal of Recovery Path via FSTNs..... | 17-100 |
| 17-56 | Split Transaction State Machine for Interrupt..... | 17-103 |
| 17-57 | Split Transaction, Isochronous Scheduling Boundary Conditions | 17-110 |
| 17-58 | siT _D Scheduling Boundary Examples | 17-112 |
| 17-59 | Split Transaction State Machine for Isochronous | 17-115 |
| 17-60 | Endpoint Queue Head Organization | 17-128 |
| 17-61 | Endpoint Queue Head Layout..... | 17-129 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 17-62 | Endpoint Transfer Descriptor (dTD)..... | 17-131 |
| 17-63 | USB 2.0 Device States | 17-136 |
| 17-64 | Endpoint Queue Head Diagram | 17-149 |
| 17-65 | Software Link Pointers..... | 17-151 |
| 17-66 | ULPI Timing | 17-162 |
| 17-67 | Sending of RX CMD..... | 17-163 |
| 17-68 | ULPI Data Transmit (NOPID)..... | 17-163 |
| 17-69 | ULPI Data Transmit (PID)..... | 17-164 |
| 17-70 | ULPI Data Receive | 17-164 |
| 17-71 | ULPI Register Write..... | 17-165 |
| 17-72 | ULPI Register Read | 17-165 |
| 18-1 | SEC Functional Modules | 18-3 |
| 18-2 | Descriptor Format | 18-17 |
| 18-3 | Header Dword | 18-18 |
| 18-4 | Pointer Dword | 18-22 |
| 18-5 | Link Table Entry | 18-23 |
| 18-6 | Descriptors, Link Tables, and Parcels | 18-25 |
| 18-7 | Fetch FIFO Enqueue Counter | 18-32 |
| 18-8 | Descriptor Finished Counter | 18-33 |
| 18-9 | Data Bytes In Counter..... | 18-33 |
| 18-10 | Data Bytes Out Counter | 18-33 |
| 18-11 | Channel Configuration Register (CCR)..... | 18-34 |
| 18-12 | Channel Status Register (CSR)..... | 18-36 |
| 18-13 | Current Descriptor Pointer Register..... | 18-38 |
| 18-14 | Fetch FIFO Enqueue Register (FFER)..... | 18-39 |
| 18-15 | Gather/Scatter Link Table Entry Format and Memory Ranges | 18-40 |
| 18-16 | EU Assignment Status Register (EUASR) | 18-45 |
| 18-17 | Interrupt Enable Register (IER)..... | 18-46 |
| 18-18 | Interrupt Status Register (ISR)..... | 18-48 |
| 18-19 | Interrupt Clear Register..... | 18-49 |
| 18-20 | ID Register | 18-49 |
| 18-21 | IP Block Revision Register | 18-50 |
| 18-22 | Master Control Register (MCR) | 18-51 |
| 18-23 | AESU Mode Register..... | 18-54 |
| 18-24 | AESU Key Size Register | 18-57 |
| 18-25 | AESU Data Size Register | 18-57 |
| 18-26 | AESU Reset Control Register..... | 18-58 |
| 18-27 | AESU Status Register | 18-59 |
| 18-28 | AESU Interrupt Status Register | 18-60 |
| 18-29 | AESU Interrupt Mask Register | 18-62 |
| 18-30 | AESU ICV Size Register | 18-63 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 18-31 | AESU End of Message Register | 18-64 |
| 18-32 | AESU CCM Context Registers | 18-67 |
| 18-33 | CRCU Mode Register | 18-80 |
| 18-34 | CRCU Key Size Register | 18-81 |
| 18-35 | CRCU Data Size Register | 18-81 |
| 18-36 | CRCU Reset Control Register | 18-82 |
| 18-37 | CRCU Control Register | 18-83 |
| 18-38 | CRCU Status Register | 18-83 |
| 18-39 | CRCU Interrupt Status Register | 18-84 |
| 18-40 | CRCU Interrupt Mask Register | 18-86 |
| 18-41 | CRCU Context Register (Write) | 18-88 |
| 18-42 | CRCU Context Register (Read-Default Mode) | 18-88 |
| 18-43 | CRCU Context Register (Read-Raw Mode) | 18-88 |
| 18-44 | CRCU Key Register | 18-89 |
| 18-45 | DEU Key Size Register | 18-90 |
| 18-46 | DEU Data Size Register | 18-91 |
| 18-47 | DEU Reset Control Register | 18-91 |
| 18-48 | DEU Status Register | 18-92 |
| 18-49 | DEU Interrupt Status Register | 18-93 |
| 18-50 | DEU Interrupt Mask Register | 18-94 |
| 18-51 | DEU End_of_Message Register | 18-96 |
| 18-52 | MDEU Mode Register in Old Configuration (NEW = 0) | 18-98 |
| 18-53 | MDEU Mode Register in New Configuration (NEW = 1) | 18-99 |
| 18-54 | MDEU Key Size Register | 18-102 |
| 18-55 | MDEU Data Size Register | 18-102 |
| 18-56 | MDEU Reset Control Register | 18-103 |
| 18-57 | MDEU Status Register | 18-103 |
| 18-58 | MDEU Interrupt Status Register | 18-105 |
| 18-59 | MDEU Interrupt Mask Register | 18-106 |
| 18-60 | MDEU ICV Size Register | 18-107 |
| 18-61 | MDEU End_of_message Register | 18-108 |
| 18-62 | MDEU Context Register | 18-109 |
| 18-63 | PKEU Mode Register | 18-112 |
| 18-64 | PKEU Key Size Register | 18-112 |
| 18-65 | PKEU AB Size Register | 18-114 |
| 18-66 | PKEU Data Size Register | 18-114 |
| 18-67 | PKEU Reset Control Register | 18-114 |
| 18-68 | PKEU Status Register | 18-115 |
| 18-69 | PKEU Interrupt Status Register | 18-116 |
| 18-70 | RNGU Mode Register | 18-120 |
| 18-71 | RNGU Data Size Register | 18-120 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 18-72 | RNGU Reset Control Register | 18-121 |
| 18-73 | RNGU Status Register | 18-122 |
| 18-74 | RNGU Interrupt Status Register | 18-123 |
| 18-75 | RNGU Interrupt Mask Register | 18-124 |
| 18-76 | RNGU End_Of_Message Register | 18-124 |
| 19-1 | eTSEC Block Diagram..... | 19-2 |
| 19-2 | TSEC_ID Register | 19-22 |
| 19-3 | TSEC_ID2 Register | 19-23 |
| 19-4 | IEVENT Register Definition | 19-25 |
| 19-5 | IMASK Register Definition | 19-28 |
| 19-6 | EDIS Register Definition | 19-30 |
| 19-7 | ECNTRL Register Definition | 19-31 |
| 19-8 | PTV Register Definition..... | 19-34 |
| 19-9 | DMACTRL Register..... | 19-34 |
| 19-10 | TBIPA Register Definition..... | 19-36 |
| 19-11 | TCTRL Register Definition | 19-36 |
| 19-12 | TSTAT Register Definition | 19-38 |
| 19-13 | DFVLAN Register Definition..... | 19-42 |
| 19-14 | TXIC Register Definition..... | 19-43 |
| 19-15 | TQUEUE Register Definition | 19-44 |
| 19-16 | TR03WT Register Definition..... | 19-45 |
| 19-17 | TR47WT Register Definition..... | 19-45 |
| 19-18 | TBPTR0–TBPTR7 Register Definition | 19-46 |
| 19-19 | TBASE Register Definition | 19-47 |
| 19-20 | TMR_TXTS _n _ID Register Definition | 19-47 |
| 19-21 | TMR_TXTS _n _H/L Register Definition | 19-48 |
| 19-22 | RCTRL Register Definition | 19-48 |
| 19-23 | RSTAT Register Definition | 19-51 |
| 19-24 | RXIC Register Definition | 19-52 |
| 19-25 | RQUEUE Register Definition..... | 19-53 |
| 19-26 | RBIFX Register Definition | 19-54 |
| 19-27 | Receive Queue Filer Table Address Register Definition | 19-56 |
| 19-28 | Receive Queue Filer Table Control Register Definition | 19-56 |
| 19-29 | Receive Queue Filer Table Property IDs 0, 2–15 Register Definition..... | 19-58 |
| 19-30 | Receive Queue Filer Table Property ID1 Register Definition | 19-58 |
| 19-31 | MRBLR Register Definition | 19-61 |
| 19-32 | RBPTR0–RBPTR7 Register Definition | 19-62 |
| 19-33 | RBASE Register Definition | 19-62 |
| 19-34 | TMR_RXTS_H/L Register Definition..... | 19-63 |
| 19-35 | MACCFG1 Register Definition | 19-66 |
| 19-36 | MACCFG2 Register Definition | 19-67 |

Figures

| Figure Number | Title | Page Number |
|------------------|---|----------------|
| 19-37 | IPGIFG Register Definition | 19-69 |
| 19-38 | Half-Duplex Register Definition | 19-70 |
| 19-39 | Maximum Frame Length Register Definition | 19-71 |
| 19-40 | MII Management Configuration Register Definition | 19-72 |
| 19-41 | MIIMCOM Register Definition | 19-72 |
| 19-42 | MIIMADD Register Definition | 19-73 |
| 19-43 | MII Mgmt Control Register Definition | 19-74 |
| 19-44 | MIIMSTAT Register Definition | 19-74 |
| 19-45 | MII Mgmt Indicator Register Definition | 19-75 |
| 19-46 | Interface Status Register Definition | 19-75 |
| 19-47 | MAC Station Address Part 1 Register Definition | 19-76 |
| 19-48 | MAC Station Address Part 2 Register Definition | 19-77 |
| 19-49 | MAC Exact Match Address <i>n</i> Part 1 Register Definition | 19-77 |
| 19-50 | MAC Exact Match Address <i>x</i> Part 2 Register Definition | 19-78 |
| 19-51 | Transmit and Receive 64-Byte Frame Register Definition | 19-79 |
| 19-52 | Transmit and Receive 65- to 127-Byte Frame Register Definition | 19-80 |
| 19-53 | Transmit and Received 128- to 255-Byte Frame Register Definition | 19-80 |
| 19-54 | Transmit and Received 256- to 511-Byte Frame Register Definition | 19-81 |
| 19-55 | Transmit and Received 512- to 1023-Byte Frame Register Definition | 19-81 |
| 19-56 | Transmit and Received 1024- to 1518-Byte Frame Register Definition | 19-82 |
| 19-57 | Transmit and Received 1519- to 1522-Byte VLAN Frame Register Definition | 19-82 |
| 19-58 | Receive Byte Counter Register Definition | 19-83 |
| 19-59 | Receive Packet Counter Register Definition | 19-83 |
| 19-60 | Receive FCS Error Counter Register Definition | 19-83 |
| 19-61 | Receive Multicast Packet Counter Register Definition | 19-84 |
| 19-62 | Receive Broadcast Packet Counter Register Definition | 19-84 |
| 19-63 | Receive Control Frame Packet Counter Register Definition | 19-85 |
| 19-64 | Receive Pause Frame Packet Counter Register Definition | 19-85 |
| 19-65 | Receive Unknown OPCode Packet Counter Register Definition | 19-86 |
| 19-66 | Receive Alignment Error Counter Register Definition | 19-86 |
| 19-67 | Receive Frame Length Error Counter Register Definition | 19-87 |
| 19-68 | Receive Code Error Counter Register Definition | 19-87 |
| 19-69 | Receive Carrier Sense Error Counter Register Definition | 19-88 |
| 19-70 | Receive Undersize Packet Counter Register Definition | 19-88 |
| 19-71 | Receive Oversize Packet Counter Register Definition | 19-89 |
| 19-72 | Receive Fragments Counter Register Definition | 19-89 |
| 19-73 | Receive Jabber Counter Register Definition | 19-90 |
| 19-74 | Receive Dropped Packet Counter Register Definition | 19-90 |
| 19-75 | Transmit Byte Counter Register Definition | 19-91 |
| 19-76 | Transmit Packet Counter Register Definition | 19-91 |
| 19-77 | Transmit Multicast Packet Counter Register Definition | 19-92 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 19-78 | Transmit Broadcast Packet Counter Register Definition | 19-92 |
| 19-79 | Transmit Pause Control Frame Counter Register Definition | 19-93 |
| 19-80 | Transmit Deferral Packet Counter Register Definition | 19-93 |
| 19-81 | Transmit Excessive Deferral Packet Counter Register Definition | 19-94 |
| 19-82 | Transmit Single Collision Packet Counter Register Definition | 19-94 |
| 19-83 | Transmit Multiple Collision Packet Counter Register Definition | 19-95 |
| 19-84 | Transmit Late Collision Packet Counter Register Definition | 19-95 |
| 19-85 | Transmit Excessive Collision Packet Counter Register Definition | 19-96 |
| 19-86 | Transmit Total Collision Counter Register Definition | 19-96 |
| 19-87 | Transmit Drop Frame Counter Register Definition | 19-97 |
| 19-88 | Transmit Jabber Frame Counter Register Definition | 19-97 |
| 19-89 | Transmit FCS Error Counter Register Definition | 19-98 |
| 19-90 | Transmit Control Frame Counter Register Definition | 19-98 |
| 19-91 | Transmit Oversized Frame Counter Register Definition | 19-99 |
| 19-92 | Transmit Undersize Frame Counter Register Definition | 19-99 |
| 19-93 | Transmit Fragment Counter Register Definition | 19-100 |
| 19-94 | Carry Register 1 (CAR1) Register Definition | 19-100 |
| 19-95 | Carry Register 2 (CAR2) Register Definition | 19-102 |
| 19-96 | Carry Mask Register 1 (CAM1) Register Definition | 19-103 |
| 19-97 | Carry Mask Register 2 (CAM2) Register Definition | 19-104 |
| 19-98 | Receive Filer Rejected Packet Counter Register Definition | 19-105 |
| 19-99 | IGADDR _n Register Definition | 19-106 |
| 19-100 | GADDR _n Register Definition | 19-107 |
| 19-101 | ATTR Register Definition | 19-107 |
| 19-102 | RQPRM Register Definition | 19-108 |
| 19-103 | RFBPTR0–RFBPTR7 Register Definition | 19-109 |
| 19-104 | TMR_CTRL Register Definition | 19-110 |
| 19-105 | TMR_TEVENT Register Definition | 19-112 |
| 19-106 | TMR_TEMASK Register Definition | 19-113 |
| 19-107 | TMR_PEVENT Register Definition | 19-114 |
| 19-108 | TMR_PEMASK Register Definition | 19-114 |
| 19-109 | TMR_STAT Register Definition | 19-115 |
| 19-110 | TMR_CNT_H Register Definition | 19-116 |
| 19-111 | TMR_ADD Register Definition | 19-116 |
| 19-112 | TMR_ACC Register Definition | 19-117 |
| 19-113 | TMR_PRSC Register Definition | 19-117 |
| 19-114 | TMROFF_H/L Register Definition | 19-118 |
| 19-115 | TMR_ALARM1-2_H/L Register Definition | 19-118 |
| 19-116 | TMR_FIPER _n Register Definition | 19-120 |
| 19-117 | TMR_ETTS1-2_H/L Register Definition | 19-120 |
| 19-118 | Control Register Definition | 19-123 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 19-119 | Status Register Definition | 19-124 |
| 19-120 | AN Advertisement Register Definition..... | 19-125 |
| 19-121 | AN Link Partner Base Page Ability Register Definition | 19-127 |
| 19-122 | AN Expansion Register Definition | 19-128 |
| 19-123 | AN Next Page Transmit Register Definition | 19-128 |
| 19-124 | AN Link Partner Ability Next Page Register Definition | 19-129 |
| 19-125 | Extended Status Register Definition | 19-130 |
| 19-126 | Jitter Diagnostics Register Definition | 19-131 |
| 19-127 | TBI Control Register Definition | 19-132 |
| 19-128 | eTSEC-MII Connection | 19-134 |
| 19-129 | eTSEC-RMII Connection | 19-135 |
| 19-130 | eTSEC-RGMII Connection..... | 19-136 |
| 19-131 | eTSEC-RTBI Connection | 19-137 |
| 19-132 | eTSEC-SGMII Connection | 19-138 |
| 19-133 | Definition of Custom Preamble Sequence | 19-146 |
| 19-134 | Definition of Received Preamble Sequence..... | 19-146 |
| 19-135 | Ethernet Address Recognition Flowchart | 19-148 |
| 19-136 | Sample C Code for Computing eTSEC Hash Table Indices..... | 19-150 |
| 19-137 | Location of Frame Control Blocks for TOE Parameters | 19-158 |
| 19-138 | Transmit Frame Control Block | 19-159 |
| 19-139 | Receive Frame Control Block..... | 19-160 |
| 19-140 | Structure of the Receive Queue Filer Table | 19-165 |
| 19-141 | 1588 Timer Design Partition | 19-178 |
| 19-142 | Ethernet Sampling Points for 1588 | 19-178 |
| 19-143 | PTP Packet Format..... | 19-180 |
| 19-144 | Buffer Format for Transmit Time-Stamp Insertion..... | 19-181 |
| 19-145 | Transmit Frame Control Block | 19-182 |
| 19-146 | Example of eTSEC Memory Structure for BDs | 19-184 |
| 19-147 | Buffer Descriptor Ring..... | 19-185 |
| 19-148 | Transmit Buffer Descriptor | 19-185 |
| 19-149 | Mapping of TxBDs to a C Data Structure..... | 19-186 |
| 19-150 | Receive Buffer Descriptor..... | 19-188 |
| 19-151 | Mapping of RxBDs to a C Data Structure | 19-189 |
| 20-1 | I ² C Block Diagram..... | 20-1 |
| 20-2 | I ² C Address Register (I2CADR)..... | 20-5 |
| 20-3 | I ² C Frequency Divider Register (I2CFDR) | 20-5 |
| 20-4 | I ² C Control Register (I2CCR)..... | 20-6 |
| 20-5 | I ² C Status Register (I2CSR) | 20-7 |
| 20-6 | I ² C Data Register (I2CDR)..... | 20-9 |
| 20-7 | I ² C Digital Filter Sampling Rate Register (I2CDFSRR)..... | 20-9 |
| 20-8 | I ² C Interface Transaction Protocol..... | 20-10 |

Figures

| Figure Number | Title | Page Number |
|---------------|---|-------------|
| 20-9 | EEPROM Contents | 20-17 |
| 20-10 | EEPROM Data Format for One Register Preload Command..... | 20-18 |
| 20-11 | Example I ² C Interrupt Service Routine Flowchart..... | 20-20 |
| 21-1 | UART Block Diagram | 21-2 |
| 21-2 | Receiver Buffer Registers (URBR1 and URBR2)..... | 21-6 |
| 21-3 | Transmitter Holding Registers (UTHR1 and UTHR2)..... | 21-6 |
| 21-4 | Divisor Most Significant Byte Registers (UDMB1 and UDMB2)..... | 21-7 |
| 21-5 | Divisor Least Significant Byte Registers (UDLB1 and UDLB2)..... | 21-7 |
| 21-6 | Interrupt Enable Registers (UIER1 and UIER2)..... | 21-8 |
| 21-7 | Interrupt ID Registers (UIIR1 and UIIR2)..... | 21-9 |
| 21-8 | FIFO Control Registers (UFCR1 and UFCR2)..... | 21-11 |
| 21-9 | Alternate Function Register (UAFR)..... | 21-11 |
| 21-10 | Line Control Register (ULCR1 and ULCR2)..... | 21-12 |
| 21-11 | Modem Control Register (UMCR1 and UMCR2)..... | 21-14 |
| 21-12 | Line Status Register (ULSR1 and ULSR2)..... | 21-15 |
| 21-13 | Modem Status Register (UMSR1 and UMSR2)..... | 21-16 |
| 21-14 | Scratch Register (USCR)..... | 21-17 |
| 21-15 | DMA Status Register (UDSR)..... | 21-17 |
| 21-16 | UART Bus Interface Transaction Protocol Example..... | 21-19 |
| 22-1 | SPI Block Diagram | 22-2 |
| 22-2 | Single-Master/Multi-Slave Configuration | 22-4 |
| 22-3 | Multiple-Master Configuration | 22-6 |
| 22-4 | SPMODE-SPI Mode Register Definition | 22-9 |
| 22-5 | SPI Transfer Format with SPMODE[CP] = 0..... | 22-11 |
| 22-6 | SPI Transfer Format with SPMODE[CP] = 1..... | 22-11 |
| 22-7 | SPIE—SPI Event Register Definition..... | 22-12 |
| 22-8 | SPIM—SPI Mask Register Definition..... | 22-13 |
| 22-9 | SPI Command Register Definition | 22-14 |
| 22-10 | SPI Transmit Data Hold Register Definition | 22-14 |
| 22-11 | SPI Receive Data Hold Register Definition..... | 22-15 |
| 22-12 | Example SPMODE[REV] = 0 SPMODE[LEN] = 7 LSB Sent First..... | 22-15 |
| 22-13 | Example SPMODE[REV] = 1 SPMODE[LEN] = 7 MSB Sent First..... | 22-15 |
| 22-14 | Example SPMODE[REV] = 1 SPMODE[LEN] = 15 MSB Sent First..... | 22-16 |
| 22-15 | Example SPMODE[REV] = 0 SPMODE[LEN] = 15 LSB Sent First..... | 22-16 |
| 23-1 | JTAG Interface Block Diagram | 23-1 |
| 24-1 | GPIO Module Block Diagram | 24-1 |
| 24-2 | GPIO Direction Register (GPDIR) | 24-3 |
| 24-3 | GPIO Open Drain Register (GPODR)..... | 24-3 |
| 24-4 | GPIO Data Register (GPDAT)..... | 24-4 |
| 24-5 | GPIO Interrupt Event Register (GPIER) | 24-4 |
| 24-6 | GPIO Interrupt Mask Register (GPIMR)..... | 24-5 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 24-7 | GPIO Interrupt Control Register (GPICR) | 24-5 |
| 25-1 | TDM Block Diagram | 25-4 |
| 25-2 | TDM Frames | 25-5 |
| 25-3 | T1 Frame | 25-6 |
| 25-4 | TDM Modules Shared Mode | 25-6 |
| 25-5 | TDM General Interface Register..... | 25-8 |
| 25-6 | TDM Receive Interface Register | 25-9 |
| 25-7 | TDM Transmit Interface Register | 25-12 |
| 25-8 | TDM Transmit Interface Register | 25-14 |
| 25-9 | TDM Transmit Frame Parameters | 25-16 |
| 25-10 | TDM Receive Channel Enable <i>n</i> | 25-17 |
| 25-11 | TDM Transmit Channel Enable <i>n</i> | 25-18 |
| 25-12 | TDM Transmit Channel Mask <i>n</i> | 25-18 |
| 25-13 | TDM Receive Control Register | 25-19 |
| 25-14 | TDM Transmit Control Register..... | 25-19 |
| 25-15 | TDM Receive Interrupt Enable Register | 25-20 |
| 25-16 | TDM Transmit Interrupt Enable Register | 25-21 |
| 25-17 | TDM Receive Event Register | 25-23 |
| 25-18 | TDM Transmit Event Register..... | 25-24 |
| 25-19 | TDM Receive Status Register | 25-26 |
| 25-20 | TDM Transmit Status Register | 25-27 |
| 25-21 | TDM Receive Data Register | 25-28 |
| 25-22 | TDM Transmit Data Register..... | 25-29 |
| 25-23 | TDM Clocking (8-Bit Words, 4 Time Slots/Frame) | 25-30 |
| 25-24 | TDM Clock Generation | 25-30 |
| 25-25 | TDM Transmit Clock Generator Block Diagram | 25-31 |
| 25-26 | TDM Transmit Frame Sync Generator Block Diagram..... | 25-31 |
| 25-27 | TDM Point-to-Point Configuration..... | 25-32 |
| 25-28 | TDM Network Configuration | 25-32 |
| 25-29 | TDM Shared Network Configuration | 25-33 |
| 25-30 | Sync Length Selection | 25-34 |
| 25-31 | Frame Sync Configurations without Sync Delays | 25-35 |
| 25-32 | Frame Sync Configurations with Sync Delays | 25-36 |
| 25-33 | Frame Sync Configuration At T1 mode..... | 25-37 |
| 25-34 | Frame Sync Polarity | 25-37 |
| 25-35 | Frame Sync Synchronization State Diagram | 25-38 |
| 25-36 | Reserve Bit Order..... | 25-39 |
| 25-37 | TDM Module Modes | 25-40 |
| 25-38 | Receive and Transmit Totally Independent..... | 25-40 |
| 25-39 | TDM Network Mode Transmit Timing with Mask Register | 25-41 |
| 25-40 | TDM Network Mode Receive Timing with Enable Register | 25-43 |

Figures

| Figure Number | Title | Page Number |
|---------------|--|-------------|
| 25-41 | TDM Rx FIFO Using RWEN Bit with Ten 8-Bit Channels | 25-45 |
| 25-42 | TDM Tx FIFO Using TWEN Bit with Ten 8-Bit Channels..... | 25-45 |
| 25-43 | Division Value 1 | 25-51 |
| 25-44 | Division Value Even..... | 25-51 |
| 25-45 | Division Value Odd..... | 25-52 |
| 25-46 | TDMCLK_DIV_VAL_RX Register | 25-52 |
| 25-47 | TDMCLK_DIV_VAL_TX Register | 25-53 |
| 25-48 | DMA Block Diagram..... | 25-54 |
| 25-49 | DMA Control Register (DMACR) | 25-58 |
| 25-50 | DMA Error Status Register (DMAES) | 25-61 |
| 25-51 | DMA Enable Request Register (DMAERQ)..... | 25-63 |
| 25-52 | DMA Enable Error Interrupt Register (DMAEEI) | 25-63 |
| 25-53 | DMA Set Enable Request Register | 25-64 |
| 25-54 | DMA Clear Enable Request Register | 25-65 |
| 25-55 | DMA Set Enable Error Interrupt Register | 25-65 |
| 25-56 | DMA Clear Enable Error Interrupt Register..... | 25-66 |
| 25-57 | DMA Clear Interrupt Request Register | 25-66 |
| 25-58 | DMA Clear Error Register..... | 25-67 |
| 25-59 | DMA Set START Bit Register..... | 25-68 |
| 25-60 | DMA Clear DONE Status Register..... | 25-68 |
| 25-61 | DMA Interrupt Request Register Low (DMAINT) | 25-69 |
| 25-62 | DMA Error Register (DMAERR)..... | 25-70 |
| 25-63 | DMA Hardware Request Status Register (DMAHRS)..... | 25-70 |
| 25-64 | DMA Hardware Request Status Register—High..... | 25-71 |
| 25-65 | DMA Clear DONE Status Register..... | 25-72 |
| 25-66 | TCD Word 0 (TCD.saddr) Field | 25-73 |
| 25-67 | TCD Word 1 (TCDn.{soff, smod, ssize, dmod, dsize}) Fields..... | 25-73 |
| 25-68 | TCD Word 2 (TCD.{smloe, dmloe, nbytes}) Field | 25-74 |
| 25-69 | TCD Word 3 (TCD.slast) Field..... | 25-75 |
| 25-70 | TCD Word 4 (TCDn.daddr) Field..... | 25-76 |
| 25-71 | TCD Word 5 (TCD.{citer, doff}) Fields | 25-76 |
| 25-72 | TCD Word 6 (TCD.dlast_sga) Field | 25-77 |
| 25-73 | TCD Word 7 (TCD.{biter, control/status}) Fields | 25-78 |
| 25-74 | DMA Operation—Part 1 | 25-82 |
| 25-75 | DMA Operation—Part 2 | 25-83 |
| 25-76 | DMA Operation—Part 3 | 25-84 |
| 25-77 | ipd_req Removal | 25-89 |

Tables

| Table Number | Title | Page Number |
|--------------|--|-------------|
| 2-1 | MPC8315E Signal Reference by Functional Block..... | 2-4 |
| 2-2 | MPC8315E Signal Reference by Alphabetical Order..... | 2-18 |
| 2-3 | Output Signal States During System Reset..... | 2-31 |
| 2-4 | Signals for Multiplexing | 2-32 |
| 3-1 | IMMR Memory Map | 3-2 |
| 4-1 | System Control Signals..... | 4-1 |
| 4-2 | External Clock Signals..... | 4-3 |
| 4-3 | Reset Causes | 4-5 |
| 4-4 | Reset Actions | 4-6 |
| 4-5 | Reset Configuration Words Source..... | 4-10 |
| 4-6 | SYS_CLK_IN Division | 4-11 |
| 4-7 | Selecting Reset Configuration Input Signals | 4-11 |
| 4-8 | RCWLR Bit Settings..... | 4-13 |
| 4-9 | System PLL VCO Division..... | 4-14 |
| 4-10 | System PLL Ratio | 4-14 |
| 4-11 | SPMF Maximum Values | 4-14 |
| 4-12 | Reset Configuration Word High Bit Settings..... | 4-15 |
| 4-13 | PCI Host/Agent Configuration..... | 4-17 |
| 4-14 | Boot Memory Space..... | 4-17 |
| 4-15 | Boot Sequencer Configuration..... | 4-18 |
| 4-16 | Boot ROM Location..... | 4-19 |
| 4-17 | eTSEC1 Mode Configuration | 4-20 |
| 4-18 | eTSEC2 Mode Configuration | 4-21 |
| 4-19 | e300 Core True Little-Endian | 4-21 |
| 4-20 | LALE Configuration | 4-22 |
| 4-21 | Local Bus Configuration EEPROM Addresses | 4-22 |
| 4-22 | Local Bus Reset Configuration Words Data Structure..... | 4-22 |
| 4-23 | Local Bus Controller Setting when Loading RCW | 4-23 |
| 4-24 | Hard Coded Reset Configuration Word Low Fields Values | 4-27 |
| 4-25 | Hard-Coded Reset Configuration Word High Field Values..... | 4-28 |
| 4-26 | Examples For Hard-Coded Reset Configuration Words Usage..... | 4-28 |
| 4-27 | Configurable Clock Units | 4-32 |
| 4-28 | Reset Configuration and Status Registers Memory Map..... | 4-33 |
| 4-29 | Reset Status Register Field Descriptions | 4-34 |
| 4-30 | RMR Field Descriptions | 4-36 |
| 4-31 | RPR Bit Descriptions | 4-37 |
| 4-32 | RCR Bit Settings..... | 4-37 |
| 4-33 | RCER Bit Settings | 4-38 |
| 4-34 | Clock Configuration Registers Memory Map..... | 4-38 |
| 4-35 | System PLL Mode Register Bit Settings | 4-39 |
| 4-36 | OCCR Bit Settings..... | 4-40 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 4-37 | SCCR Bit Descriptions | 4-41 |
| 5-1 | Local Access Windows Target Interface | 5-1 |
| 5-2 | Local Access Windows Example | 5-2 |
| 5-3 | Format of Window Definitions | 5-3 |
| 5-4 | Local Access Register Memory Map | 5-5 |
| 5-5 | IMMRBAR Bit Settings | 5-7 |
| 5-6 | ALTCBAR Bit Settings | 5-8 |
| 5-7 | LBLAWBAR0–LBLAWBAR3 Bit Settings | 5-8 |
| 5-8 | LBLAWBAR0[BASE_ADDR] Reset Value | 5-9 |
| 5-9 | LBLAWAR0–LBLAWAR3 Bit Settings | 5-9 |
| 5-10 | LBLAWAR0[EN] Reset Value | 5-10 |
| 5-11 | PCILAWBAR0–PCILAWBAR1 Bit Settings | 5-10 |
| 5-12 | PCILAWBAR0[BASE_ADDR] Reset Value | 5-10 |
| 5-13 | PCILAWAR0–PCILAWAR1 Bit Settings | 5-11 |
| 5-14 | PCILAWAR0[EN] Reset Value | 5-11 |
| 5-15 | PCIEXP1LAWBAR Bit Settings | 5-12 |
| 5-16 | PCIEXP1LAWAR Bit Settings | 5-13 |
| 5-17 | PCIEXP2LAWBAR Bit Settings | 5-13 |
| 5-18 | PCIEXP2LAWAR Bit Settings | 5-14 |
| 5-19 | DDRLAWBAR0–DDRLAWBAR1 Bit Settings | 5-15 |
| 5-20 | DDRLAWBAR0[BASE_ADDR] Reset Value | 5-15 |
| 5-21 | DDRLAWAR0–DDRLAWAR1 Bit Settings | 5-16 |
| 5-22 | DDRLAWAR0[EN] Reset Value | 5-16 |
| 5-23 | Overlapping Local Access Windows | 5-16 |
| 5-24 | System Configuration Register Memory Map | 5-19 |
| 5-25 | SGPRL Bit Settings | 5-20 |
| 5-26 | SGPRH Bit Settings | 5-20 |
| 5-27 | SPRIDR Bit Settings | 5-21 |
| 5-28 | PARTID Coding | 5-21 |
| 5-29 | REVID Coding | 5-21 |
| 5-30 | SPCR Bit Settings | 5-22 |
| 5-31 | SICRL Bit Settings | 5-24 |
| 5-32 | SICRH Bit Settings | 5-27 |
| 5-33 | SICRH[27–31] Bit Settings | 5-28 |
| 5-34 | DDRCDR Field Descriptions | 5-30 |
| 5-35 | DDRDSR Field Descriptions | 5-31 |
| 5-36 | PECR Field Description | 5-32 |
| 5-37 | WDT Register Address Map | 5-34 |
| 5-38 | SWCRR Bit Settings | 5-35 |
| 5-39 | SWCNR Bit Settings | 5-36 |
| 5-40 | SWSRR Bit Settings | 5-37 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 5-41 | RTC Signal Properties..... | 5-41 |
| 5-42 | RTC External Signal—Detailed Signal Description | 5-41 |
| 5-43 | RTC Register Address Map | 5-41 |
| 5-44 | RTCNr Bit Settings..... | 5-42 |
| 5-45 | RTLDR Bit Settings | 5-43 |
| 5-46 | RTPSR Bit Settings..... | 5-43 |
| 5-47 | RTCTR Bit Settings | 5-44 |
| 5-48 | RTEVR Bit Settings | 5-44 |
| 5-49 | RTALR Bit Settings | 5-45 |
| 5-50 | PIT Signal Properties | 5-48 |
| 5-51 | PIT External Signal—Detailed Signal Descriptions | 5-48 |
| 5-52 | PIT Register Address Map..... | 5-49 |
| 5-53 | PTCNr Bit Settings | 5-49 |
| 5-54 | PTLDR Bit Settings | 5-50 |
| 5-55 | PTPSR Bit Settings | 5-50 |
| 5-56 | PTCTR Bit Settings..... | 5-51 |
| 5-57 | PTEVR Bit Settings | 5-51 |
| 5-58 | GTM Signal Properties..... | 5-56 |
| 5-59 | GTM External Signals—Detailed Signal Descriptions..... | 5-57 |
| 5-60 | GTM Register Address Map | 5-58 |
| 5-61 | GTCFR1 Bit Settings | 5-59 |
| 5-62 | GTCFR2 Bit Settings | 5-61 |
| 5-63 | GTMDR Bit Settings..... | 5-62 |
| 5-64 | GTRFR Bit Settings | 5-63 |
| 5-65 | GTCPR _n Bit Settings | 5-64 |
| 5-66 | GTCNR Bit Settings..... | 5-64 |
| 5-67 | GTEVR _n Bit Settings..... | 5-65 |
| 5-68 | GTPSR _n Bit Settings..... | 5-65 |
| 5-69 | MPC8315E Power States Supported..... | 5-70 |
| 5-70 | System Control Signals—Detailed Signal Descriptions | 5-70 |
| 5-71 | Power Management Controller Registers Memory Map | 5-71 |
| 5-72 | PMCCR Bit Settings | 5-71 |
| 5-73 | PMCER Bit Settings | 5-72 |
| 5-74 | PMCMR Bit Settings | 5-75 |
| 5-75 | PMCCR1 Bit Settings | 5-76 |
| 5-76 | PMCCR2 Bit Settings | 5-78 |
| 5-77 | MPC8315E Power Rail Nomenclature | 5-80 |
| 5-78 | Software-Controller Power-Down States—Basic Description | 5-80 |
| 5-79 | Power State Mapped to OS State | 5-81 |
| 5-80 | PCI Defined Power Management State Support..... | 5-83 |
| 5-81 | PCI Bus Power Management State Support..... | 5-83 |

Tables

| Table Number | Title | Page Number |
|--------------|--|-------------|
| 5-82 | Software-Controller Power-Down States—Basic Description | 5-84 |
| 5-83 | MPC8315E Agent Mode Wake-Up Support | 5-88 |
| 5-84 | MPC8315E Host Mode Wake-Up Support | 5-90 |
| 6-1 | Arbiter Register Map | 6-2 |
| 6-2 | ACR Field Descriptions | 6-3 |
| 6-3 | ATR Field Descriptions | 6-5 |
| 6-4 | AEER Bit Settings | 6-6 |
| 6-5 | AER Field Descriptions | 6-7 |
| 6-6 | AIDR Field Descriptions | 6-8 |
| 6-7 | AMR Field Descriptions | 6-9 |
| 6-8 | AEATR Field Descriptions | 6-10 |
| 6-9 | AEADR Field Descriptions | 6-11 |
| 6-10 | AERR Field Descriptions | 6-12 |
| 6-11 | Address Only Transaction Type Encoding | 6-16 |
| 6-12 | Reserved Transaction Type Encoding | 6-17 |
| 6-13 | Illegal Transaction Type Encoding | 6-17 |
| 7-1 | Device Revision Level Cross-Reference | 7-13 |
| 7-2 | MSR Bit Descriptions | 7-18 |
| 7-3 | e300 HID0 Bit Descriptions | 7-22 |
| 7-4 | Using HID0[ECLK] and HID0[SBCLK] to Configure <i>clk_out</i> | 7-24 |
| 7-5 | HID1 Bit Descriptions | 7-25 |
| 7-6 | e300HID2 Bit Descriptions | 7-25 |
| 7-7 | Interrupt Classifications | 7-32 |
| 7-8 | Exceptions and Interrupts | 7-33 |
| 7-9 | Differences Between e300 and G2_LE Cores | 7-40 |
| 8-1 | IPIC Signal Properties | 8-5 |
| 8-2 | IPIC External Signals—Detailed Signal Descriptions | 8-5 |
| 8-3 | IPIC Register Address Map | 8-6 |
| 8-4 | SICFR Field Descriptions | 8-8 |
| 8-5 | SIVCR Field Descriptions | 8-9 |
| 8-6 | IVEC/CVEC/MVEC Field Definition | 8-10 |
| 8-7 | SIPNR_H/SIFCR_H/SIMSR_H Bit Assignments | 8-12 |
| 8-8 | SIPNR_H Field Descriptions | 8-13 |
| 8-9 | SIPNR_L/SIFCR_L/SIMSR_L Bit Assignments | 8-14 |
| 8-10 | SIPNR_L Field Descriptions | 8-15 |
| 8-11 | SIPRR_A Field Descriptions | 8-15 |
| 8-12 | SIPRR_B Field Descriptions | 8-16 |
| 8-13 | SIPRR_C Field Descriptions | 8-17 |
| 8-14 | SIPRR_D Field Descriptions | 8-18 |
| 8-15 | SIMSR_H Field Descriptions | 8-19 |
| 8-16 | SIMSR_L Field Descriptions | 8-19 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 8-17 | SICNR Field Descriptions | 8-20 |
| 8-18 | SEPNR Field Descriptions | 8-22 |
| 8-19 | SMPRR_A Field Descriptions | 8-22 |
| 8-20 | SMPRR_B Field Descriptions | 8-23 |
| 8-21 | SEMSR Field Descriptions | 8-24 |
| 8-22 | SECNR Field Descriptions | 8-25 |
| 8-23 | SERSR/SERM/R/SERFR Bit Assignments | 8-26 |
| 8-24 | SERSR Field Descriptions | 8-26 |
| 8-25 | SERM Field Descriptions | 8-27 |
| 8-26 | SERCR Field Descriptions | 8-28 |
| 8-27 | SEPCR Field Descriptions | 8-28 |
| 8-28 | SIFCR_H Field Descriptions | 8-29 |
| 8-29 | SIFCR_L Field Descriptions | 8-29 |
| 8-30 | SEFCR Field Descriptions | 8-30 |
| 8-31 | SERFR Field Descriptions | 8-31 |
| 8-32 | SCVCR Field Descriptions | 8-31 |
| 8-33 | SMVCR Field Descriptions | 8-32 |
| 8-34 | Interrupt Source Priority Levels | 8-35 |
| 8-35 | Message Shared Registers Address Map | 8-40 |
| 8-36 | MSIRs Field Descriptions | 8-41 |
| 8-37 | MSIMR Field Descriptions | 8-42 |
| 8-38 | MSISR Field Descriptions | 8-42 |
| 8-39 | MSIIR Field Descriptions | 8-43 |
| 9-1 | DDR Memory Interface Signal Summary | 9-3 |
| 9-2 | Memory Address Signal Mappings | 9-4 |
| 9-3 | Memory Interface Signals—Detailed Signal Descriptions | 9-5 |
| 9-4 | Clock Signals—Detailed Signal Descriptions | 9-7 |
| 9-5 | DDR Memory Controller Memory Map | 9-8 |
| 9-6 | CS _n _BNDS Field Descriptions | 9-10 |
| 9-7 | CS _n _CONFIG Field Descriptions | 9-11 |
| 9-8 | TIMING_CFG_3 Field Descriptions | 9-12 |
| 9-9 | TIMING_CFG_0 Field Descriptions | 9-13 |
| 9-10 | TIMING_CFG_1 Field Descriptions | 9-14 |
| 9-11 | TIMING_CFG_2 Field Descriptions | 9-17 |
| 9-12 | DDR_SDRAM_CFG Field Descriptions | 9-19 |
| 9-13 | DDR_SDRAM_CFG_2 Field Descriptions | 9-21 |
| 9-14 | DDR_SDRAM_MODE Field Descriptions | 9-23 |
| 9-15 | DDR_SDRAM_MODE_2 Field Descriptions | 9-24 |
| 9-16 | DDR_SDRAM_MD_CNTL Field Descriptions | 9-25 |
| 9-17 | Settings of DDR_SDRAM_MD_CNTL Fields | 9-26 |
| 9-18 | DDR_SDRAM_INTERVAL Field Descriptions | 9-27 |

Tables

| Table Number | Title | Page Number |
|--------------|---|-------------|
| 9-19 | DDR_DATA_INIT Field Descriptions | 9-27 |
| 9-20 | DDR_SDRAM_CLK_CNTL Field Descriptions | 9-28 |
| 9-21 | DDR_INIT_ADDR Field Descriptions | 9-28 |
| 9-22 | DDR_IP_REV1 Field Descriptions | 9-29 |
| 9-23 | DDR_IP_REV2 Field Descriptions | 9-29 |
| 9-24 | ERR_DETECT Field Descriptions | 9-30 |
| 9-25 | ERR_DISABLE Field Descriptions..... | 9-30 |
| 9-26 | ERR_INT_EN Field Descriptions | 9-31 |
| 9-27 | Byte Lane to Data Relationship | 9-35 |
| 9-28 | Supported DDR1 SDRAM Device Configurations | 9-35 |
| 9-29 | Supported DDR2 SDRAM Device Configurations | 9-36 |
| 9-30 | DDR1 Address Multiplexing for 32-Bit Data Bus with Interleaving Disabled..... | 9-37 |
| 9-31 | DDR1 Address Multiplexing for 16-Bit Data Bus..... | 9-38 |
| 9-32 | DDR2 Address Multiplexing for 16-Bit Data Bus..... | 9-39 |
| 9-33 | DDR2 Address Multiplexing | 9-39 |
| 9-34 | Example of Address Multiplexing for 32-Bit Data Bus Interleaving between Two Banks with Partial Array Self Refresh Disabled..... | 9-40 |
| 9-35 | DDR SDRAM Command Table..... | 9-42 |
| 9-36 | DDR SDRAM Interface Timing Intervals | 9-43 |
| 9-37 | DDR SDRAM Power-Saving Modes Refresh Configuration..... | 9-50 |
| 9-38 | Memory Controller—Data Beat Ordering | 9-52 |
| 9-39 | Memory Interface Configuration Register Initialization Parameters..... | 9-53 |
| 9-40 | Programming Differences between Memory Types | 9-54 |
| 10-1 | Signal Properties—Summary..... | 10-4 |
| 10-2 | Enhanced Local Bus Controller Detailed Signal Descriptions | 10-5 |
| 10-3 | Enhanced Local Bus Controller Registers | 10-8 |
| 10-4 | BR _n Field Descriptions..... | 10-10 |
| 10-5 | Reset value of OR0 Register | 10-11 |
| 10-6 | Memory Bank Sizes in Relation to Address Mask | 10-12 |
| 10-7 | OR _n —GPCM Field Descriptions | 10-13 |
| 10-8 | OR _n —FCM Field Descriptions | 10-16 |
| 10-9 | OR _n —UPM Field Descriptions | 10-18 |
| 10-10 | MAR Field Descriptions | 10-19 |
| 10-11 | M _x MR Field Descriptions..... | 10-20 |
| 10-12 | MRTPR Field Descriptions | 10-22 |
| 10-13 | MDR Field Description..... | 10-23 |
| 10-14 | LSOR Field Description..... | 10-24 |
| 10-15 | LURT Field Descriptions | 10-25 |
| 10-16 | LTESR Field Descriptions | 10-26 |
| 10-17 | LTEDR Field Descriptions..... | 10-27 |
| 10-18 | LTEIR Field Descriptions | 10-28 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 10-19 | LTEATR Field Descriptions..... | 10-29 |
| 10-20 | LTEAR Field Descriptions..... | 10-30 |
| 10-21 | LBCR Field Descriptions..... | 10-31 |
| 10-22 | LCRR Field Descriptions..... | 10-32 |
| 10-23 | FMR Field Descriptions..... | 10-33 |
| 10-24 | FIR Field Descriptions..... | 10-35 |
| 10-25 | FCR Field Descriptions..... | 10-36 |
| 10-26 | FBAR Field Descriptions..... | 10-36 |
| 10-27 | FPAR Field Descriptions, Small Page Device (OR _x [PGS] = 0)..... | 10-37 |
| 10-28 | FPAR Field Descriptions, Large Page Device (OR _x [PGS] = 1)..... | 10-37 |
| 10-29 | FBCR Field Descriptions..... | 10-38 |
| 10-30 | GPCM Read Control Signal Timing..... | 10-45 |
| 10-31 | GPCM Write Control Signal Timing..... | 10-46 |
| 10-32 | Boot Bank Field Values after Reset for GPCM as Boot Controller..... | 10-55 |
| 10-33 | FCM Chip-Select to First Command Timing..... | 10-64 |
| 10-34 | FCM Command, Address, and Write Data Timing Parameters..... | 10-65 |
| 10-35 | FCM Read Data Timing Parameters..... | 10-68 |
| 10-36 | Boot Bank Field Values after Reset for FCM as Boot Controller..... | 10-69 |
| 10-37 | UPM Routines Start Addresses..... | 10-72 |
| 10-38 | RAM Word Field Descriptions..... | 10-78 |
| 10-39 | MxMR Loop Field Use..... | 10-84 |
| 10-40 | UPM Address Multiplexing..... | 10-85 |
| 10-41 | Data Bus Drive Requirements For Read Cycles..... | 10-92 |
| 10-42 | FCM Register Settings for Soft Reset (OR _n [PGS] = 1)..... | 10-93 |
| 10-43 | FCM Register Settings for Status Read (OR _n [PGS] = 1)..... | 10-93 |
| 10-44 | FCM Register Settings for ID Read (OR _n [PGS] = 1)..... | 10-94 |
| 10-45 | FCM Register Settings for Page Read (OR _n [PGS] = 1)..... | 10-94 |
| 10-46 | FCM Register Settings for Block Erase (OR _n [PGS] = 1)..... | 10-95 |
| 10-47 | FCM Register Settings for Page Program (OR _n [PGS] = 1)..... | 10-96 |
| 11-1 | Sequencer Memory Map..... | 11-2 |
| 11-2 | POTAR _n Field Descriptions..... | 11-3 |
| 11-3 | POBAR _n Field Descriptions..... | 11-4 |
| 11-4 | POCMR _n Field Descriptions..... | 11-4 |
| 11-5 | PMCR Field Descriptions..... | 11-5 |
| 11-6 | DTCR Field Descriptions..... | 11-6 |
| 12-1 | DMA Interface Signals—Detailed Signal Descriptions..... | 12-2 |
| 12-2 | Module Memory Map..... | 12-3 |
| 12-3 | OMISR Field Descriptions..... | 12-5 |
| 12-4 | OMIMR Field Descriptions..... | 12-5 |
| 12-5 | IMR0 and IMR1 Field Descriptions..... | 12-6 |
| 12-6 | OMR0 and OMR1 Field Descriptions..... | 12-6 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 12-7 | ODR Field Descriptions | 12-7 |
| 12-8 | IDR Field Descriptions | 12-8 |
| 12-9 | IMISR Field Descriptions | 12-8 |
| 12-10 | IMIMR Field Descriptions | 12-9 |
| 12-11 | DMAMR _n Field Descriptions | 12-10 |
| 12-12 | DMASR _n Field Descriptions | 12-12 |
| 12-13 | DMACDAR _n Field Descriptions | 12-13 |
| 12-14 | DMASAR _n Field Descriptions | 12-14 |
| 12-15 | DMASAR _n Field Descriptions | 12-14 |
| 12-16 | DMABCR _n Field Descriptions | 12-15 |
| 12-17 | DMANDAR _n Field Descriptions | 12-15 |
| 12-18 | DMA Segment Descriptor Fields | 12-20 |
| 13-1 | PCI Controller Modes | 13-3 |
| 13-2 | Signal Properties | 13-4 |
| 13-3 | PCI Interface Signals—Detailed Signal Descriptions | 13-5 |
| 13-4 | PCI Configuration Access Registers | 13-11 |
| 13-5 | PCI Memory-Mapped Registers | 13-11 |
| 13-6 | PCI_CONFIG_ADDRESS Field Descriptions | 13-13 |
| 13-7 | PCI_CONFIG_DATA Field Descriptions | 13-15 |
| 13-8 | PCI_ESR Field Descriptions | 13-16 |
| 13-9 | PCI_ECDR Field Descriptions | 13-17 |
| 13-10 | PCI_EER Field Descriptions | 13-17 |
| 13-11 | PCI_EATCR Field Descriptions | 13-18 |
| 13-12 | PCI_EACR Field Description | 13-19 |
| 13-13 | PCI_EEACR Field Description | 13-20 |
| 13-14 | PCI_EDLCR Field Description | 13-20 |
| 13-15 | PCI_GCR Field Descriptions | 13-21 |
| 13-16 | PCI_ECR Field Descriptions | 13-22 |
| 13-17 | PCI_GSR Field Descriptions | 13-22 |
| 13-18 | PITAR _n Field Descriptions | 13-23 |
| 13-19 | PIBAR _n Field Descriptions | 13-23 |
| 13-20 | PIEBAR _n Field Descriptions | 13-24 |
| 13-21 | PIWAR _n Field Descriptions | 13-24 |
| 13-22 | PCI Configuration Space Registers | 13-26 |
| 13-23 | Vendor ID Configuration Register Field Descriptions | 13-27 |
| 13-24 | Device ID Configuration Register Field Descriptions | 13-27 |
| 13-25 | PCI Command Configuration Register Field Descriptions | 13-28 |
| 13-26 | PCI Status Configuration Register Field Descriptions | 13-29 |
| 13-27 | Revision ID Configuration Register Field Descriptions | 13-30 |
| 13-28 | Standard Programming Interface Configuration Register Field Descriptions | 13-30 |
| 13-29 | Subclass Code Configuration Register Field Descriptions | 13-31 |

Tables

| Table Number | Title | Page Number |
|--------------|---|-------------|
| 13-30 | Class Code Configuration Register Field Descriptions | 13-31 |
| 13-31 | Cache Line Size Configuration Register Field Descriptions | 13-32 |
| 13-32 | Latency Timer Configuration Register Field Descriptions | 13-32 |
| 13-33 | PIMMR Base Address Configuration Register Field Descriptions | 13-33 |
| 13-34 | GPL Base Address Register 0 Field Descriptions | 13-34 |
| 13-35 | GPL Base Address Register 1,2 Field Descriptions | 13-35 |
| 13-36 | GPL Extended Base Address Registers 1–2 Field Descriptions | 13-35 |
| 13-37 | Subsystem Vendor ID Configuration Register Field Descriptions | 13-36 |
| 13-38 | Subsystem Device ID Configuration Register Field Descriptions | 13-36 |
| 13-39 | Interrupt Line Configuration Register Field Descriptions | 13-37 |
| 13-40 | PCI Function Configuration Register Field Descriptions | 13-38 |
| 13-41 | PCI Arbiter Control Register (PCIACR) Field Descriptions | 13-39 |
| 13-42 | Hot Swap Register Block Field Descriptions | 13-40 |
| 13-43 | PCIPMR0 Field Descriptions | 13-41 |
| 13-44 | PCIPMR1 Field Descriptions | 13-42 |
| 13-45 | PCI Command Definitions | 13-46 |
| 13-46 | Special Cycle Commands | 13-56 |
| 14-1 | PCI Express Interface Signals—Detailed Signal Descriptions | 14-5 |
| 14-2 | PCI Express Controller Register Groups | 14-6 |
| 14-3 | PCI Express Memory Map | 14-6 |
| 14-4 | PCI Express Vendor ID Register Field Description | 14-15 |
| 14-5 | PCI Express Device ID Register Field Description | 14-16 |
| 14-6 | PCI Express Command Register Field Descriptions | 14-16 |
| 14-7 | PCI Express Status Register Field Descriptions | 14-18 |
| 14-8 | PCI Express Revision ID Register Field Descriptions | 14-19 |
| 14-9 | PCI Express Class Code Register Field Descriptions | 14-19 |
| 14-10 | PCI Express Bus Cache Line Size Register Field Descriptions | 14-20 |
| 14-11 | PCI Express Latency Timer Register Field Descriptions | 14-20 |
| 14-12 | PCI Express Header Type Register Field Descriptions | 14-21 |
| 14-13 | BAR0 and BAR1 Register Field Descriptions | 14-22 |
| 14-14 | BAR2 and BAR4 Register Field Descriptions | 14-23 |
| 14-15 | BAR3 and BAR5 Register Field Descriptions | 14-23 |
| 14-16 | PCI Express Subsystem Vendor ID Register Field Descriptions | 14-24 |
| 14-17 | PCI Express Subsystem ID Register Field Descriptions | 14-24 |
| 14-18 | PCI Express Capabilities Pointer Register Field Descriptions | 14-25 |
| 14-19 | PCI Express Interrupt Line Register Field Descriptions | 14-25 |
| 14-20 | PCI Express Interrupt Pin Register Field Descriptions | 14-26 |
| 14-21 | PCI Express MInimum Grant Register Field Descriptions | 14-26 |
| 14-22 | PCI Express Maximum Latency Register Field Descriptions | 14-27 |
| 14-23 | PCI Express Primary Bus Number Register Field Descriptions | 14-28 |
| 14-24 | PCI Express Secondary Bus Number Register Field Descriptions | 14-28 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 14-25 | PCI Express Subordinate Bus Number Register Field Descriptions | 14-28 |
| 14-26 | PCI Express I/O Base Register Field Descriptions | 14-29 |
| 14-27 | PCI Express I/O Limit Register Field Descriptions | 14-30 |
| 14-28 | PCI Express Secondary Status Register Field Descriptions | 14-30 |
| 14-29 | PCI Express Memory Base Register Field Descriptions | 14-31 |
| 14-30 | PCI Express Memory Limit Register Field Descriptions | 14-31 |
| 14-31 | PCI Express Prefetchable Memory Base Register Field Descriptions | 14-32 |
| 14-32 | PCI Express Prefetchable Memory Limit Register Field Descriptions | 14-32 |
| 14-33 | PCI Express Prefetchable Base Upper 32-Bit Register Field Descriptions | 14-33 |
| 14-34 | PCI Express Prefetchable Limit Upper 32-Bit Register Field Descriptions | 14-33 |
| 14-35 | PCI Express I/O Base Upper 16-Bit Register Field Descriptions | 14-34 |
| 14-36 | PCI Express I/O Limit Upper 16-Bit Register Field Descriptions | 14-34 |
| 14-37 | PCI Express Capabilities Pointer Register Field Descriptions | 14-34 |
| 14-38 | PCI Express Interrupt Line Register Field Descriptions | 14-35 |
| 14-39 | PCI Express Interrupt Pin Register Field Descriptions | 14-35 |
| 14-40 | PCI Express Bridge Control Register Field Descriptions | 14-36 |
| 14-41 | PCI Express Power Management Capability ID Register Field Descriptions | 14-38 |
| 14-42 | PCI Express Power Management Next Capabilities Pointer Field Descriptions | 14-38 |
| 14-43 | PCI Express Power Management Capabilities Register Field Descriptions | 14-39 |
| 14-44 | PCI Express Power Management Status and Control Register Field Descriptions | 14-39 |
| 14-45 | PCI Express Power Management Data Register Field Descriptions | 14-40 |
| 14-46 | PCI Express Capability ID Register Field Descriptions | 14-41 |
| 14-47 | PCI Express Next Capabilities Pointer Field Descriptions | 14-41 |
| 14-48 | PCI Express Capabilities Register Field Descriptions | 14-41 |
| 14-49 | PCI Express Device Capabilities Register Field Descriptions | 14-42 |
| 14-50 | PCI Express Device Control Register Field Descriptions | 14-43 |
| 14-51 | PCI Express Device Status Register Field Descriptions | 14-44 |
| 14-52 | PCI Express Link Capabilities Register Field Descriptions | 14-45 |
| 14-53 | PCI Express Link Control Register Field Descriptions | 14-45 |
| 14-54 | PCI Express Link Status Register Field Descriptions | 14-46 |
| 14-55 | PCI Express Slot Capabilities Register Field Descriptions | 14-47 |
| 14-56 | PCI Express Slot Control Register Field Descriptions | 14-47 |
| 14-57 | PCI Express Slot Status Register Field Descriptions | 14-48 |
| 14-58 | PCI Express Root Control Register Field Descriptions | 14-49 |
| 14-59 | PCI Express Root Status Register Field Descriptions | 14-49 |
| 14-60 | PCI Express Capability ID Register Field Descriptions | 14-50 |
| 14-61 | PCI Express MSI Message Control Register Field Descriptions | 14-50 |
| 14-62 | PCI Express MSI Message Address Register Field Descriptions | 14-51 |
| 14-63 | PCI Express MSI Message Upper Address Register Field Descriptions | 14-51 |
| 14-64 | PCI Express MSI Message Data Register Field Descriptions | 14-51 |
| 14-65 | PCI Express Advanced Error Reporting Capability ID Register Field Descriptions | 14-53 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 14-66 | PCI Express Uncorrectable Error Status Register Field Descriptions | 14-53 |
| 14-67 | PCI Express Uncorrectable Error Mask Register Field Descriptions | 14-54 |
| 14-68 | PCI Express Uncorrectable Error Severity Register Field Descriptions..... | 14-55 |
| 14-69 | PCI Express Correctable Error Status Register Field Descriptions | 14-56 |
| 14-70 | PCI Express Correctable Error Mask Register Field Descriptions | 14-57 |
| 14-71 | PCI Express Advanced Error Capabilities and Control Register Field Descriptions | 14-57 |
| 14-72 | PCI Express Header Log Register Field Descriptions | 14-59 |
| 14-73 | PCI Express Root Error Command Register Field Descriptions | 14-59 |
| 14-74 | PCI Express Root Error Command Register Field Descriptions | 14-60 |
| 14-75 | PCI Express Error Source Identification Register Field Descriptions | 14-60 |
| 14-76 | PEX_LTSSM_STAT Field Descriptions | 14-61 |
| 14-77 | PEX_LTSSM_STAT Status Codes..... | 14-61 |
| 14-78 | PCI Express N_FTS Control Register Field Descriptions | 14-63 |
| 14-79 | PCI Express ACK Replay Timeout Register Field Descriptions | 14-64 |
| 14-80 | PEX_GCLK_RATIO Field Descriptions | 14-65 |
| 14-81 | PEX_PM_TIMER Field Descriptions | 14-65 |
| 14-82 | PEX_PME_TIMEOUT Field Descriptions | 14-66 |
| 14-83 | PCI Express ASPM Request Timer Register Field Descriptions..... | 14-67 |
| 14-84 | PEX_SSVID_UPDATE Field Descriptions..... | 14-67 |
| 14-85 | PCI Express Device Capabilities Update Register Field Descriptions | 14-68 |
| 14-86 | PCI Express Link Capabilities Update Register Field Descriptions..... | 14-69 |
| 14-87 | PCI Express Slot Capabilities Update Register Field Descriptions | 14-70 |
| 14-88 | PEX_CFG_READY Field Descriptions | 14-71 |
| 14-89 | PEX_BAR_ENABLE Field Descriptions..... | 14-72 |
| 14-90 | PEX_BAR_SIZEL Field Descriptions..... | 14-73 |
| 14-91 | PEX_BAR_SEL Field Descriptions | 14-73 |
| 14-92 | PEX_BAR_PF Field Descriptions | 14-74 |
| 14-93 | PEX_PME_TO_ACK_TOR Field Descriptions..... | 14-75 |
| 14-94 | PEX_PME_TO_ACK_SR Field Descriptions | 14-75 |
| 14-95 | PEX_SS_INTR_MASK Field Descriptions | 14-76 |
| 14-96 | PEX_CSB_CTRL Register Field Descriptions..... | 14-77 |
| 14-97 | PEX_DMA_DSTMR Field Descriptions..... | 14-78 |
| 14-98 | PEX_CSB_STAT Register Field Descriptions..... | 14-79 |
| 14-99 | PEX_CSB_OBCTRL Register Field Descriptions | 14-80 |
| 14-100 | PEX_CSB_OBSTAT Register Field Descriptions | 14-80 |
| 14-101 | PEX_CSB_IBCTRL Register Field Descriptions..... | 14-81 |
| 14-102 | PEX_CSB_IBSTAT Register Field Descriptions..... | 14-82 |
| 14-103 | PEX_WDMA_CTRL Register Field Descriptions | 14-82 |
| 14-104 | PEX_WDMA_ADDR Register Field Descriptions..... | 14-83 |
| 14-105 | PEX_WDMA_STAT Register Field Descriptions | 14-83 |
| 14-106 | PEX_RDMA_CTRL Register Field Descriptions | 14-84 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 14-107 | PEX_RDMA_ADDR Register Field Descriptions | 14-85 |
| 14-108 | PEX_RDMA_STAT Register Field Descriptions | 14-85 |
| 14-109 | PEX_OMBCR Register Field Descriptions | 14-87 |
| 14-110 | PEX_OMBDR Register Field Descriptions | 14-87 |
| 14-111 | PEX_IMBCR Register Field Descriptions | 14-88 |
| 14-112 | PEX_IMBDR Register Field Descriptions | 14-88 |
| 14-113 | PEX_HIER Register Field Descriptions | 14-89 |
| 14-114 | PEX_HISR Register Field Descriptions | 14-90 |
| 14-115 | PEX_HOPIVR Register Field Descriptions | 14-91 |
| 14-116 | PEX_HIPIVR Register Field Descriptions | 14-91 |
| 14-117 | PEX_HWDIVR Register Field Descriptions | 14-92 |
| 14-118 | PEX_HRDIVR Register Field Descriptions | 14-92 |
| 14-119 | PEX_HMIVR Register Field Descriptions | 14-93 |
| 14-120 | PEX_CSPIER Register Field Descriptions | 14-93 |
| 14-121 | PEX_CSWDIER Register Field Descriptions | 14-94 |
| 14-122 | PEX_CSRDIER Register Field Descriptions | 14-95 |
| 14-123 | PEX_CSMIER Register Field Descriptions | 14-96 |
| 14-124 | PEX_CSPISR Register Field Descriptions | 14-97 |
| 14-125 | PEX_CSWDISR Register Field Descriptions | 14-98 |
| 14-126 | PEX_CSRDISR Register Field Descriptions | 14-98 |
| 14-127 | PEX_CSMISR Register Field Descriptions | 14-99 |
| 14-128 | PEX_PM_CTRL Register Field Descriptions | 14-101 |
| 14-129 | PCI Express Slot Control Misc Register Field Descriptions | 14-102 |
| 14-130 | PEX_OWAR0–PEX_OWAR3 Register Field Descriptions | 14-102 |
| 14-131 | PEX_OWBAR n Register Field Descriptions | 14-103 |
| 14-132 | PEX_OWTARL n Register Field Descriptions | 14-104 |
| 14-133 | PEX_OWTARH n Register Field Descriptions | 14-104 |
| 14-134 | EP Inbound Base and Translation Address Registers Correspondence | 14-105 |
| 14-135 | PEX_EPIWTAR n Register Field Descriptions | 14-106 |
| 14-136 | PEX_RCIWAR n Register Field Descriptions | 14-106 |
| 14-137 | PEX_RCIWTAR n Register Field Descriptions | 14-107 |
| 14-138 | PEX_RCIWBARL n Register Field Descriptions | 14-108 |
| 14-139 | PEX_RCIWBARH n Register Field Descriptions | 14-108 |
| 14-140 | Address Translation Window Combinations | 14-110 |
| 14-141 | Configuration Address Mapping | 14-115 |
| 14-142 | PCI Express RC Inbound Message Handling | 14-119 |
| 14-143 | PCI Express EP Inbound Message Handling | 14-120 |
| 14-144 | Initial Credit Advertisement | 14-123 |
| 14-145 | Power Management State Supported | 14-125 |
| 14-146 | DMA Descriptor Bit Field Descriptions | 14-128 |
| 15-1 | SATA Register Summary | 15-4 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 15-2 | CQR Field Descriptions | 15-6 |
| 15-3 | CAR Field Descriptions | 15-6 |
| 15-4 | CCR Field Descriptions | 15-7 |
| 15-5 | CER Field Descriptions | 15-8 |
| 15-6 | DER Field Descriptions | 15-8 |
| 15-7 | CHBA Field Descriptions | 15-9 |
| 15-8 | HStatus Field Descriptions | 15-9 |
| 15-9 | HControl Field Descriptions | 15-12 |
| 15-10 | CQPMP Field Descriptions | 15-14 |
| 15-11 | SIG Register Field Descriptions | 15-14 |
| 15-12 | ICC Field Descriptions | 15-15 |
| 15-13 | SStatus Field Descriptions | 15-15 |
| 15-14 | SError Field Descriptions | 15-16 |
| 15-15 | SControl Field Descriptions | 15-19 |
| 15-16 | SNotification Field Descriptions | 15-20 |
| 15-17 | TransCfg Field Descriptions | 15-20 |
| 15-18 | TransStatus Field Descriptions | 15-21 |
| 15-19 | LinkCfg Field Descriptions | 15-21 |
| 15-20 | LinkCfg1 Field Descriptions | 15-23 |
| 15-21 | LinkCfg2 Field Descriptions | 15-23 |
| 15-22 | LinkStatus Field Descriptions | 15-24 |
| 15-23 | LinkStatus1 Field Descriptions | 15-25 |
| 15-24 | PhyCtrlCfg1 Field Descriptions | 15-26 |
| 15-25 | CommandStatus Field Descriptions | 15-28 |
| 15-26 | PhyCtrlCfg2 Field Descriptions | 15-29 |
| 15-27 | SYSPR Field Descriptions | 15-31 |
| 15-28 | Command Header Word 0—Data Base Address | 15-32 |
| 15-29 | Command Header Word 1—FIS_LEN | 15-32 |
| 15-30 | Command Header Word 2—Data Base Address | 15-33 |
| 15-31 | Command Header Word 3—Description Information | 15-33 |
| 15-32 | PRDT Word 0—Data Base Address | 15-36 |
| 15-33 | PRDT Word 3—Description Information | 15-36 |
| 15-34 | Vendor BIST Test—Command Header | 15-37 |
| 15-35 | Vendor BIST Test—Command Descriptor | 15-37 |
| 15-36 | Read DMA Command—Command Header | 15-45 |
| 15-37 | Read DMA Command—Command Descriptor | 15-45 |
| 15-38 | Read DMA Command—PRD Entries | 15-45 |
| 16-1 | SerDes External Signals—Detailed Signal Descriptions | 16-3 |
| 16-2 | SerDes PHY Block Memory Map | 16-4 |
| 16-3 | SRDSCR0 Field Descriptions | 16-5 |
| 16-4 | SRDSCR1 Field Descriptions | 16-8 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 16-5 | SRDSCR2 Field Descriptions | 16-10 |
| 16-6 | SRDSCR3 Field Descriptions | 16-11 |
| 16-7 | SRDSCR4 Field Descriptions | 16-12 |
| 16-8 | SRDSRSTCTL Field Descriptions | 16-13 |
| 17-1 | USB External Signals..... | 17-3 |
| 17-2 | UTMI Signal Descriptions | 17-4 |
| 17-3 | ULPI Signal Descriptions | 17-5 |
| 17-4 | USB Interface Memory Map..... | 17-7 |
| 17-5 | CAPLENGTH Register Field Descriptions | 17-9 |
| 17-6 | HCIVERSION Register Field Descriptions..... | 17-9 |
| 17-7 | HCSPARAMS Register Field Descriptions | 17-10 |
| 17-8 | HCCPARAMS Register Field Descriptions..... | 17-11 |
| 17-9 | DCIVERSION Register Field Descriptions..... | 17-12 |
| 17-10 | DCCPARAMS Register Field Descriptions..... | 17-12 |
| 17-11 | USBCMD Register Field Descriptions | 17-13 |
| 17-12 | USBSTS Register Field Descriptions | 17-16 |
| 17-13 | USBINTR Register Field Descriptions..... | 17-18 |
| 17-14 | FRINDEX Register Field Descriptions..... | 17-20 |
| 17-15 | FRINDEX N Values..... | 17-20 |
| 17-16 | PERIODICLISTBASE Register Field Descriptions | 17-21 |
| 17-17 | DEVICEADDR Register Field Descriptions..... | 17-21 |
| 17-18 | ASYNCLISTADDR Register Field Descriptions | 17-22 |
| 17-19 | ENDPOINTLISTADDR Register Field Descriptions | 17-23 |
| 17-20 | BURSTSIZE Register Field Descriptions..... | 17-23 |
| 17-21 | TXFILLTUNING Register Field Descriptions | 17-24 |
| 17-22 | ULPI VIEWPORT Field Descriptions | 17-26 |
| 17-23 | CONFIGFLAG Register Field Descriptions..... | 17-27 |
| 17-24 | PORTSC Register Field Descriptions | 17-28 |
| 17-25 | OTGSC Register Field Descriptions..... | 17-33 |
| 17-26 | USBMODE Register Field Descriptions | 17-35 |
| 17-27 | ENDPTSETUPSTAT Register Field Descriptions..... | 17-36 |
| 17-28 | ENDPTPRIME Register Field Descriptions | 17-37 |
| 17-29 | ENDPTFLUSH Register Field Descriptions..... | 17-37 |
| 17-30 | ENDPTSTATUS Register Field Descriptions | 17-38 |
| 17-31 | ENDPTCOMPLETE Register Field Descriptions | 17-39 |
| 17-32 | ENDPTCTRL0 Register Field Descriptions | 17-39 |
| 17-33 | ENDPTCTRL n Register Field Descriptions | 17-40 |
| 17-34 | SNOOP n Register Field Descriptions..... | 17-42 |
| 17-35 | AGE_CNT_THRESH Register Field Descriptions | 17-44 |
| 17-36 | PRI_CTRL Register Field Descriptions | 17-45 |
| 17-37 | SI_CTRL Register Field Descriptions | 17-45 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 17-38 | CONTROL Register Field Descriptions | 17-46 |
| 17-39 | Supported PHY Interfaces | 17-49 |
| 17-40 | Typ Field Encodings | 17-51 |
| 17-41 | Next Schedule Element Pointer | 17-52 |
| 17-42 | iTD Transaction Status and Control..... | 17-53 |
| 17-43 | Buffer Pointer Page 0 (Plus) | 17-54 |
| 17-44 | iTD Buffer Pointer Page 1 (Plus)..... | 17-54 |
| 17-45 | Buffer Pointer Page 2 (Plus) | 17-54 |
| 17-46 | Buffer Pointer Page 3–6..... | 17-55 |
| 17-47 | Next Link Pointer | 17-55 |
| 17-48 | Endpoint and Transaction Translator Characteristics | 17-56 |
| 17-49 | Microframe Schedule Control..... | 17-56 |
| 17-50 | siTD Transfer Status and Control..... | 17-57 |
| 17-51 | siTD Buffer Pointer Page 0 (Plus) | 17-58 |
| 17-52 | siTD Buffer Pointer Page 1 (Plus) | 17-58 |
| 17-53 | siTD Back Link Pointer | 17-58 |
| 17-54 | qTD Next Element Transfer Pointer (DWord 0) | 17-60 |
| 17-55 | qTD Alternate Next Element Transfer Pointer (DWord 1) | 17-60 |
| 17-56 | qTD Token (DWord 2)..... | 17-61 |
| 17-57 | qTD Buffer Pointer | 17-64 |
| 17-58 | Queue Head DWord 0 | 17-65 |
| 17-59 | Endpoint Characteristics: Queue Head DWord 1..... | 17-66 |
| 17-60 | Endpoint Capabilities: Queue Head DWord 2 | 17-67 |
| 17-61 | Current qTD Link Pointer | 17-68 |
| 17-62 | Host-Controller Rules for Bits in Overlay (DWords 5, 6, 8, and 9) | 17-69 |
| 17-63 | FSTN Normal Path Pointer | 17-70 |
| 17-64 | FSTN Back Path Link Pointer | 17-70 |
| 17-65 | Behavior During Wake-Up Events..... | 17-74 |
| 17-66 | Operation of FRINDEX and SOFV (SOF Value Register)..... | 17-79 |
| 17-67 | Example Periodic Reference Patterns for Interrupt Transfers | 17-92 |
| 17-68 | Ping Control State Transition Table | 17-93 |
| 17-69 | Interrupt IN/OUT Do Complete Split State Execution Criteria..... | 17-107 |
| 17-70 | Initial Conditions for OUT siTD TP and T-Count Fields | 17-116 |
| 17-71 | Transaction Position (TP)/Transaction Count (T-Count) Transition Table..... | 17-116 |
| 17-72 | Summary siTD Split Transaction State..... | 17-120 |
| 17-73 | Example Case 2a—Software Scheduling siTDs for an IN Endpoint..... | 17-121 |
| 17-74 | Summary of Transaction Errors | 17-124 |
| 17-75 | Summary Behavior on Host System Errors | 17-127 |
| 17-76 | Endpoint Capabilities/Characteristics | 17-129 |
| 17-77 | Current dTD Pointer..... | 17-130 |
| 17-78 | Multiple Mode Control | 17-131 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 17-79 | Next dTD Pointer | 17-131 |
| 17-80 | dTD Token | 17-132 |
| 17-81 | Buffer Pointer Page 0 | 17-132 |
| 17-82 | Buffer Pointer Page 1 | 17-133 |
| 17-83 | Buffer Pointer Pages 2–4 | 17-133 |
| 17-84 | Device Controller State Information Bits | 17-136 |
| 17-85 | Device Controller Endpoint Initialization | 17-139 |
| 17-86 | Device Controller Stall Response Matrix | 17-140 |
| 17-87 | Variable Length Transfer Protocol Example (ZLT = 0) | 17-142 |
| 17-88 | Variable Length Transfer Protocol Example (ZLT = 1) | 17-142 |
| 17-89 | Interrupt/Bulk Endpoint Bus Response Matrix | 17-144 |
| 17-90 | Control Endpoint Bus Response Matrix | 17-146 |
| 17-91 | Isochronous Endpoint Bus Response Matrix | 17-149 |
| 17-92 | Device Error Matrix | 17-154 |
| 17-93 | Error Descriptions | 17-154 |
| 17-94 | Interrupt Handling Order | 17-154 |
| 17-95 | Low Frequency Interrupt Events | 17-155 |
| 17-96 | Error Interrupt Events | 17-155 |
| 17-97 | Functional Differences Between EHCI and EHCI with Embedded TT | 17-157 |
| 17-98 | Emulated Handshakes | 17-158 |
| 17-99 | ULPI Timing | 17-162 |
| 18-1 | Example Descriptor | 18-4 |
| 18-2 | SEC Address Map | 18-11 |
| 18-3 | SEC Address Map | 18-12 |
| 18-4 | Header Dword Bit Definitions | 18-18 |
| 18-5 | Header Dword Writeback Bit Definitions | 18-19 |
| 18-6 | EU_SEL0 and EU_SEL1 Values | 18-20 |
| 18-7 | Descriptor Types | 18-20 |
| 18-8 | Pointer Dword Field Definitions | 18-22 |
| 18-9 | Link Table Field Definitions | 18-24 |
| 18-10 | Descriptor Format Summary | 18-27 |
| 18-11 | CCR Bit Definitions | 18-34 |
| 18-12 | CSR Bit Descriptions | 18-36 |
| 18-13 | CSR Error Field Definitions | 18-37 |
| 18-14 | Current Descriptor Pointer Register Signals | 18-38 |
| 18-15 | Fetch FIFO Field Descriptions | 18-39 |
| 18-16 | Channel Assignment Value | 18-45 |
| 18-17 | Fields in Interrupt Enable, Interrupt Status, and Interrupt Clear Registers | 18-46 |
| 18-18 | IP Block Revision Register Fields | 18-50 |
| 18-19 | MCR Bit Descriptions | 18-51 |
| 18-20 | AESU Mode Register Field Descriptions | 18-54 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 18-21 | AES Cipher Modes | 18-56 |
| 18-22 | Use of Data Size Register | 18-57 |
| 18-23 | AESU Reset Control Register Field Descriptions | 18-58 |
| 18-24 | AESU Status Register Field Descriptions..... | 18-59 |
| 18-25 | AESU Interrupt Status Register Field Descriptions..... | 18-60 |
| 18-26 | AESU Interrupt Mask Register Field Descriptions..... | 18-62 |
| 18-27 | AESU Context Registers..... | 18-64 |
| 18-28 | AESU Context Registers..... | 18-65 |
| 18-29 | GCM Cipher Mode Auxiliary Bit Definitions | 18-70 |
| 18-30 | GCM Encryption Context | 18-73 |
| 18-31 | GCM Decryption Context | 18-74 |
| 18-32 | GCM with ICV Context..... | 18-75 |
| 18-33 | GCM-GHASH Context..... | 18-76 |
| 18-34 | CRCU Mode Register Bit Definitions | 18-80 |
| 18-35 | CRCU Reset Control Register Field Descriptions..... | 18-82 |
| 18-36 | CRCU Status Register Bit Definitions | 18-83 |
| 18-37 | CRCU Interrupt Status Register Bit Definitions..... | 18-85 |
| 18-38 | CRCU Interrupt Mask Register Bit Definitions..... | 18-86 |
| 18-39 | DEU Mode Register Field Descriptions | 18-89 |
| 18-40 | DEU Key Size Register Field Descriptions | 18-90 |
| 18-41 | DEU Reset Control Register Field Descriptions..... | 18-91 |
| 18-42 | DEU Status Register | 18-92 |
| 18-43 | DEU Interrupt Status Register Field Descriptions | 18-93 |
| 18-44 | DEU Interrupt Mask Register Field Descriptions..... | 18-95 |
| 18-45 | MDEU Mode Register in Old Configuration..... | 18-98 |
| 18-46 | MDEU Mode Register in New Configuration | 18-100 |
| 18-47 | Mode Register—HMAC or SSL-MAC Generated by Single Descriptor..... | 18-101 |
| 18-48 | Mode Register—HMAC Generated across a Sequence of Descriptors..... | 18-101 |
| 18-49 | MDEU Reset Control Register Field Descriptions | 18-103 |
| 18-50 | MDEU Status Register Field Descriptions | 18-104 |
| 18-51 | MDEU Interrupt Status Register Field Descriptions | 18-105 |
| 18-52 | MDEU Interrupt Mask Register Field Descriptions | 18-106 |
| 18-53 | ROUTINE Field Description | 18-112 |
| 18-54 | PKEU Reset Control Register Field Descriptions | 18-114 |
| 18-55 | PKEU Status Register Field Descriptions..... | 18-115 |
| 18-56 | PKEU Interrupt Status Register Field Descriptions..... | 18-117 |
| 18-57 | PKEU Interrupt Mask Register Field Descriptions..... | 18-118 |
| 18-58 | RNGU Reset Control Register Field Descriptions | 18-121 |
| 18-59 | RNGU Status Register Field Descriptions..... | 18-122 |
| 18-60 | RNGU Interrupt Status Register Field Descriptions..... | 18-123 |
| 18-61 | RNGU Interrupt Mask Register Field Descriptions..... | 18-124 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 19-1 | eTSEC _n Network Interface Signal Properties | 19-6 |
| 19-2 | eTSEC Signals—Detailed Signal Descriptions | 19-8 |
| 19-3 | Module Memory Map Summary | 19-12 |
| 19-4 | Module Memory Map | 19-12 |
| 19-5 | TSEC_ID Field Descriptions | 19-23 |
| 19-6 | TSEC_ID2 Field Descriptions | 19-23 |
| 19-7 | TSEC_ID2[TSEC_INT] Field Settings | 19-23 |
| 19-8 | IEVENT Field Descriptions | 19-25 |
| 19-9 | IMASK Field Descriptions | 19-29 |
| 19-10 | EDIS Field Descriptions | 19-30 |
| 19-11 | ECNTRL Field Descriptions | 19-32 |
| 19-12 | eTSEC Interface Configurations | 19-33 |
| 19-13 | PTV Field Descriptions | 19-34 |
| 19-14 | DMACTRL Field Descriptions | 19-34 |
| 19-15 | TBIPA Field Descriptions | 19-36 |
| 19-16 | TCTRL Field Descriptions | 19-36 |
| 19-17 | TSTAT Field Descriptions | 19-39 |
| 19-18 | DFVLAN Field Descriptions | 19-42 |
| 19-19 | TXIC Field Descriptions | 19-43 |
| 19-20 | TQUEUE Field Descriptions | 19-44 |
| 19-21 | TR03WT Field Descriptions | 19-45 |
| 19-22 | TR47WT Field Descriptions | 19-46 |
| 19-23 | TBPTR _n Field Descriptions | 19-46 |
| 19-24 | TBASE0–TBASE7 Field Descriptions | 19-47 |
| 19-25 | TMR_TXTS _n _ID Register Field Descriptions | 19-47 |
| 19-26 | TMR_TXTS _n _H/L Register Field Descriptions | 19-48 |
| 19-27 | RCTRL Field Descriptions | 19-49 |
| 19-28 | RSTAT Field Descriptions | 19-51 |
| 19-29 | RXIC Field Descriptions | 19-53 |
| 19-30 | RQUEUE Field Descriptions | 19-53 |
| 19-31 | RBIFX Field Descriptions | 19-55 |
| 19-32 | RQFAR Field Descriptions | 19-56 |
| 19-33 | RQFCR Field Descriptions | 19-56 |
| 19-34 | RQFPR Field Descriptions | 19-58 |
| 19-35 | MRBLR Field Descriptions | 19-61 |
| 19-36 | RBPTR _n Field Descriptions | 19-62 |
| 19-37 | RBASE0–RBASE7 Field Descriptions | 19-62 |
| 19-38 | TMR_RXTS_H/L Register Field Descriptions | 19-63 |
| 19-39 | MACCFG1 Field Descriptions | 19-66 |
| 19-40 | MACCFG2 Field Descriptions | 19-68 |
| 19-41 | IPGIFG Field Descriptions | 19-69 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 19-42 | HAFDUP Field Descriptions | 19-70 |
| 19-43 | MAXFRM Field Descriptions | 19-71 |
| 19-44 | MIIMCFG Field Descriptions..... | 19-72 |
| 19-45 | MIIMCOM Descriptions..... | 19-73 |
| 19-46 | MIIMADD Field Descriptions..... | 19-73 |
| 19-47 | MIIMCON Field Descriptions..... | 19-74 |
| 19-48 | MIIMSTAT Field Descriptions | 19-74 |
| 19-49 | MIIMIND Field Descriptions | 19-75 |
| 19-50 | IFSTAT Field Descriptions | 19-75 |
| 19-51 | MACSTNADDR1 Field Descriptions | 19-76 |
| 19-52 | MACSTNADDR2 Field Descriptions | 19-77 |
| 19-53 | MAC _n ADDR1 Field Descriptions..... | 19-78 |
| 19-54 | MAC01ADDR2–MAC15ADDR2 Field Descriptions | 19-78 |
| 19-55 | TR64 Field Descriptions | 19-79 |
| 19-56 | TR127 Field Descriptions | 19-80 |
| 19-57 | TR255 Field Descriptions | 19-80 |
| 19-58 | TR511 Field Descriptions | 19-81 |
| 19-59 | TR1K Field Descriptions | 19-81 |
| 19-60 | TRMAX Field Descriptions..... | 19-82 |
| 19-61 | TRMGV Field Descriptions..... | 19-82 |
| 19-62 | RBYT Field Descriptions..... | 19-83 |
| 19-63 | RPKT Field Descriptions | 19-83 |
| 19-64 | RFCS Field Descriptions | 19-84 |
| 19-65 | RMCA Field Descriptions | 19-84 |
| 19-66 | RBCA Field Descriptions | 19-84 |
| 19-67 | RXCF Field Descriptions..... | 19-85 |
| 19-68 | RXPF Field Descriptions | 19-85 |
| 19-69 | RXUO Field Descriptions..... | 19-86 |
| 19-70 | RALN Field Descriptions | 19-86 |
| 19-71 | RFLR Field Descriptions | 19-87 |
| 19-72 | RCDE Field Descriptions..... | 19-87 |
| 19-73 | RCSE Field Descriptions | 19-88 |
| 19-74 | RUND Field Descriptions..... | 19-88 |
| 19-75 | ROVR Field Descriptions | 19-89 |
| 19-76 | RFRG Field Descriptions..... | 19-89 |
| 19-77 | RJBR Field Descriptions..... | 19-90 |
| 19-78 | RDRP Field Descriptions..... | 19-90 |
| 19-79 | TBYT Field Descriptions..... | 19-91 |
| 19-80 | TPKT Field Descriptions | 19-91 |
| 19-81 | TMCA Field Descriptions..... | 19-92 |
| 19-82 | TBCA Field Descriptions..... | 19-92 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 19-83 | TXPF Field Descriptions | 19-93 |
| 19-84 | TDFR Field Descriptions | 19-93 |
| 19-85 | TEDF Field Descriptions | 19-94 |
| 19-86 | TSCL Field Descriptions | 19-94 |
| 19-87 | TMCL Field Descriptions | 19-95 |
| 19-88 | TLCL Field Descriptions | 19-95 |
| 19-89 | TXCL Field Descriptions | 19-96 |
| 19-90 | TNCL Field Descriptions | 19-96 |
| 19-91 | TDRP Field Descriptions | 19-97 |
| 19-92 | TJBR Field Descriptions | 19-97 |
| 19-93 | TFCS Field Descriptions | 19-98 |
| 19-94 | TXCF Field Descriptions | 19-98 |
| 19-95 | TOVR Field Descriptions | 19-99 |
| 19-96 | TUND Field Descriptions | 19-99 |
| 19-97 | TFRG Field Descriptions | 19-100 |
| 19-98 | CAR1 Field Descriptions | 19-100 |
| 19-99 | CAR2 Field Descriptions | 19-102 |
| 19-100 | CAM1 Field Descriptions | 19-103 |
| 19-101 | CAM2 Field Descriptions | 19-104 |
| 19-102 | RREJ Field Descriptions | 19-105 |
| 19-103 | IGADDR _n Field Descriptions | 19-106 |
| 19-104 | GADDR _n Field Descriptions | 19-107 |
| 19-105 | ATTR Field Descriptions | 19-108 |
| 19-106 | RQPRM Field Descriptions | 19-109 |
| 19-107 | RFBPTR0–RFBPTR7 Field Descriptions | 19-109 |
| 19-108 | TMR_CTRL Register Field Descriptions | 19-110 |
| 19-109 | TMR_TEVENT Register Field Descriptions | 19-112 |
| 19-110 | TMR_TEMASK Register Field Descriptions | 19-113 |
| 19-111 | TMR_PEVENT Register Field Descriptions | 19-114 |
| 19-112 | TMR_PEMASK Register Field Descriptions | 19-115 |
| 19-113 | TMR_STAT Register Field Descriptions | 19-115 |
| 19-114 | TMR_CNT_H/L Register Field Descriptions | 19-116 |
| 19-115 | TMR_ADD Register Field Descriptions | 19-117 |
| 19-116 | TMR_ACC Register Field Descriptions | 19-117 |
| 19-117 | TMR_PRSC Register Field Descriptions | 19-118 |
| 19-118 | TMROFF_H/L Register Field Descriptions | 19-118 |
| 19-119 | TMR_ALARM _n _H/L Register Field Descriptions | 19-119 |
| 19-120 | TMR_FIPER Register Field Descriptions | 19-120 |
| 19-121 | TMR_ETTS1-2_H Register Field Descriptions | 19-120 |
| 19-122 | TBI MII Register Set | 19-122 |
| 19-123 | CR Field Descriptions | 19-123 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 19-124 | SR Descriptions..... | 19-124 |
| 19-125 | ANA Field Descriptions..... | 19-125 |
| 19-126 | PAUSE Priority Resolution..... | 19-126 |
| 19-127 | ANLPBPA Field Descriptions | 19-127 |
| 19-128 | ANEX Field Descriptions | 19-128 |
| 19-129 | ANNPT Field Descriptions | 19-129 |
| 19-130 | ANLPANP Field Descriptions | 19-129 |
| 19-131 | EXST Field Descriptions | 19-130 |
| 19-132 | JD Field Descriptions..... | 19-131 |
| 19-133 | TBICON Field Descriptions | 19-132 |
| 19-134 | MII and RMII Signals Multiplexing | 19-138 |
| 19-135 | RGMII and RTBI Signals Multiplexing..... | 19-139 |
| 19-136 | SGMII Signalling..... | 19-140 |
| 19-137 | Shared Signals..... | 19-140 |
| 19-138 | Steps for Minimum Register Initialization..... | 19-141 |
| 19-139 | Custom Preamble Field Descriptions..... | 19-146 |
| 19-140 | Received Preamble Field Descriptions | 19-147 |
| 19-141 | Flow Control Frame Structure | 19-151 |
| 19-142 | Non-Error Transmit Interrupts | 19-153 |
| 19-143 | Non-Error Receive Interrupts..... | 19-153 |
| 19-144 | Interrupt Coalescing Timing Threshold Ranges | 19-154 |
| 19-145 | Transmission Errors | 19-155 |
| 19-146 | Reception Errors | 19-156 |
| 19-147 | Tx Frame Control Block Description..... | 19-159 |
| 19-148 | Rx Frame Control Block Descriptions..... | 19-161 |
| 19-149 | Supported Stack L2 Ethernet Headers | 19-163 |
| 19-150 | Special Filer Rules | 19-167 |
| 19-151 | Receive Queue Filer Interrupt Events | 19-167 |
| 19-152 | Filer Table Example—802.1p Priority Filing | 19-168 |
| 19-153 | Filer Table Example—IP Diff-Serv Code Points Filing | 19-169 |
| 19-154 | Filer Table Example—TCP and UDP Port Filing..... | 19-170 |
| 19-155 | Filer Example—Interrupt from Deep Sleep..... | 19-170 |
| 19-156 | PTP Payload Special Fields..... | 19-179 |
| 19-157 | Time-Stamp Insertion Programming Requirements | 19-180 |
| 19-158 | Tx Frame Control Block Description..... | 19-182 |
| 19-159 | Transmit Data Buffer Descriptor (TxBD) Field Descriptions | 19-186 |
| 19-160 | Receive Buffer Descriptor Field Descriptions | 19-189 |
| 19-161 | MII Interface Mode Signal Configuration | 19-191 |
| 19-162 | Shared MII Signals..... | 19-192 |
| 19-163 | MII Mode Register Initialization Steps..... | 19-192 |
| 19-164 | RGMII Interface Mode Signal Configuration..... | 19-194 |

Tables

| Table Number | Title | Page Number |
|-----------------|---|----------------|
| 19-165 | Shared RGMII Signals | 19-195 |
| 19-166 | RGMII Mode Register Initialization Steps | 19-195 |
| 19-167 | RMII Interface Mode Signal Configuration..... | 19-198 |
| 19-168 | Shared RMII Signals | 19-198 |
| 19-169 | RMII Mode Register Initialization Steps | 19-199 |
| 19-170 | RTBI Interface Mode Signal Configuration..... | 19-202 |
| 19-171 | Shared RTBI Signals | 19-202 |
| 19-172 | RTBI Mode Register Initialization Steps | 19-203 |
| 19-173 | SGMII Interface Signal Configuration (4-Wire)..... | 19-205 |
| 19-174 | SGMII Mode Register Initialization Steps..... | 19-206 |
| 20-1 | I ² C Interface Signal Descriptions | 20-3 |
| 20-2 | I ² C Interface Signals—Detailed Signal Descriptions | 20-3 |
| 20-3 | I ² C Memory Map | 20-4 |
| 20-4 | I2CADR Field Descriptions | 20-5 |
| 20-5 | I2C FDR Field Descriptions | 20-6 |
| 20-6 | I2CCR Field Descriptions | 20-7 |
| 20-7 | I2CSR Field Descriptions | 20-8 |
| 20-8 | I2CDR Field Descriptions..... | 20-9 |
| 20-9 | I2CDFSRR Field Descriptions..... | 20-9 |
| 21-1 | DUART Signal Overview | 21-3 |
| 21-2 | DUART Signals—Detailed Signal Descriptions | 21-3 |
| 21-3 | DUART Register Summary | 21-5 |
| 21-4 | URBR Field Descriptions | 21-6 |
| 21-5 | UTHR Field Descriptions | 21-7 |
| 21-6 | UDMB Field Descriptions | 21-7 |
| 21-7 | UDLB Field Descriptions | 21-7 |
| 21-8 | Baud Rate Examples | 21-8 |
| 21-9 | UIER Field Descriptions..... | 21-9 |
| 21-10 | UIIR Field Descriptions | 21-10 |
| 21-11 | UIIR IID Bits Summary | 21-10 |
| 21-12 | UFCR Field Descriptions..... | 21-11 |
| 21-13 | UAFR Field Descriptions..... | 21-12 |
| 21-14 | ULCR Field Descriptions..... | 21-13 |
| 21-15 | Parity Selection Using ULCR[PEN], ULCR[SP], and ULCR[EPS] | 21-14 |
| 21-16 | UMCR Field Descriptions | 21-14 |
| 21-17 | ULSR Field Descriptions | 21-15 |
| 21-18 | UMSR Field Descriptions..... | 21-16 |
| 21-19 | USCR Field Descriptions..... | 21-17 |
| 21-20 | UDSR Field Descriptions..... | 21-17 |
| 21-21 | UDSR[TXRDY] Set Conditions | 21-18 |
| 21-22 | UDSR[TXRDY] Cleared Conditions..... | 21-18 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 21-23 | UDSR[RXRDY] Set Conditions..... | 21-18 |
| 21-24 | UDSR[RXRDY] Cleared..... | 21-18 |
| 22-1 | Signal Properties | 22-7 |
| 22-2 | Detailed Signal Descriptions..... | 22-7 |
| 22-3 | SPI Register Summary | 22-9 |
| 22-4 | SPMODE Field Descriptions | 22-9 |
| 22-5 | SPIE Field Descriptions | 22-12 |
| 22-6 | SPIM Field Descriptions | 22-13 |
| 22-7 | SPCOM Field Descriptions..... | 22-14 |
| 22-8 | SPI Transmit Data Hold Field Descriptions..... | 22-14 |
| 22-9 | SPI Receive Data Hold Field Descriptions | 22-15 |
| 23-1 | JTAG Test Signals Summary | 23-2 |
| 23-2 | JTAG Test—Detailed Signal Descriptions..... | 23-2 |
| 24-1 | GPIO External Signals—Detailed Signal Descriptions | 24-2 |
| 24-2 | GPIO Register Address Map..... | 24-2 |
| 24-3 | GPDIR Bit Settings | 24-3 |
| 24-4 | GPODR Bit Settings | 24-3 |
| 24-5 | GP _n DAT Bit Settings..... | 24-4 |
| 24-6 | GPIER Bit Settings | 24-4 |
| 24-7 | GPIMR Bit Settings | 24-5 |
| 24-8 | GPICR Bit Settings | 24-5 |
| 25-1 | TDM Signal Properties | 25-2 |
| 25-2 | TDM External Signals—Detailed Signal Descriptions | 25-2 |
| 25-3 | TDM Module Memory Map Summary..... | 25-7 |
| 25-4 | TDM SB Memory-Mapped Registers..... | 25-7 |
| 25-5 | TDMGIR Bit Descriptions..... | 25-9 |
| 25-6 | Pins Utilized for TDM Modules Based on RTS | 25-9 |
| 25-7 | TDMRIR Bit Descriptions | 25-10 |
| 25-8 | Received Data Delay for Receive Frame Sync | 25-11 |
| 25-9 | TDMTIR Bit Descriptions | 25-12 |
| 25-10 | Transmit Data Delay for Transmit Frame Sync | 25-13 |
| 25-11 | TDMTIR Bit Descriptions | 25-15 |
| 25-12 | TDMTFP Bit Descriptions..... | 25-16 |
| 25-13 | TDMRCEN _n Bit Descriptions | 25-17 |
| 25-14 | TDMTCEN _n Bit Descriptions | 25-18 |
| 25-15 | TDMTCMA _n Bit Descriptions | 25-19 |
| 25-16 | TDMRCR Bit Descriptions..... | 25-19 |
| 25-17 | TDMTCR Bit Descriptions | 25-20 |
| 25-18 | TDMRIER Bit Descriptions..... | 25-20 |
| 25-19 | TDMTIER Bit Descriptions..... | 25-21 |
| 25-20 | TDMRER Bit Descriptions | 25-23 |

Tables

| Table Number | Title | Page Number |
|-----------------|--|----------------|
| 25-21 | TDMTER Bit Descriptions | 25-25 |
| 25-22 | TDMRSR Bit Descriptions | 25-26 |
| 25-23 | TDMTSR Bit Descriptions | 25-27 |
| 25-24 | TDM AHB Memory-Mapped Registers | 25-28 |
| 25-25 | TDMRDREG Bit Descriptions | 25-29 |
| 25-26 | TDMTDREG Bit Descriptions | 25-29 |
| 25-27 | Clock Summary..... | 25-30 |
| 25-28 | Notes for Transmit Timing with Mask Register in Figure 25-39 | 25-42 |
| 25-29 | Notes for Receive Timing with Mask Register in Figure 25-40 | 25-43 |
| 25-30 | DMA Request Channels for the TDM | 25-46 |
| 25-31 | TDM Clock Memory Map | 25-52 |
| 25-32 | TDMCLK_DIV_VAL_RX Register | 25-53 |
| 25-33 | TDMCLK_DIV_VAL_TX Register | 25-53 |
| 25-34 | DMAC Register Summary | 25-56 |
| 25-35 | DMA Control Register (DMACR) Field Descriptions | 25-58 |
| 25-36 | DMAES Field Descriptions | 25-61 |
| 25-37 | DMAERQ Field Descriptions | 25-63 |
| 25-38 | DMAEEI Field Descriptions..... | 25-64 |
| 25-39 | DMASERQ Field Descriptions..... | 25-64 |
| 25-40 | DMACERQ Field Descriptions | 25-65 |
| 25-41 | DMASEEI Field Descriptions | 25-65 |
| 25-42 | DMACEEI Field Descriptions | 25-66 |
| 25-43 | DMACINT Field Descriptions..... | 25-67 |
| 25-44 | DMACERR Field Descriptions | 25-67 |
| 25-45 | DMASSRT Field Descriptions..... | 25-68 |
| 25-46 | DMACDNE Field Descriptions | 25-68 |
| 25-47 | DMAINT Field Descriptions | 25-69 |
| 25-48 | DMAERR Field Descriptions | 25-70 |
| 25-49 | DMAHRS Field Descriptions | 25-70 |
| 25-50 | DMAGPOR Field Descriptions | 25-71 |
| 25-51 | DCHPRI _n Field Descriptions..... | 25-72 |
| 25-52 | TCD 32-Bit Memory Structure | 25-73 |
| 25-53 | TCD Word 0 (TCD.saddr) Field Description..... | 25-73 |
| 25-54 | TCD Word 1 (TCD.{smod, ssize, dmod, dsize, soff}) Field Descriptions | 25-74 |
| 25-55 | TCD Word 2 (TCD.{smloe, dmloe, nbytes}) Description..... | 25-75 |
| 25-56 | TCD Word 3 (TCD.slast) Field Descriptions..... | 25-75 |
| 25-57 | TCD Word 4 (TCD.daddr) Field Description | 25-76 |
| 25-58 | TCD Word 5 (TCD.{citer, doff} Field Descriptions..... | 25-76 |
| 25-59 | TCD Word 6 (TCD.dlast_sga) Field Descriptions..... | 25-77 |
| 25-60 | TCD Word 7 (TCD.{biter, control/status}) Field Descriptions | 25-78 |

About This Book

This reference manual defines the functionality of the MPC8315E. It is written from the perspective of the MPC8315E, which is the superset device, and unless otherwise noted, the information applies also to the MPC8315, MPC8314E, and MPC8314. Note that the MPC8315 and MPC8314 do not support a security engine. Also, the MPC8314E and MPC8314 do not support dual SATA controllers.

The MPC8315E is a cost-effective, highly integrated host processor that addresses the requirements of several printing and imaging, consumer, and industrial applications, including main CPUs and I/O processors in printing systems, networking switches and line cards, wireless LANs (WLANs), network access servers (NAS), VPN routers, intelligent NIC, and industrial controllers. The MPC8315E extends the PowerQUICC family, adding higher CPU performance, additional functionality, and faster interfaces while addressing the requirements related to time-to-market, price, power consumption, and package size.

Audience

It is assumed that the reader understands operating systems, microprocessor system design, and the basic principles of RISC processing.

Organization

Following is a summary and a brief description of the major parts of this reference manual:

- [Chapter 1, “Overview,”](#) provides a high-level description of features and functionality of the MPC8315E integrated host processor. It describes the device, its interfaces, and the programming model. The functional operation of the device, with emphasis on peripheral functions, is also described.
- [Chapter 2, “Signal Descriptions,”](#) provides a listing of all the external signals, cross-references for signals that serve multiple functions, their functional blocks, and I/O states. Also, these signals are listed by alphabetical order.
- [Chapter 3, “Memory Map,”](#) describes the memory map of the device. An overview of the local address map is provided. Next, a complete listing of all memory-mapped registers is provided, with cross references to the sections detailing descriptions of each.
- [Chapter 4, “Reset, Clocking, and Initialization,”](#) describes the hard and soft resets, the power-on reset (POR) sequence, power-on reset configuration, clocking, and initialization of the device.
- [Chapter 5, “System Configuration,”](#) provides an overview of several functions that control the local access windows, system configuration, software watchdog, real time clock, periodic and general purpose timers, power management, protection, and general utilities.
- [Chapter 6, “Arbiter and Bus Monitor,”](#) provides an overview of the arbiter in the device. Also, it describes the configuration, control, and status registers of the arbiter.

- [Chapter 7, “e300 Processor Core Overview,”](#) provides an overview of the basic functionality of the processor core and briefly describes how the functional units interact.
- [Chapter 8, “Integrated Programmable Interrupt Controller \(IPIC\),”](#) describes the IPIC interrupt protocol, various types of interrupt sources controlled by the IPIC unit, and the IPIC registers with some programming guidelines. It also provides a definition of the external interrupt signals and their functions. In addition, the interrupt configuration, control, and status registers are described in this chapter.
- [Chapter 9, “DDR Memory Controller,”](#) describes the 32-bit DDR SDRAM memory controller of the device. This fully programmable controller supports most DDR memories available today, including both buffered and unbuffered devices. Dynamic power management and auto-precharge modes simplify memory system design.
- [Chapter 10, “Enhanced Local Bus Controller,”](#) describes the enhanced local bus controller (eLBC) of the device. It describes the external signals and the memory-mapped registers as well as a functional description of the general-purpose chip-select machine (GPCM), Flash control machine (FCM), and user-programmable machines (UPMs) of the eLBC. Also, it includes an initialization and applications information section with many specific examples of its use.
- [Chapter 11, “Sequencer,”](#) describes how the I/O sequencer (IOS) switches transactions among its ports, using a buffer pool to minimize blocking. It also provides address translation on outbound PCI transactions.
- [Chapter 12, “DMA/Messaging Unit,”](#) describes the four-channel high speed general-purpose DMA controller of the device. The channels share buffer space in the IOS to facilitate the gathering and sending of data. The DMA/messaging unit supports communication between two processors on different buses, for example, a local processor and a processor on a PCI bus. This communication unit operates with generic messages and doorbell registers. This block also provides a DMA controller that transfers blocks of data independent of the local processor or PCI hosts.
- [Chapter 13, “PCI Bus Interface,”](#) describes the PCI interface, which complies with the *PCI Local Bus Specification*, Rev. 2.3. This chapter provides a basic description of PCI bus operations. The specific emphasis is directed at how this device implements the PCI specification.
- [Chapter 14, “PCI Express Interface Controller,”](#) describes the PCI Express interface controller, which connects the CSB to the PCI Express bus, a 2.5 GHz serial interface that supports up to a x2 lane. As both a master (initiator) and a target device, the PCI Express interface is capable of high bandwidth data transfer and is designed to support the next generation I/O devices.
- [Chapter 15, “SATA Controller,”](#) describes the Serial ATA controller, a high-performance SATA solution incorporating some of the latest SATA-IO extensions. It is designed to operate in a system that supports command queuing and, in particular, switching scheme, based on a frame information structure (FIS) using port multipliers. Overviews on the command, transport, link, and PHY control layers are provided.
- [Chapter 16, “SerDes PHY,”](#) describes the block which includes the serializer/deserializer PHY, the protocol converter per protocol, the protocol mux, and the control registers and control logic. It supports two x1 PCI Express and two x1 SGMII.
- [Chapter 17, “Universal Serial Bus Interface,”](#) describes the universal serial bus (USB) interface. The USB DR module is a USB 2.0-compliant serial interface engine for implementing a USB

interface. The registers and data structures are based on the *Enhanced Host Controller Interface Specification for Universal Serial Bus* (EHCI) from Intel Corporation. The USB DR module can act as a host, as a device, or as an on-the-go (OTG) negotiable host/device on the USB bus.

- [Chapter 18, “Security Engine \(SEC\) 3.3,”](#) describes the SEC 3.3, which is designed to offload computationally intensive security functions, such as key generation and exchange, authentication, and bulk encryption from the e300 core of the device. It is optimized to process all the algorithms associated with IPsec. The SEC 3.0 (implemented in the device) is derived from integrated security cores found in other members of the PowerQUICC family, including SEC 1.0, the version implemented in the MPC8272/MPC8248. Note that the MPC8315 and MPC8314 do not support a security engine.
- [Chapter 19, “Enhanced Three-Speed Ethernet Controllers,”](#) describes the two enhanced three-speed Ethernet controllers on the device. These controllers provide 10/100/1Gb Ethernet support with a set of media-independent interface options including MII, RMII, GMII, and RTBI. The controllers provide two full-duplex FIFO interface modes and quality of service support. They are backward compatible with PowerQUICC III TSEC controllers.
- [Chapter 20, “I²C Interface,”](#) describes the inter-IC (IIC or I²C) bus controllers of the device. These synchronous, serial, bidirectional, multiple-master buses allow two-wire connection of devices, such as microcontrollers, EEPROMs, real-time clock devices, A/D converters, and LCDs. The device powers up in boot sequencer mode, which allows the I²C controllers to initialize configuration registers.
- [Chapter 21, “DUART,”](#) describes the (dual) universal asynchronous receiver/transmitters (UARTs) which feature a PC16552D-compatible programming model. These independent UARTs are provided specifically to support system debugging.
- [Chapter 22, “Serial Peripheral Interface,”](#) describes the MPC8315E serial peripheral interface (SPI) that allows the exchange of data between MPC83xx family devices. The SPI can also be used to communicate with peripheral devices such as EEPROMs, real-time clocks, A/D converters, and ISDN devices.
- [Chapter 23, “JTAG/Testing Support,”](#) describes the joint test action group (JTAG) interface of the MPC8315E to facilitate boundary-scan testing. The JTAG interface complies to the IEEE 1149.1 boundary-scan specification.
- [Chapter 24, “General Purpose I/O \(GPIO\),”](#) describes the general purpose I/O (GPIO) module in the MPC8315E device, including a definition of the external signals and functions they serve. Additionally, interrupt capabilities, pin description, and register settings are described.
- [Chapter 25, “Time Division Multiplexing \(TDM\) Interface,”](#) is a full-duplex, serial port typically used to transfer samples in a periodic manner. It consists of independent transmitter and receiver sections with independent clock generation and frame synchronization, and is composed of three identical and independent TDM modules. This chapter discusses the architecture, the programming model, the operating modes, and the initialization of the TDM.

Suggested Reading

This section lists additional reading that provides background for the information in this manual as well as general information about the architecture.

General Information

The following documentation, published by Morgan-Kaufmann Publishers, 340 Pine Street, Sixth Floor, San Francisco, CA, provides useful information about the PowerPC architecture and computer architecture in general:

- *The PowerPC Architecture: A Specification for a New Family of RISC Processors*, Second Edition, by International Business Machines, Inc.
- *Computer Architecture: A Quantitative Approach*, Third Edition, by John L. Hennessy and David A. Patterson.
- *Computer Organization and Design: The Hardware/Software Interface*, Second Edition, by David A. Patterson and John L. Hennessy.

Related Documentation

Freescale documentation is available from the sources listed on the back cover of this manual; the document order numbers are included in parentheses for ease in ordering:

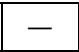
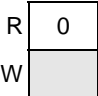

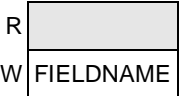
- *e300 Core Reference Manual*—This book provides a more detailed description of the e300 core.
- Reference manuals (formerly called user's manuals)—These books provide details about individual implementations.
- Addenda/errata to reference or user's manuals—Because some processors have follow-on devices, an addendum may be provided that describes the additional features and functionality changes. These addenda are intended for use with the corresponding reference or user's manuals.
- Hardware specifications—Hardware specifications provide specific data regarding bus timing, signal behavior, and AC, DC, and thermal characteristics, as well as other design considerations.
- Product briefs—Each device has a product brief that provides an overview of its features.
- Application notes—These short documents address specific design issues useful to programmers and engineers working with Freescale processors.

Additional literature is published as new processors become available. For a current list of documentation, refer to <http://www.freescale.com>.

Conventions

This document uses the following notational conventions:

| | |
|------------------|--|
| cleared/set | When a bit takes the value zero, it is said to be cleared; when it takes a value of one, it is said to be set. |
| mnemonics | Instruction mnemonics are shown in lowercase bold |
| <i>italics</i> | Italics indicate variable command parameters, for example, bcctrx |
| | Book titles in text are set in italics |
| | Internal signals are set in lowercase italics, for example, <u>core_int</u> |
| 0x0 | Prefix to denote hexadecimal number |
| 0b0 | Prefix to denote binary number |

| | |
|---|---|
| rA, rB | Instruction syntax used to identify a source GPR |
| rD | Instruction syntax used to identify a destination GPR |
| REG[FIELD] | Abbreviations for registers are shown in uppercase text. Specific bits, fields, or ranges appear in brackets. For example, MSR[LE] refers to the little-endian mode enable bit in the machine state register. |
| <i>x</i> | In some contexts, such as signal encodings, an unitalicized <i>x</i> indicates a don't care |
| <i>x</i> | An italicized <i>x</i> indicates an alphanumeric variable |
| <i>n</i> | An italicized <i>n</i> indicates a numeric variable |
| ¬ | NOT logical operator |
| & | AND logical operator |
| | OR logical operator |
| | Concatenation, for example TCR[WP] TCR[WPEXT] |
|  | Indicates a reserved bit field in a register. Although these bits can be written to as ones or zeros, they are always read as zeros. |
|  | Indicates a reserved bit field in a memory-mapped register. Although these bits can be written to as ones or zeros, they are always read as zeros. |
|  | Indicates a read-only bit field in a memory-mapped register. |
|  | Indicates a write-only bit field in a memory-mapped register. Although these bits can be written to as ones or zeros, they are always read as zeros. |

Signal Conventions

| | |
|-----------------------------|---|
| $\overline{\text{OVERBAR}}$ | An overbar indicates that a signal is active-low. |
| <i>lowercase_italics</i> | Lowercase italics is used to indicate internal signals. |
| lowercase_plaintext | Lowercase plain text is used to indicate signals that are used for configuration. |

Acronyms and Abbreviations

Table i contains acronyms and abbreviations used in this document.

Table i. Acronyms and Abbreviated Terms

| Term | Meaning |
|------|-----------------------------------|
| AESU | Advanced encryption standard unit |
| AFEU | ARC four execution unit |
| BD | Buffer descriptor |

Table i. Acronyms and Abbreviated Terms (continued)

| Term | Meaning |
|------------------|---|
| BIST | Built-in self test |
| CD | Collision detect |
| COL | Collision |
| CPM | Communication processor module |
| CRC | Cyclic redundancy check |
| CRS | Carrier sense |
| CSB | Coherent system bus |
| CSMA | Carrier-sense multiple access |
| DDR | Double data rate |
| DEU | Data encryption standard execution unit |
| DMA | Direct memory access |
| DRAM | Dynamic random access memory |
| DTLB | Data translation lookaside buffers |
| DUART | Dual universal asynchronous receiver/transmitter |
| EA | Effective address |
| ECC | Error checking and correction |
| EHCI | Enhanced host controller interface |
| EHPI | Enhanced host port interface |
| EPROM | Erasable programmable read-only memory |
| FS | Full-speed |
| FCS | Frame-check sequence |
| GMII | Gigabit media independent interface |
| GPCM | General-purpose chip-select machine |
| GPIO | General-purpose I/O |
| GPR | General-purpose register |
| GTM | General purpose timers |
| IAD | Internet access device |
| I ² C | Inter-integrated circuit |
| IEEE | Institute of Electrical and Electronics Engineers |
| IOS | I/O sequencer |
| IPG | Interpacket gap |
| ISDN | Integrated services digital network |
| ITLB | Instruction translation lookaside buffer |

Table i. Acronyms and Abbreviated Terms (continued)

| Term | Meaning |
|-------------|---|
| IU | Integer unit |
| JTAG | Joint Test Action Group |
| LALE | LBC external address latch enable |
| LBC | Local bus controller |
| LRU | Least recently used |
| LSB | Least-significant byte |
| lsb | Least-significant bit |
| LSU | Load/store unit |
| MAC | Multiply accumulate, media access control |
| MCP | Machine-check interrupt |
| MDI | Medium-dependent interface |
| MDEU | Message digest execution unit |
| MIB | Management information base |
| MII | Media independent interface |
| MMU | Memory management unit |
| MPH | Multi-port host |
| MSB | Most-significant byte |
| msb | Most-significant bit |
| OSI | Open systems interconnection |
| PCI | Peripheral component interconnect |
| PCS | Physical coding sublayer |
| PIC | Programmable interrupt controller |
| PIT | Periodic interval timer |
| PKEU | Public key execution unit |
| PMA | Physical medium attachment |
| PMD | Physical medium dependent |
| POR | Power-on reset |
| PRI | Primary rate interface |
| RGMII | Reduced gigabit media independent interface |
| RISC | Reduced instruction set computing |
| RMON | Remote monitoring |
| RMW | Read-modify-write |
| RNG | Random number generator |

Table i. Acronyms and Abbreviated Terms (continued)

| Term | Meaning |
|-------------|---|
| RTBI | Reduced ten-bit interface |
| RTC | Real time clock module |
| Rx | Receive |
| RxBD | Receive buffer descriptor |
| SCL | Serial clock |
| SDA | Serial data |
| SFD | Start frame delimiter |
| SI | Serial interface |
| SPI | Serial peripheral interface |
| SPR | Special-purpose register |
| SRAM | Static random access memory |
| TAP | Test access port |
| TBI | Ten-bit interface |
| TLB | Translation lookaside buffer |
| TSEC | Three-speed Ethernet controller |
| Tx | Transmit |
| TxBD | Transmit buffer descriptor |
| UART | Universal asynchronous receiver/transmitter |
| ULPI | USB low-pin count interface |
| UPM | User-programmable machine |
| USB | Universal serial bus |
| UTMI | USB transceiver macrocell interface |
| UTP | Unshielded twisted pair |
| WDT | Watchdog timer |
| ZBT | Zero bus turnaround |

Chapter 1

Overview

This chapter provides an overview of the MPC8315 PowerQUICC II Pro processor features, including a block diagram showing the major functional components. The MPC8315E is a cost-effective, low-power, highly integrated host processor that addresses the requirements of several storage, consumer, and industrial applications, including main CPUs and I/O processors in network access storage (NAS), voice over IP (VoIP) router/gateway, and intelligent wireless LAN (WLAN), set top boxes, industrial controllers, and wireless access points. The MPC8315E extends the PowerQUICC family, adding higher CPU performance, additional functionality, and faster interfaces while addressing the requirements related to time-to-market, price, power consumption, and package size. This manual is written from the perspective of the MPC8315E, and unless otherwise noted, the information applies also to the MPC8315, MPC8314, and MPC8314E. Note that the MPC8315 and MPC8314 do not support a security engine. In addition, the MPC8314E and MPC8314 do not support dual SATA controllers.

1.1 MPC8315 PowerQUICC II Pro Processor Overview

Figure 1-1 shows the major functional units within the MPC8315E. The e300 core in the MPC8315E, with its 16 Kbytes of instruction and 16 Kbytes of data cache, implements the user instruction set architecture and provides hardware and software debugging support. In addition, the MPC8315E offers dual PCI Express controllers, dual three-speed 10, 100, 1000 Mbps Ethernet controllers (eTSECs), a DDR1/DDR2 SDRAM memory controller, a dedicated security engine, a TDM interface, a 32-bit PCI-2.3 controller with PME, a USB 2.0 host and device controller with an on-chip high-speed PHY, dual serial ATA (SATA) controllers, a SerDes block, an enhanced local bus controller (eLBC), an integrated programmable interrupt controller (IPIC), a 4-channel DMA controller, an I²C controller, dual UART (DUART), GPIOs, general purpose timers, and an SPI controller. The high level of integration in the MPC8315E helps simplify board design and offers significant bandwidth and performance.




Figure 1-1. MPC8315E Block Diagram

The major features of this device are as follows:

- e300c3 Power Architecture® processor core
 - Enhanced version of the MPC603e core
 - High-performance, superscalar processor core with a four-stage pipeline and low-interrupt latency times
 - Floating-point, dual integer units, load/store, system register, and branch processing units
 - 16-Kbyte instruction cache and 16-Kbyte data cache with lockable capabilities
 - Capable of completing two MACs every three cycles
 - Dynamic power management
 - Enhanced hardware program debug features
 - Software-compatible with Freescale processor families implementing Power Architecture technology
 - Separate PLL that is clocked by the system bus clock
 - Performance monitor
- Security engine with the following features:
 - AESA—Advanced Encryption Standard Hardware Accelerator
 - Implements the Rijndael symmetric key cipher
 - Modes providing data confidentiality
 - Modes providing data authentication: 128-, 192-, and 256-bit key lengths (only 128-bit keys in XCBC-MAC)
 - ICV checking in CCM, GCM, CMAC, and XCBC-MAC mode
 - XOR operations on 2–6 sources for RAID
 - CRCA—Cyclical Redundancy Check Hardware Accelerator

- Implements CRC32C as required for iSCSI header and payload checksums, CRC32 as required for IEEE Std 802™ packets, as well as for programmable 32-bit CRC polynomials
- ICV checking
- DESA—Data Encryption Standard Hardware Accelerator
 - DES, 3DES
 - Two key (K1, K2, K1) or three key (K1, K2, K3)
 - ECB, CBC, CFB-64, and OFB-64 modes for both DES and 3DES
- MDHA—Message Digest Hardware Accelerator
 - SHA with 160-, 224-, 256-, 384-, and 512-bit message digest
 - MD5 with 128-bit message digest
 - HMAC computation with either message digest algorithm
 - SSL MAC computation
 - ICV checking
- PKHA—Public Key Hardware Accelerator that supports the following:
 - RSA and Diffie-Hellman
 - Programmable field size up to 4096 bits (increased from 2048)
 - Elliptic curve cryptography
 - F2m and Fp modes
 - Programmable field size up to 1023 bits (increased from 511)
 - Run time equalization to protect against timing and power attacks
- RNGB—Random Number Generator Hardware Accelerator
 - Combines a True Random Number Generator (TRNG) and a NIST-approved Pseudo-Random Number Generator (PRNG) (as described in Annex C of FIPS140-2 and ANSI X9.62)
- Four-channel operation, where each channel:
 - Queue of commands (descriptor pointers) to be executed
 - Dynamic arbitration for needed crypto-execution units
 - Manages up to two execution units (one ciphering and one hashing), and configures for any required data transfers from one to another
 - Flow-control management of buffer FIFOs on the inputs and outputs of execution units
 - Supports scatter/gather of input and output data, enabling concatenation of multiple segments of memory when reading or writing data
 - Masters data bursts on 32-byte boundaries to optimize throughput
- Master and slave interfaces, with DMA capability
 - 32-bit address, 64-bit data
 - Up to 133 MHz operation
 - Master interface allows pipelined requests
 - DMA data blocks can start and end on any byte boundary

- DDR SDRAM memory controller
 - Programmable timing supporting both DDR1 and DDR2 SDRAM
 - 16-/32-bit data interface, up to 266-MHz data rate
 - The following SDRAM configurations are supported:
 - Up to two physical banks (chip selects), each bank up to 1 Gbyte independently addressable
 - 64-Mbit to 2-Gbit (for DDR1) and to 4-Gbit (for DDR2) devices with x8/x16/x32 data ports (no direct x4 support)
 - One 16-bit device or two 8-bit devices on a 16-bit bus, or one 32-bit device or two 16-bit devices or four 8-bit devices on a 32-bit bus
 - Support for up to 8 simultaneous open pages for DDR1 (up to 16 pages for DDR2)
 - Dual chip selects
 - Sleep-mode support for SDRAM self refresh
 - Supports auto refresh
 - On-the-fly power management using CKE
 - Registered DIMM support
 - 2.5-V SSTL2 compatible I/O for DDR1, 1.8-V SSTL_18 compatible I/O for DDR2
- Two enhanced three-speed Ethernet controllers (eTSECs)
 - Backward compatible with MPC8349E (PowerQUICC II Pro) TSEC
 - Three-speed support (10/100/1000 Mbps)
 - Two SGMII interfaces, two RGMII/RTBI/MII/RMII interfaces
 - Two controllers designed to comply with IEEE Std 802.3®, 802.3u®, 802.3x®, 802.3z®, 802.3ac®, and 802.3ab®
 - Support for IEEE Std 1588™
 - Support for Wake-on-Magic Packet™, a method to bring the device from standby to full operating mode
 - TCP/IP acceleration and QoS features available
 - IP v4 and IP v6 header recognition on receive
 - IP v4 header checksum verification and generation
 - TCP and UDP checksum verification and generation
 - Per-packet configurable acceleration
 - Recognition of VLAN, stacked (queue in queue) VLAN, 802.2, PPPoE session, MPLS stacks, and ESP/AH IP-security headers
 - Transmission from up to eight physical queues
 - Reception to up to eight physical queues
 - Full- and half-duplex Ethernet support (1000 Mbps supports only full-duplex):
 - IEEE Std. 802.3 full-duplex flow control (automatic PAUSE frame generation or software-programmed PAUSE frame generation and recognition)

- Programmable maximum frame length supports jumbo frames (up to 9.6 Kbytes) and IEEE Std 802.1Q virtual local area network (VLAN) tags and priority
- VLAN insertion and deletion
 - Per-frame VLAN control word or default VLAN for each eTSEC
 - Extracted VLAN control word passed to software separately
- Retransmission following a collision
- CRC generation and verification of inbound/outbound packets
- Programmable Ethernet preamble insertion and extraction of up to 7 bytes
- MAC address recognition:
 - Exact match on primary and virtual 48-bit unicast addresses
 - VRRP and HSRP support for seamless router fail-over
 - Up to 16 exact-match MAC addresses supported
 - Broadcast address (accept/reject)
 - Hash table match on up to 512 multicast addresses
 - Promiscuous mode
 - 10-K packet buffers to enable check-summing for jumbo packets
- Buffer descriptors backward compatible with MPC8260 and MPC860T 10/100 Ethernet programming models
- RMON statistics support
- MII management interface for control and status
- SerDes block with two lanes
 - Support for two x1 PCI Express and two x1 SGMII
 - Link-layer interfaces to IP controller
 - SerDes power-down/reset state machine for cold (power-on) or warm (software-initiated) reset of SerDes, PHY, and controllers
- PCI interface
 - Designed to comply with *PCI Local Bus Specification, Revision 2.3*
 - 32-bit PCI interface operating at up to 66 MHz
 - PCI 3.3-V compatible
 - Not 5-V compatible
 - Support for host and agent modes
 - Support for PCI-to-memory and memory-to-PCI streaming
 - Memory prefetching of PCI read accesses and support for delayed read transactions
 - Support for posting of processor-to-PCI and PCI-to-memory writes
 - On-chip arbitration, supporting three masters on PCI
 - Arbiter support for two-level priority request/grant signal pairs
 - Support for accesses to all PCI address spaces
 - Support for parity

- Selectable hardware-enforced coherency
- Address translation units for address mapping between host and peripheral
- Mapping from an external 32-/64-bit address space to the internal 32-bit local space
- Support for dual address cycle (DAC) (as a target only)
- Internal configuration registers accessible from PCI
- Support for PCI Power Management 1.2
- Selectable snooping for inbound transactions
- Four outbound Translation Address Windows
 - Support for mapping 32-bit internal local memory space to an external 32-bit PCI address space and translating that address within the PCI space
- Four inbound Translation Address Windows corresponding to defined PCI BARs
 - The first BAR is 32 bits and dedicated to on-chip register access
 - The second BAR is 32 bits for general use
 - The remaining two BARs may be 32 or 64 bits and are also for general use
- PCI Express
 - Supports two interfaces supporting x1 width
 - Compatible with the *PCI Express 1.0a Specification*
 - Selectable operation as root complex or endpoint
 - 32- and 64-bit addressing
 - 128-byte maximum payload size
 - Virtual channel 0 only
 - Traffic class 0–7
 - Full 64-bit decode with 32-bit wide windows
 - Four outbound Translation Address Windows
 - Support for mapping 32-bit internal local memory space to an external 32- or 64-bit address space and translating that address within the PCI Express space
 - Four inbound Translation Address Windows corresponding to defined PCI Express BARs
 - The first BAR is 32-bits can be programmed to use on-chip register access
 - The second BAR is 32-bits for general use
 - The remaining two BARs may be 32- or 64-bits and are also for general use
- Serial ATA (SATA) controller
 - Two ports
 - Controller is designed to be compliant to the *Serial ATA 2.5 Specification*
 - 1.5 and 3.0 Gbps operation for SATA and eSATA
 - Support for Gen1i, Gen1m, Gen2i, and Gen2m electricals per *Serial ATA 2.5 Specification*
 - Native command queuing
 - Staggered spin-up
 - Port multiplier support

- Hot plug including asynchronous signal recovery
- Support in each controller for 16-entry command queue
- On-chip PHY
- Universal serial bus (USB) dual-role controller
 - Designed to comply with *Universal Serial Bus Revision 2.0 Specification*
 - Supports operation as a stand-alone USB host controller
 - Supports USB root hub with one downstream-facing port
 - Enhanced host controller interface (EHCI) compatible
 - Supports operation as a stand-alone USB device
 - Supports one upstream-facing port
 - Supports three programmable bidirectional USB endpoints
 - Supports high-speed (480-Mbps), full-speed (12-Mbps), and low-speed (1.5-Mbps) operations. Low speed is only supported in host mode.
 - Host mode direct connect of full- and low-speed devices
 - Supports USB on-the-go mode when using an external ULPI (UTMI+ low-pin interface) PHY, which includes both device and host functionality
 - On-chip USB-2.0 full-/high-speed PHY with ULPI (UTMI+ low-pin interface) and serial interface
 - Host and device support
- Enhanced local bus controller (eLBC)
 - Multiplexed 26-bit address and 8-/16-bit data operating at up to 66 MHz
 - Four chip selects supporting four external slaves
 - Variable memory block sizes (32 Kbytes to 4 Gbytes in FCM mode, 32 Kbytes to 64 Mbytes in UPM mode, and 32 Kbytes to 128 Mbytes in GPCM mode)
 - Supports boot from NOR Flash and NAND Flash
 - Supports programmable clock ratio dividers
 - Up to eight-beat burst transfers
 - 16- and 8-bit ports
 - Three protocol engines available on a per chip select basis:
 - General-purpose chip select machine (GPCM)
 - Three user programmable machines (UPMs)
 - NAND Flash control machine (FCM)
 - Default boot ROM chip select with configurable bus width (8 or 16)
- Integrated programmable interrupt controller (IPIC)
 - Functional and programming compatibility with the MPC8260 interrupt controller
 - Support for external and internal discrete interrupt sources
 - Programmable highest priority request
 - Six groups of interrupts with programmable priority

- External and internal interrupts directed to host processor
- Redirects interrupts to external $\overline{\text{PCI_INTA}}$ signal when in core disable mode
- Supports MSI functionality for PCI Express
- Unique vector number for each interrupt source
- I²C interface
 - Two-wire interface
 - Multiple-master support
 - Master or slave I²C mode support
 - On-chip digital filtering rejects spikes on the bus
 - System initialization data is optionally loaded from I²C EPROM by boot sequencer embedded hardware
- Time division multiplexing (TDM) interface
 - Asynchronous receive and transmit, each having one data line, one clock line, and one frame sync line
 - Frame sync line and/or clock line can be shared between receive transmit within a single TDM
 - Independent or shared transmit and receive sections with separate or shared clocks and frame syncs
 - Frame sync and data signals can be programmed as active low or active high
 - TDM (network) mode operation allowing multiple devices to share the port up to 128 time slots
 - Single channel mode operation using frame sync
 - Each channel can be programmed to be active or inactive
 - Program options for frame sync and clock generation
 - Selectable delay (0–3 bits) between the frame sync signal and the beginning of the frame
 - TDM transmitter sync signal (TxTFS) and TDM transmitter clock signal (TxTCK) can be configured as either input or output
 - End of frame interrupt
 - Programmable internal clock divider
 - Programmable word length (8 or 16 bits)
 - Programmable to MSB or LSB first
 - A-law/ μ -law hardware conversion is supported for 8-bit channels
 - Glueless interface to E1/T1 frames and MVIP, SCAS, and H.110 buses
 - Loopback mode
 - Support for the DMA engine with the following features:
 - Two DMA channels
 - All data movement via dual-address transfers: read from source, write to destination
 - Transfer control descriptor (TCD) organized to support two-deep, nested transfer operations
 - Channel activation via one of three methods (for all three methods, one activation per execution of the minor loop is required):

- Explicit software initiation
 - Initiation via a channel-to-channel linking mechanism for continuous transfers (independent channel linking at end of minor loop and/or major loop)
 - Peripheral initiated hardware requests (one per channel)
 - Support for fixed-priority and round-robin channel arbitration
 - Channel completion reported via optional interrupt requests
- Support for scatter/gather DMA processing
- I/O sequencer
- DMA (Direct memory access) controller
 - Four independent fully programmable DMA channels
 - Concurrent execution across multiple channels with programmable bandwidth control
 - Handshaking (external control) signals supported for three channels: $\overline{\text{DMA_DREQ}}[0:1]$, $\overline{\text{DMA_DACK}}[0:1]$, $\overline{\text{DMA_DDONE}}[0:1]$
 - Misaligned transfer capability for source/destination address
 - Data chaining and direct mode
 - Interrupt on completed segment, error and chain
- DUART
 - Two 4-wire interfaces (RxD, TxD, $\overline{\text{RTS}}$, $\overline{\text{CTS}}$)
 - Programming model compatible with the original 16450 UART and the PC16550D
- Serial peripheral interface (SPI)
 - Master or slave support
- Power management controller (PMC)
 - Provides power management when the device is used in both host and agent modes
 - Low power standby mode
 - Supports PCI Power Management 1.2 D0, D1, D2, D3hot, D3warm, and D3cold states
 - On-chip split power supply controlled through external power switch for minimum standby power
 - Supports PME generation in PCI agent mode and PME detection in PCI host mode in D3hot (not in D3warm)
 - Supports wake-up from Ethernet Magic Packet (one eTSEC only), GPIO, general purpose timer, and IRQ input assertion; as well as wake-up from USB, both eTSECs, and PCI in D3 hot
 - Supports MPC8349E backward-compatibility mode
- Parallel I/O
 - General-purpose I/O (GPIO)
 - 32 parallel I/O pins multiplexed on various chip interfaces
 - Supports eight non-multiplexed I/O pins
 - Supports 24 additional GPIO pins multiplexed with eTSEC, TDM, SPI, and DMA
 - Open drain capability

- Interrupt capability
- System timers
 - Periodic interrupt timer
 - Real-time clock
 - Software watchdog timer
 - Four general-purpose timers
- IEEE Std. 1149.1™ compliant JTAG boundary scan
- Integrated PCI bus and SDRAM clock generation

1.2 MPC8315E Architecture Overview

The following sections describe the major functional units of this device.

1.2.1 Power Architecture Core

The device contains the e300c3 Power Architecture processor core, which is an enhanced version of the MPC603e core (used in previous generations of PowerQUICC II processors). Enhancements include integrated parity checking, dual integer units, and other performance-enhancing features. The e300 core is upward software-compatible with existing MPC603e core-based products.

For detailed information regarding the processor core refer to the following:

- The *e300 Power Architecture™ Core Family Reference Manual* (chapters describing the programming model, cache model, memory management model, exception model, and instruction timing) (Document No. E300CORERM)
- The *Programming Environments Manual for 32-Bit Implementations of the PowerPC™ Architecture* (Document No. MPCFPE32B)

The e300 core is a low-power implementation of the family of microprocessors that implements Power Architecture technology. The core implements the 32-bit portion of the architecture, which provides 32-bit effective addresses, integer data types of 8, 16, and 32 bits, and floating-point data types of 32 and 64 bits.

The core is a superscalar processor that can issue three instructions (two plus a branch) and completes and retires as many as two instructions per clock cycle. Instructions can execute out of order for increased performance; however, the core makes completion appear sequential.

The e300c3 core integrates six execution units—two integer units (IU1 and IU2) with full multiply and divides, a floating-point unit (FPU), a branch processing unit (BPU) with static branch prediction, a load/store unit (LSU) for data transfers, a performance monitor, and a system register unit (SRU). The ability to execute five instructions in parallel and the use of simple instructions with rapid execution times yield high efficiency and throughput. Most integer instructions execute in one clock cycle; two integer instructions may be executed at the same time with the dual integer units. The FPU is pipelined so a single-precision multiply-add instruction can be issued and completed every clock cycle.

The e300c3 core provides independent on-chip, 16-Kbyte, eight-way set-associative, physically addressed instruction and data caches with parity and integrated way lock capabilities. The processor also features independent on-chip instruction and data memory management units (MMUs). The MMUs contain

64-entry, two-way set-associative, data and instruction translation lookaside buffers (DTLB and ITLB) that provide support for demand-paged virtual memory address translation. The caches use a pseudo least recently used (PLRU) replacement algorithm; the TLBs use a least recently used (LRU) replacement algorithm. The processor also supports block address translation through the use of two independent instruction and data block address translation (IBAT and DBAT) arrays of eight entries each. Effective addresses are compared simultaneously with all eight entries in the BAT array during block translation. In accordance with the architecture, if an effective address hits in both the TLB and BAT array, the BAT translation takes priority.

As an added feature to the e300 core, the device can lock the contents of three of the four ways in the instruction and data cache (or an entire cache). For example, this allows embedded applications to lock interrupt routines or other important (time-sensitive) instruction sequences into the instruction cache. It allows data to be locked into the data cache, which may be important to code that must have deterministic execution.

The e300 core has high-performance 64-bit data bus and 32-bit address bus interfaces to the rest of the device. The e300 core supports single-beat and burst data transfers for memory accesses, and memory-mapped I/O operations.

[Figure 1-2](#) provides a block diagram of the e300 core that shows how the execution units (IU1, IU2, FPU, BPU, LSU, and SRU) operate independently and in parallel. Note that this is a conceptual diagram and does not attempt to show how these features are physically implemented on the chip.




Figure 1-2. MPC8315E Integrated e300c3 Core Block Diagram

1.2.2 Security Engine

A hardware encryption block is also integrated in the device. It supports many encryption algorithms allowing for high performance data encryption and authentication as required in today's SoHo/RoBo routers. The encryption block is compatible with the corresponding block in the MPC8280.

The SEC 3.3 is designed to off-load computationally intensive security functions, such as key generation and exchange, authentication, and bulk encryption from the processor core of the device. It is optimized to process cryptographic algorithms associated with IPsec, IKE, SSL/TLS, iSCSI, SRTP, 802.11i, 802.16 (WiMAX), and 802.1ae (MACSec).

The security engine includes six different execution units (EUs). Where data flows in and out of an EU, each has buffer FIFOs of at least 256 bytes. A block diagram of the security engine's internal architecture is shown in Figure 1-3.

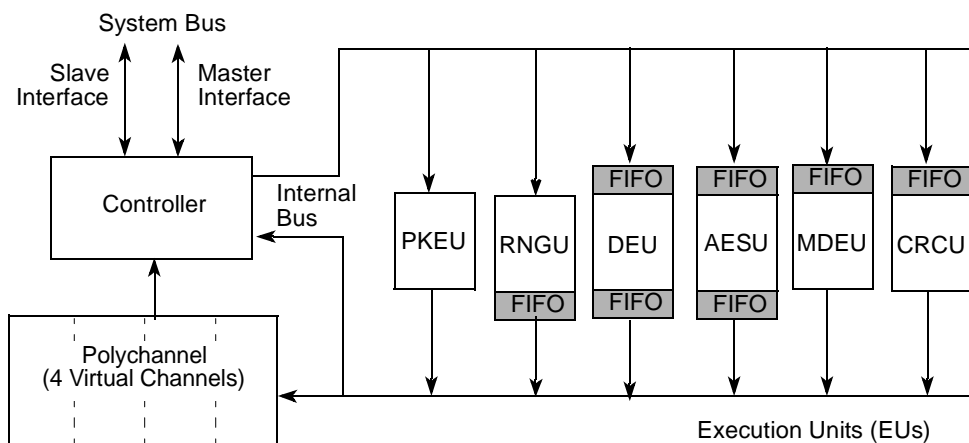


Figure 1-3. Integrated Security Engine Functional Blocks

1.2.3 DDR Memory Controller

This fully programmable DDR SDRAM controller supports most JEDEC standard x8 or x16 DDR1 or DDR2 memories available today, including buffered and unbuffered DIMMs. However, mixing nonregistered and registered DIMMs in the same system is not supported. Dynamic power management and auto-precharge modes simplify memory system design.

The DDR memory controller includes the following features:

- Support for DDR1 and DDR2 SDRAM
- 16- or 32-bit SDRAM data bus
- Programmable settings for meeting all SDRAM timing parameters
- Many different SDRAM configurations supported
 - Support for as many as two physical banks (chip selects), each bank independently addressable
 - Support for 64-Mbit to 1-Gbit devices with x8/x16/x32 data ports. Some 2-Gbit devices are supported depending on the internal device configuration.
 - Support for unbuffered and registered DIMMs

- Support for data mask signals and read-modify-write operations for sub-double word writes
- Four-entry input request queue
- Open page management (dedicated entry for each sub-bank)
- Memory controller clock frequency of two times the SDRAM clock with support for sleep power management mode

1.2.4 Dual Enhanced Three-Speed Ethernet Controllers

The MPC8315E has two on-chip enhanced three-speed Ethernet controllers. The eTSECs incorporate a media access control (MAC) sublayer that supports 10- and 100-Mbps and 1-Gbps Ethernet/IEEE Std. 802.3 networks with MII, RGMII, SGMII, and RTBI physical interfaces. The eTSECs include 2-Kbyte receive and 10-Kbyte transmit FIFOs and DMA functions. They also support IEEE Std. 1588.

The buffer descriptors are based on the MPC8260 and MPC860T 10/100 Ethernet programming models. Each eTSEC can emulate a PowerQUICC III TSEC, allowing existing driver software to be re-used with minimal change.

The MPC8315E eTSECs support programmable CRC generation and checking, RMON statistics, and jumbo frames of up to 9.6 Kbytes. Frame headers and buffer descriptors can be forced into the L2 cache to speed classification or other frame processing.

Each eTSEC provides hardware support for accelerating TCP/IP packet transmission and reception. By default, TCP/IP acceleration is not enabled, and the eTSEC processes frames as pure Ethernet frames.

TCP/IP acceleration can be performed at a number of levels. The eTSEC can parse frames at layer 2 of the stack only (Ethernet headers and switching headers), layers 2 to 3 (including IP v4 or IP v6), or layers 2 to 4 (including TCP and UDP).

On receive, the eTSEC provides protocol header recognition, header verification (IP v4 header checksum verification), and TCP/UDP payload checksum verification including verification of associated pseudo-header checksums. On transmit, the eTSEC provides IP v4 and TCP/UDP header checksum generation. The eTSEC does not checksum transmitted packets with IP header options or IP fragments.

To provide for quality of service, transmission from up to eight queues is supported with priority-based queue selection. Arbitration is a modified weighted round-robin queue selection with fair bandwidth allocation.

On receive, packets may be distributed to any of the 64 virtual receive queues overlaid onto the 8 physical receive queues. A table-oriented queue filing strategy is provided based on 16 header fields or flags. Frame rejection is supported for filtering applications.

Filing can be based on Ethernet, IP, and TCP/UDP properties, including VLAN fields, Ether-type, IP protocol type, IP TOS or differentiated services, IP source and destination addresses, TCP/UDP port numbers, or user-defined bit fields.

1.2.5 SerDes PHY

The SerDes PHY block includes the SerDes PHY, the protocol converter per protocol, the protocol mux, and the control registers and control logic.

The SerDes PHY block has the following features:

- Support for two x1 PCI Express and two x1 SGMII
- Link-layer interfaces to IP controller
- Memory-mapped registers with 256-byte address region
- SerDes power-down/reset state machine for cold (power-on) or warm (software-initiated) reset of SerDes, PHY, and controllers

The SerDes PHY block supports the following modes of operation:

- Two lanes running x1 SGMII at 1.25 Gbps
- Two lanes running x1 PCI Express at 2.5 Gbps

1.2.6 PCI Controller

The 32-bit PCI controller is compatible with the *PCI Local Bus Specification, Rev. 2.3*. The PCI interface can function as a host bridge interface. The PCI interface can optionally function as an agent device. The PCI controller supports 32-bit addressing and 32-bit data buses.

As a host, the device supports read and write operations to the PCI memory space, the PCI I/O space, and the PCI configuration space. Also, the device can generate PCI special-cycle and interrupt acknowledge commands. As an agent, the device supports read and write operations to system memory, as well as PCI configuration space and the on-chip memory-mapped configuration space.

The device PCI controller includes the following distinctive features:

- Address stepping on configuration transactions
- Fast back-to-back transactions
- Data streaming
- Supports mapping from an external 32- or 64-bit address space to the internal 32-bit local space
- Supports dual address cycle (DAC) 64-bit addressing mode as target only
- When in host mode, the PCI controller supports external signal isolation, thus enabling power shut off to external devices
- Supports PCI Power Management 1.2
- Supports PME generation (agent) and Wake on PME in D3 hot mode

1.2.6.1 PCI Bus Arbitration Unit

The PCI controller contains a PCI bus arbitration unit, which eliminates the need for an external unit, thus lowering system complexity and cost. It has the following features:

- Supports three $\overline{\text{REQ}}/\overline{\text{GNT}}$ signal pairs, thus supporting three external masters. The device PCI controller is the fourth member of the arbitration pool.
- The bus arbitration unit allows fairness as well as a priority mechanism.
- A two-level round-robin scheme is used in which each device can be programmed within a pool of a high- or low-priority arbitration. One member of the low-priority pool is promoted to the high-priority pool. As soon as it is granted the bus, it returns to the low-priority pool.

Overview

- The unit can be disabled to allow a remote arbitration unit to be used.
- The unit can be isolated to allow power shut off of external devices.

The Serial ATA controller is a high-performance SATA solution incorporating some of the latest SATA-IO extensions. The SATA may also be referred to as a host bus adapter (HBA). The SATA controller is designed to operate in a system that supports command queuing and, in particular, a switching scheme based on a frame information structure (FIS) using port multipliers.

The SATA controller has the following features:

- Designed to be compliant to the *Serial ATA 2.5 Specification*
- Supports speeds: 1.5 Gbps (first-generation SATA), 3 Gbps (second-generation SATA and eSATA)
- Supports advanced technology attachment packet interface (ATAPI) devices
- Contains high-speed descriptor-based DMA controller
- Supports native command queuing (NCQ) commands
- Supports port multiplier operation
- Supports hot plug including asynchronous signal recovery

There are four layers in the SATA architecture: application, transport, link, and PHY. The application layer is responsible for overall ATA command execution, including controlling command block register accesses. The transport layer is responsible for placing control information and data to be transferred between the host and device in a packet/frame, known as a frame information structure (FIS). The link layer is responsible for taking data from the constructed frames, inserting control characters, and moving data to the PHY layer. The PHY layer is responsible for 8B/10B encoding/decoding, then transmitting and receiving the encoded information as a serial data stream on the wire.

1.2.7 Universal Serial Bus (USB) 2.0

The USB 2.0 controller offers operation as a host or device. The USB controller provides point-to-point connectivity, which complies with the *Universal Serial Bus Revision 2.0 Specification*. The USB controllers can be configured to operate as a stand-alone host or stand-alone device. See [Figure 1-4](#) for more information.

The host and device functions are both configured to support the following four types of USB transfers:

- Bulk
- Control
- Interrupt
- Isochronous

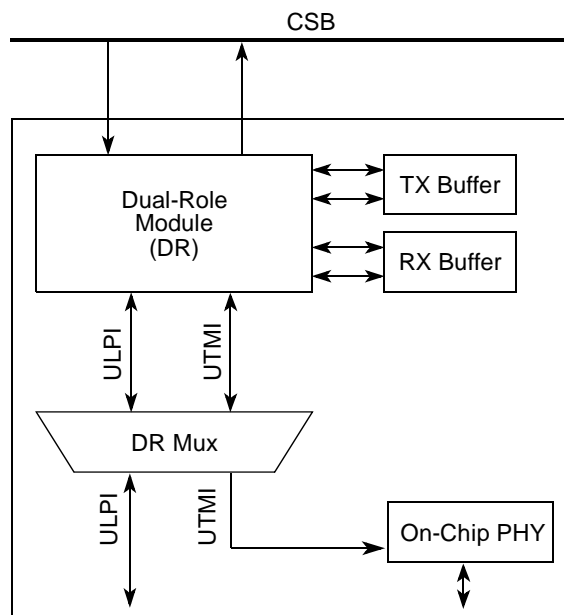


Figure 1-4. USB Controllers Port Configuration

1.2.7.1 USB Dual-Role Controller

- Designed to comply with *Universal Serial Bus Revision 2.0 Specification*
- Supports operation as a stand-alone USB host controller
 - Supports USB root hub with one downstream-facing port
 - Enhanced host controller interface (EHCI) compatible
- Supports operation as a stand-alone USB device
 - Supports one upstream-facing port
 - Supports three programmable bidirectional USB endpoints
- Supports high-speed (480-Mbps), full-speed (12-Mbps), and low-speed (1.5-Mbps) operations. Low speed is only supported in host mode.
- Host mode direct connect of full-speed and low-speed devices
- Supports USB on-the-go mode when using an external ULPI (UTMI+ low-pin interface) PHY, which includes both device and host functionality
- On-chip USB-2.0 full-/high-speed PHY with ULPI (UTMI+ low-pin interface) and serial interface
- Host and device support

1.2.8 Enhanced Local Bus Controller (eLBC)

The main component of the enhanced local bus controller (eLBC) is its memory controller, which provides a seamless interface to many types of memory devices and peripherals. The memory controller is responsible for controlling four memory banks shared by a NAND Flash control machine (FCM), a general-purpose chip-select machine (GPCM), and up to three user-programmable machines (UPMs). As such, it supports a minimal glue logic interface to SRAM, EPROM, NOR Flash EPROM, NAND Flash

EPROM, Flash EPROM, burstable RAM, and other peripherals. The eLBC external address latch enable (LALE) signal allows multiplexing of addresses with data signals to reduce the device pin count.

The enhanced local bus controller also includes a number of data checking and protection features such as data parity generation and checking, write protection, and a bus monitor to ensure that each bus cycle is terminated within a user-specified period.

The main features of the enhanced local bus controller (eLBC) are as follows:

- Memory controller with four memory banks (chip selects)
 - 32-bit address decoding with mask
 - Variable memory block sizes (32 Kbytes to 2 Gbytes in FCM mode, 32 Kbytes to 64 Mbytes in UPM mode, and 32 Kbytes to 128 Mbytes in GPCM mode)
 - Selection of control signal generation on a per-bank basis
 - Data buffer controls activated on a per-bank basis
 - Up to 256-byte bursts, arbitrarily aligned
 - Automatic segmentation of large transactions into memory accesses optimized for bus width and addressing capability
 - Odd/even parity checking including read-modify-write (RMW) parity for single accesses
 - Write-protection capability
 - Atomic operation
- General-purpose chip-select machine (GPCM)
 - Compatible with SRAM, EPROM, NOR Flash EEPROM, FEPRM, and peripherals
 - Global (boot) chip-select available at system reset
 - Boot chip-select support for 8- and 16-bit devices
 - Minimum three-clock access to external devices
 - Two byte-write-enable signals ($\overline{\text{LWE}}[0:1]$)
 - Output enable signal ($\overline{\text{LOE}}$)
 - External access termination signal ($\overline{\text{LGTA}}$)
- NAND Flash control machine (FCM)
 - Compatible with small (512 + 16 bytes) and large (2048 + 64 bytes) page parallel NAND Flash EEPROM
 - Global (boot) chip-select available at system reset, with 4-Kbyte boot block buffer for execute-in-place boot loading
 - Boot chip-select support for 8-bit devices
 - Dual 2-Kbyte/eight 512-byte buffers allow simultaneous data transfer during Flash reads and programming
 - Interrupt-driven block transfer for reads and writes
 - Programmable command and data transfer sequences of up to eight steps supported
 - Generic command and address registers support proprietary Flash interfaces
 - Block write locking to ensure system security and integrity

- Three user-programmable machines (UPMs)
 - Programmable-array-based machine controls external signal timing with a granularity of up to one quarter of an external bus clock period
 - User-specified control-signal patterns run when an internal master requests a single-beat or burst read or write access
 - UPM refresh timer runs a user-specified control signal pattern to support refresh
 - User-specified control-signal patterns can be initiated by software
 - Each UPM can be defined to support DRAM devices with depths of 64, 128, 256, and 512 Kbytes, and 1, 2, 4, 8, 16, 32, 64, 128, and 256 Mbytes
 - Support for 8- and 16-bit devices
 - Page mode support for successive transfers within a burst
 - Internal address multiplexing supporting 64-, 128-, 256-, and 512-Kbyte, and 1-, 2-, 4-, 8-, 16-, 32-, 64-, 128-, and 256-Mbyte page banks
- Optional monitoring of transfers between local bus internal masters and local bus slaves (local bus error reporting)

1.2.9 Integrated Programmable Interrupt Controller (IPIC)

The IPIC implements the necessary functions to provide a flexible solution for general-purpose interrupt control. The IPIC includes the following features:

- Functional and programming models are compatible with the MPC8260 interrupt controller
- Support for external and internal discrete interrupt sources
- Support for one external (optional) and seven internal machine checkstop interrupt sources
- Programmable highest priority request
- Two programmable priority mixed groups of four on-chip and four external interrupt signals with two priority schemes for each group: grouped and spread
- Four programmable priority internal groups of eight on-chip interrupt signals with two priority schemes for each group: grouped and spread
- Priority interrupts can be programmed to support a critical (\overline{cint}) or system management (\overline{smi}) interrupt type
- External and internal interrupts directed to a host processor
- Unique vector number for each interrupt source
- Ability to redirect interrupts to external $\overline{PCI_INTA}$ pin when in core disable mode

1.2.10 Time Division Multiplexing (TDM) Interface

The TDM is a full-duplex, serial port designed to allow digital signal processors (DSPs) to communicate with a variety of serial devices, and typically transfers samples in a periodic manner. The TDM consists of independent transmitter and receiver sections with independent clock generation and frame synchronization.

The TDM interface supports 128 channels running at up to 50 Mbps with 8- and 16-bit word size. The TDM bus connects gluelessly to most T1/E1 frames as well as to common buses such as the H.110, SCAS, and MVIP. The TDM also supports an I2S mode and operates in independent or shared mode when receiving or transmitting data.

TDM capabilities include:

- Independent (asynchronous) or shared (synchronous) transmit and receive sections with separate or shared internal/external clocks and frame syncs
- TDM (network) mode operation allowing multiple devices to share the port with as many as 128 time slots
- Single-channel mode operation using frame sync
- Time slot enable registers (receive and transmit)
- End-of-frame interrupt
- Programmable internal clock divider
- Programmable word length (8 or 16 bits)
- Program options for frame sync and clock generation
- Programmable to MSB or LSB first
- TDM power-down feature
- A-law/ μ -law hardware conversion is supported for 8-bit channels
- Loopback mode

1.2.11 I²C Interface

The inter-IC (IIC or I²C) bus is a two-wire—serial data (SDA) and serial clock (SCL)—bidirectional serial bus that provides a simple, efficient method of data exchange between the system and other devices, such as microcontrollers, EEPROMs, real-time clock devices, A/D converters, and LCDs. The two-wire bus minimizes the interconnections between devices. The synchronous, multi-master bus of the I²C allows the connection of additional devices to the bus for expansion and system development.

The I²C controller is a true multi-master bus which includes collision detection and arbitration that prevents data corruption if two or more masters attempt to control the bus simultaneously. This feature allows for complex applications with multiprocessor control. The I²C controller consists of a transmitter/receiver unit, clocking unit, and control unit. The I²C unit supports general broadcast mode and on-chip filtering rejects spikes on the bus.

The I²C interface includes the following features:

- Two-wire interface
- Multi-master operational
- Arbitration lost interrupt with automatic mode switching from master to slave
- Calling address identification interrupt
- Bus busy detection
- Software-programmable clock frequency

- Software-selectable acknowledge bit
- On-chip filtering for spikes on the bus
- Address broadcasting supported
- System initialization data is optionally loaded from I²C EPROM by boot sequencer embedded hardware

1.2.12 DMA Controller

The DMA engine is capable of transferring blocks of data from any legal address range to any other legal address range. Therefore, it can perform a DMA transfer between any of its I/O or memory ports, or even between two devices or locations on the same port.

The DMA controller offers the following features:

- Four high-speed/high-bandwidth channels accessible by local and remote masters
- Handshaking (external control) signals supported for three channels: $\overline{\text{DREQ}}[0:1]$, $\overline{\text{DACK}}[0:1]$, $\overline{\text{DDONE}}[0:1]$
- Basic DMA operation modes (direct and extended chaining)
- Data transfer between LBC, PCI Express, DDR, and PCI
- Support for misaligned transfers
- Programmable bandwidth control between channels
- Interrupt on error and completed segment or chain
- Uses round-robin algorithm for channel arbitration

1.2.13 Dual Universal Asynchronous Receiver/Transmitter (DUART)

The device includes a DUART intended for use in maintenance, bring up, and debug systems. The device provides a standard four-wire handshake (TXD, RXD, $\overline{\text{RTS}}$, $\overline{\text{CTS}}$) for each port. The DUART is a slave interface. An interrupt is provided to the interrupt controller or optionally steered externally to allow device handshakes. Interrupts are generated for transmit, receive, and line status.

The DUART supports full-duplex operation. It is compatible with the PC16450 and PC16550 programming models. The transmitter and receiver both support 16-byte FIFOs.

Software programmable baud rate generators divide the system clock to generate a 16x clock. Serial interface data formats (data length, parity, 1/1.5/2 STOP bit, baud rate) are also software selectable.

The DUART includes the following features:

- Full-duplex operation
- Programming model compatible with the original PC16450 UART and the PC16550D (an improved version of the PC16450 that also operates in FIFO mode)
- PC16450 register reset values
- FIFO mode for both transmitter and receiver, providing 16-byte FIFOs
- Serial data encapsulation and decapsulation with standard asynchronous communication bits (START, STOP, and parity)

- Maskable transmit, receive, and line status interrupts
- Software-programmable baud rate generators that divide the system clock by 1 to $(2^{16} - 1)$ and generate a 16x clock for the transmitter and receiver engines
- Clear to send ($\overline{\text{CTS}}$) and ready to send ($\overline{\text{RTS}}$) MODEM control functions
- Software-selectable serial-interface data format (data length, parity, 1/1.5/2 STOP bit, baud rate)
- Line status registers
- Line-break detection and generation
- Internal diagnostic support, local loopback, and break functions
- Prioritized interrupt reporting
- Overrun, parity, and framing error detection

1.2.14 System Timers

The system includes the following timers:

- Periodic interrupt timer
- Real time clock
- Software watchdog timer
- Two general-purpose timer blocks, each supporting four 16-bit programmable timers, two cascaded 32-bit timers, or one cascaded 64-bit counter

1.3 Applications

The internal features of the MPC8315E make it suitable for a wide variety of network communication applications. It addresses the requirements of several storage, consumer, and industrial applications, including main CPUs and I/O processors in network attached storage (NAS), voice over IP (VoIP) router/gateway, intelligent wireless LAN (WLAN), set top boxes, industrial controllers, and wireless access points.

1.3.1 Media Server/NAS

Figure 1-5 shows how the MPC8315E can be configured as a media server.

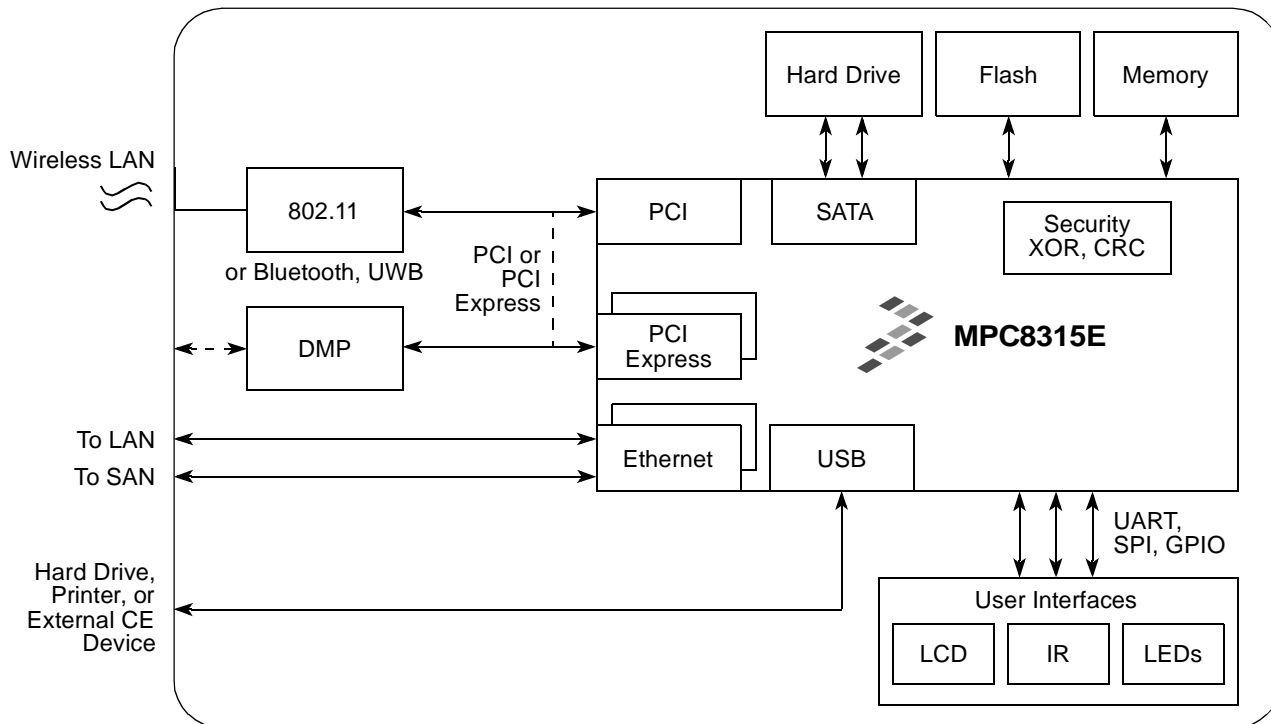


Figure 1-5. MPC8315E as a Media Server

Multimedia home networking emphasizes both audio and video streaming around the house. Since digital audio takes up relatively little bandwidth, almost any current home network can stream digital audio. Ultimately, it is video that is the test of a multimedia home network. As consumers begin to demand standard and high definition video streams and content providers require Digital Rights Management of encrypted streams, multimedia system and network demands grow significantly.

The MPC8315E offers significant processor and memory performance coupled with high levels of SoC integration. The dual integrated SATA 3 Gb/s controllers offer connectivity to the amount of storage necessary in a NAS application while dual Ethernet controllers offer connectivity to both a LAN and a storage area network (SAN). Wireless LAN is provided via either PCI Express for next generation IEEE Std. 802.11n chipsets and PCI for legacy IEEE standards 802.11a, 802.11b, 802.11g. The security engine provides acceleration for IPsec as well as DTCP-IP applications.

1.3.2 Low-End Voice Gateway

Figure 1-6 illustrates how the MPC8314E can perform the function of a complete low-end voice gateway system.

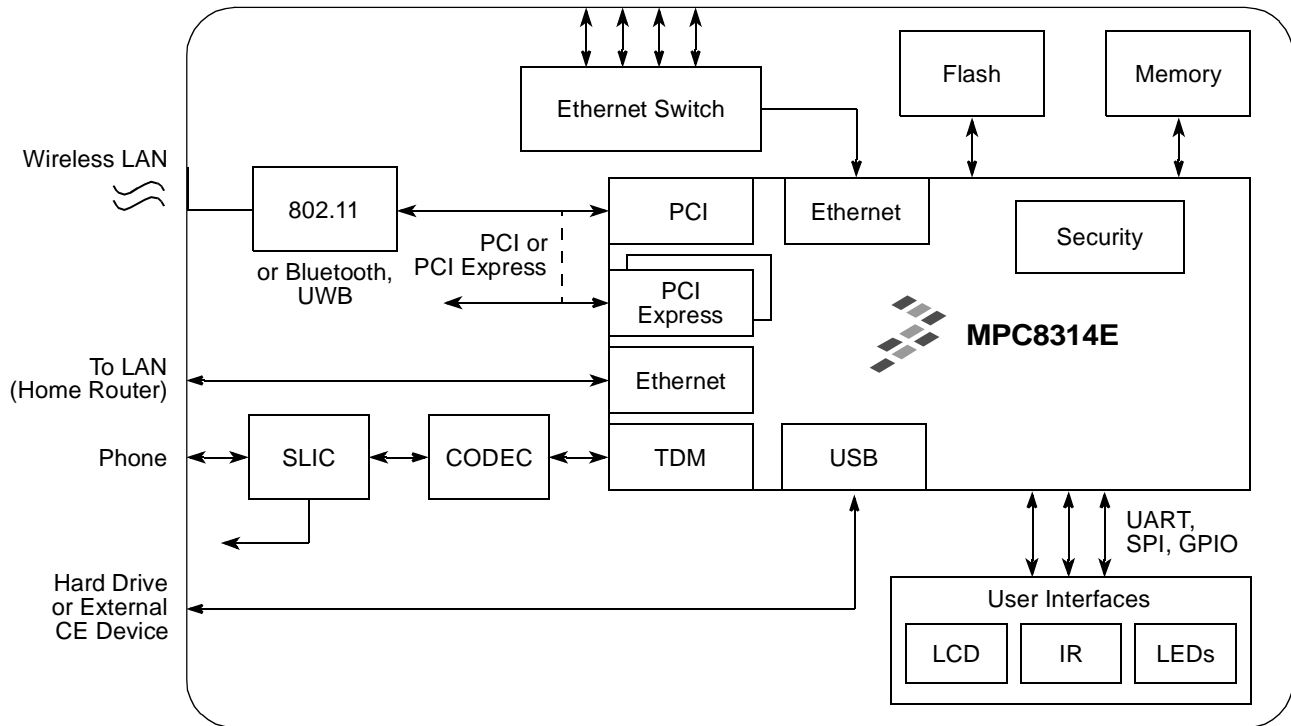


Figure 1-6. MPC8314E Serving as a Low-End Voice Gateway Application

In this application, the TDM interface provides connectivity to the codec and SLIC for easy integration of voice capabilities. The Ethernet interfaces provide wired access to the home router/LAN and Ethernet switch while wireless connectivity can be provided using PCI or PCI Express. The USB 2.0 interface with integrated PHY provides connectivity to either a hard drive or external CE device. The security core is available to accelerate protocols such as IPsec or IEEE Std. 802.11i for wireless LAN.

1.3.3 802.11n WLAN Access Point

Figure 1-7 illustrates the MPC8315E acting as an IEEE Std. 802.11n WLAN access point.

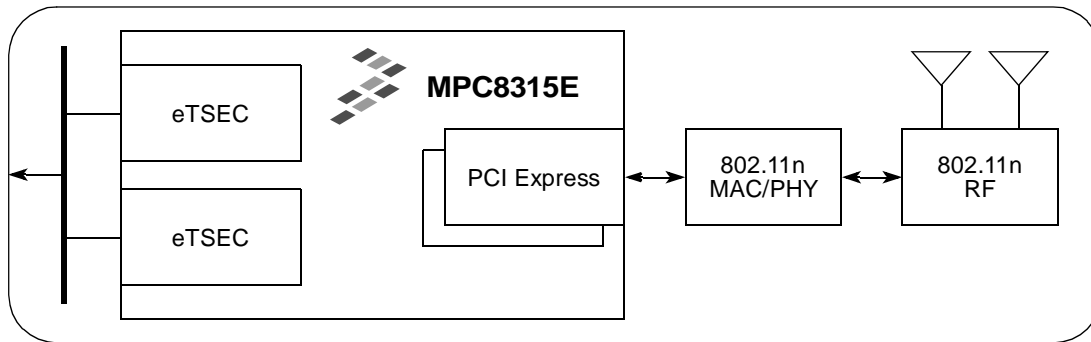


Figure 1-7. MPC8315E as a WLAN Access Point

Current systems are being designed for IEEE Std. 802.11b and 802.11a/g as well as combination radios. The WiFi chipsets use PCI in current systems. The demand on system performance continues to grow as the industry migrates to IEEE Std. 802.11n chipsets using PCI Express and the IT community requires additional management features.

WLAN access points (WAPs) require low power and are often powered exclusively by Power over Ethernet (POE). Each Ethernet line can supply about 12 W. After accounting for radio and other board power needs, this often leaves < 2.5 W for the embedded processor.

When not using PCI Express, available SGMII interfaces make it possible to connect to low-power Gigabit Ethernet PHYs. SGMII support on the Gigabit Ethernet PHYs lowers overall power consumption. The MPC8315E also has superior PCI to memory performance.

Chapter 2

Signal Descriptions

This chapter describes the external signals of the device. It is organized into the following sections:

- Overview of signals and cross references for signals that serve multiple functions, including two lists: one ordered by functional block and one alphabetical.
- List of reset configuration signals
- List of output signal states at reset

NOTE

A bar over a signal name indicates that the signal is active low, such as $\overline{\text{IRQ_OUT}}$ (interrupt out). Active-low signals are referred to as asserted (active) when they are low and negated when they are high. Signals that are not active low, such as IRQ (interrupt input), are referred to as asserted when they are high and negated when they are low.

Internal signals throughout this document are shown as lower case and in italics. For example, *sys_logic_clk* is an internal signal. These are referenced only as necessary for understanding of the external functionality of the device.

2.1 Signals Overview

The signals are grouped as follows:

- DDR memory interface signals
- PCI interface signals
- DUART interface signals
- I²C interface signals
- Serial peripheral interface signals
- Ethernet management interface signals
- eTSEC1 and USB interface signals
- eTSEC2 interface signals
- SerDes interface signals
- TDM signals
- Enhanced local bus interface signals
- USB PHY signals
- Global timers/USB interface signals
- PIC interface signals
- DMA interface signals
- SPI, JTAG, PMC, configuration, system control signals
- SGMII PHY interface signals
- SATA PHY signals (MPC8315/E only)
- Clock signals

Figure 2-1 and Figure 2-2 show the external signals of the device and how the signals are grouped. Refer to the *MPC8315E Integrated Processor Hardware Specifications* for a pinout diagram showing pin numbers and a listing of all the electrical and mechanical specifications.

Signal Descriptions

Note that individual chapters of this document provide details for each signal, describing each signal's behavior when asserted and negated and when the signal is an input or an output.

NOTE

The MPC8314E and MPC8314 do not support dual SATA controllers and any relevant signals do not apply.

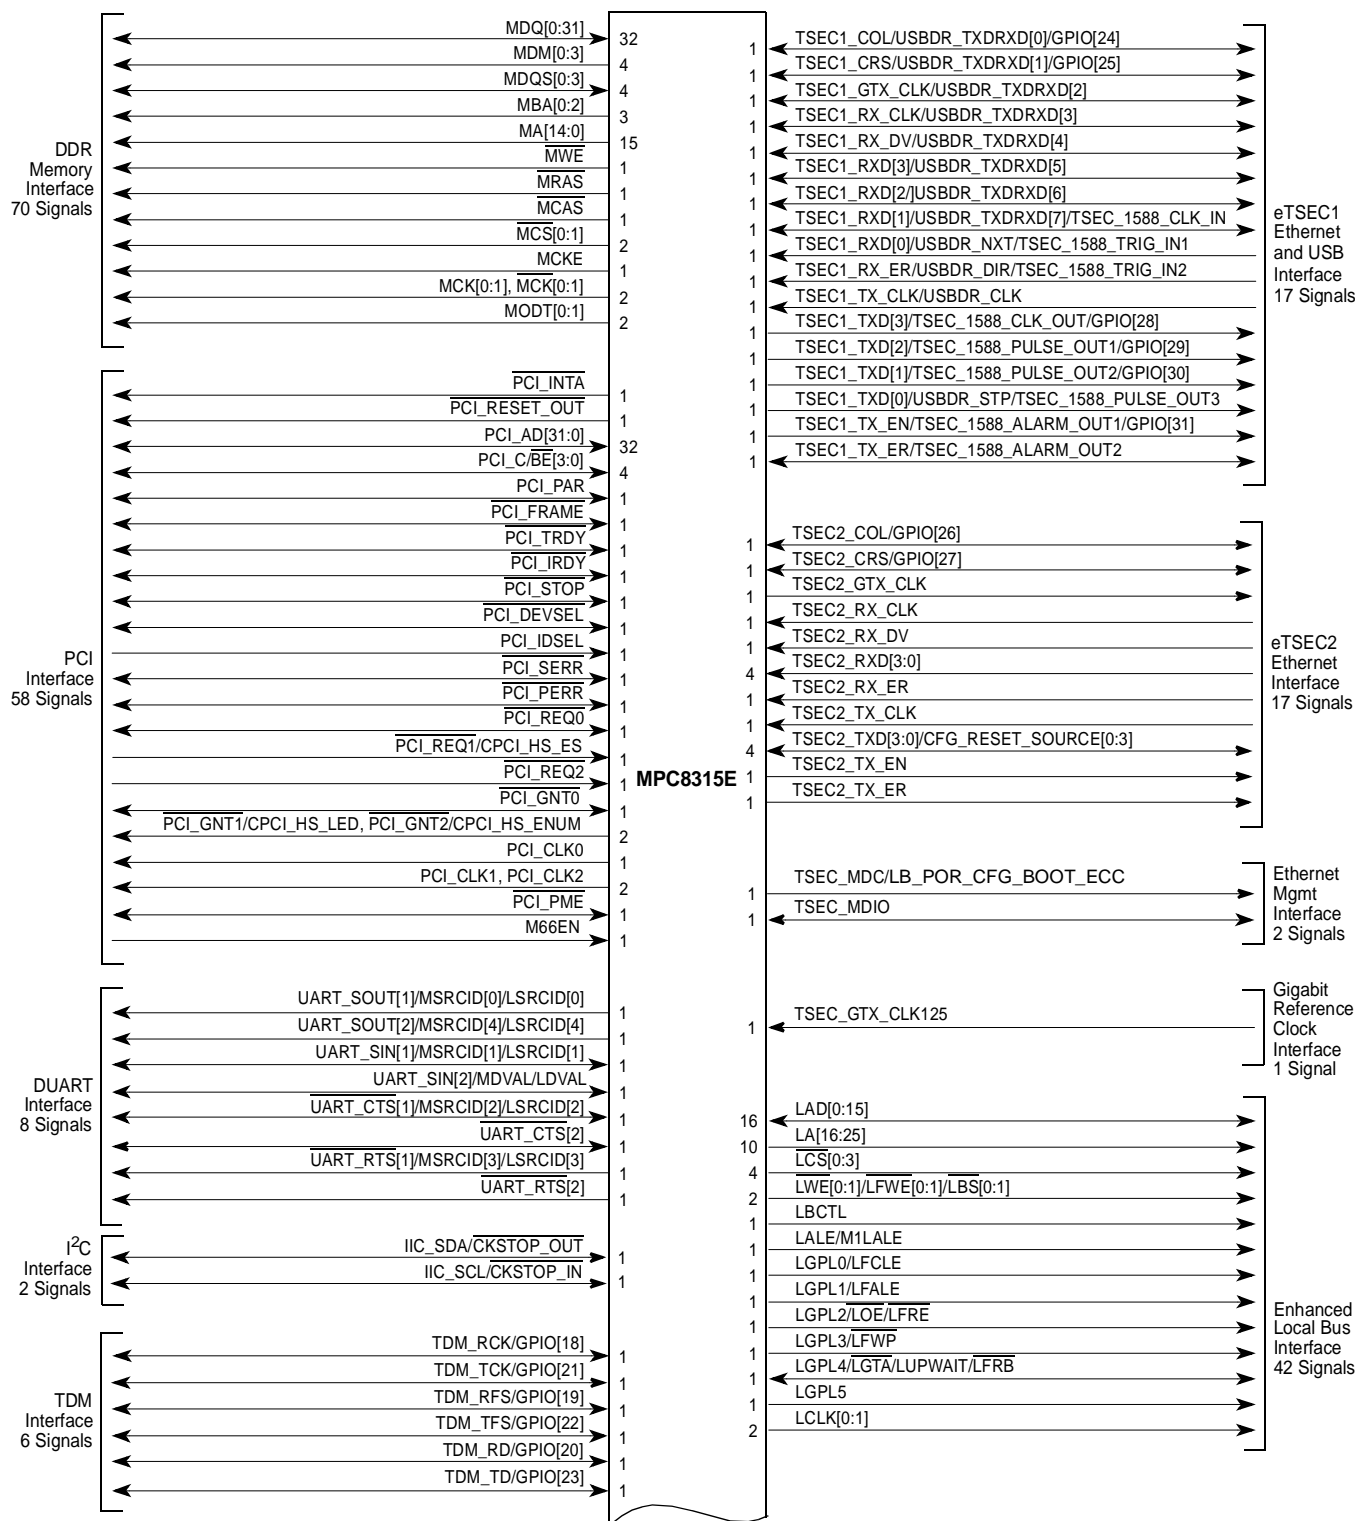


Figure 2-1. MPC8315E Signal Groupings (1 of 2)

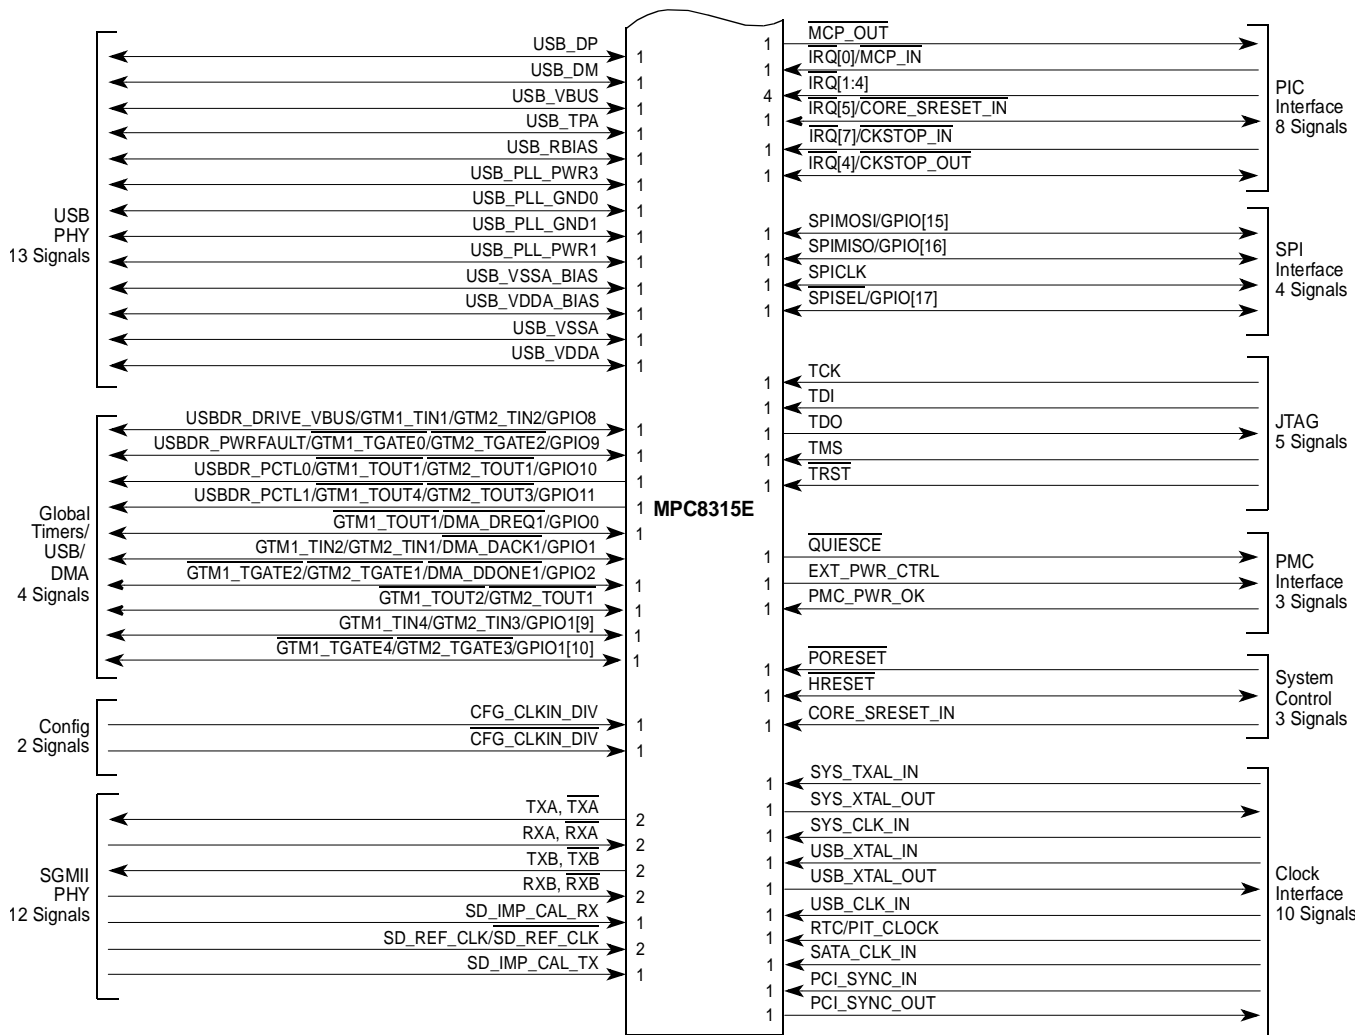


Figure 2-2. MPC8315E Signal Groupings (2 of 2)

The following tables provide summaries of signal functions. [Table 2-1](#) provides a summary of the signals grouped by function, and [Table 2-2](#) provides a summary of the signals grouped alphabetically. These tables detail the signal name, interface, alternate functions, number of signals, and whether the signal is an input, output, or bidirectional. Finally, the table provides a pointer to the table where the signal function is described.

Table 2-1. MPC8315E Signal Reference by Functional Block

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|----------|---------------------------|------------------|----------------|-----|-------------------------|-----------------------|-------------|
| MA[14:0] | DDR address | DDR | 15 | O | 9-1/9-3 | — | — |
| MBA[0:2] | DDR bank select | DDR | 3 | O | 9-1/9-3 | — | — |
| MCAS | DDR column address strobe | DDR | 1 | O | 9-1/9-3 | — | — |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|---|---------------------------|------------------|----------------|-----|-------------|-----------------------|-------------|
| MCK[0:1], $\overline{\text{MCK}}$ [0:1] | DDR differential clocks | DDR | 4 | O | 9-1/9-3 | — | — |
| MCKE | DDR clock enable | DDR | 1 | O | 9-1/9-3 | — | — |
| MCS[0:1] | DDR chip select (2/DIMM) | DDR | 2 | O | 9-1/9-3 | — | — |
| MDM[0:3] | DDR data mask | DDR | 4 | O | 9-1/9-3 | — | — |
| MDQ[0:31] | DDR data | DDR | 32 | I/O | 9-1/9-3 | — | — |
| MDQS[0:3] | DDR data strobe | DDR | 4 | I/O | 9-1/9-3 | — | — |
| MODT[0:1] | DRAM on-die termination | DDR | 2 | O | 9-1/9-3 | — | — |
| $\overline{\text{MRAS}}$ | DDR row address strobe | DDR | 1 | O | 9-1/9-3 | — | — |
| $\overline{\text{MWE}}$ | DDR write enable | DDR | 1 | O | 9-1/9-3 | — | — |
| PCI_AD[31:0] | PCI address/data | PCI | 14 | I/O | 13-2/13-4 | — | — |
| PCI_AD[14] | PCI address/data | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_AD[31:15] | PCI address/data | PCI | 17 | I/O | 13-2/13-4 | — | — |
| PCI_C/ $\overline{\text{BE}}$ [3:0] | PCI command/byte enable | PCI | 4 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_FRAME}}$ | PCI frame | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_INTA}}$ | PCI interrupt output | PCI | 1 | O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_IRDY}}$ | PCI initiator ready | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_PAR | PCI parity | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_RESET_OUT}}$ | PCI reset | PCI | 1 | O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_STOP}}$ | PCI stop | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_TRDY}}$ | PCI target ready | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_DEVSEL}}$ | PCI device select | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_IDSEL | PCI initial device select | PCI | 1 | I | 13-2/13-4 | — | — |
| $\overline{\text{PCI_SERR}}$ | PCI system error | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_PERR}}$ | PCI parity error | PCI | 1 | I/O | 13-2/13-4 | — | — |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|----------------------|------------------------------------|------------------|----------------|-----|-------------|-------------------------------|----------------------|
| PCI_REQ0 | PCI request 0 | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_REQ1 | PCI request 1 | PCI | 1 | I | 13-2/13-4 | CPCI_HS_ES | 13-2/13-4 |
| PCI_REQ2 | PCI request 2 | PCI | 1 | I | 13-2/13-4 | — | — |
| PCI_GNT0 | PCI grant 0 | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_GNT1 | PCI grant 1 | PCI | 1 | O | 13-2/13-4 | CPCI_HS_LED | 13-2/13-4 |
| PCI_GNT2 | PCI grant 2 | PCI | 1 | O | 13-2/13-4 | CPCI_HS_ENUM | 13-2/13-4 |
| PCI_CLK0 | PCI clock 0 | PCI | 1 | O | 4-2/4-3 | — | — |
| PCI_CLK1 | PCI clock 1 | PCI | 1 | O | 4-2/4-3 | — | — |
| PCI_CLK2 | PCI clock 2 | PCI | 1 | O | 4-2/4-3 | — | — |
| PCI_PME | PCI PME assertion request | PCI | 1 | I/O | 13-2/13-4 | — | — |
| M66EN | 66-MHz system configuration | PCI | 1 | I | 13-2/13-4 | — | — |
| CPCI_HS_ES | CompactPCI hot swap ejector switch | PCI | 1 | I | 13-2/13-4 | PCI_REQ1 | 13-2/13-4 |
| CPCI_HS_LED | CompactPCI hot swap LED | PCI | 1 | O | 13-2/13-4 | PCI_GNT1 | 13-2/13-4 |
| CPCI_HS_ENUM | CompactPCI hot swap enumerator | PCI | 1 | O | 13-2/13-4 | PCI_GNT2 | 13-2/13-4 |
| TSEC_1588_CLK_IN | 1588 clock-in | eTSEC | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD [7]/TSEC1_RXD[1] | 17-1/17-3, 19-2/19-8 |
| TSEC_1588_CLK_OUT | 1588 clock-out | eTSEC | 1 | I/O | 19-2/19-8 | TSEC1_TXD[3]/GPIO[28] | 19-2/19-8, 24-1/24-2 |
| TSEC_1588_TRIG_IN1 | 1588 trigger-in 1 | eTSEC | 1 | I | 19-2/19-8 | USBDR_NXT/TSEC1_RXD[0] | 17-1/17-3, 19-2/19-8 |
| TSEC_1588_TRIG_IN2 | 1588 trigger-in 2 | eTSEC | 1 | I/O | 19-2/19-8 | USBDR_DIR/TSEC1_RX_ER | 17-1/17-3, 19-2/19-8 |
| TSEC_1588_ALARM_OUT1 | 1588 timer alarm-out 1 | eTSEC | 1 | I/O | 19-2/19-8 | TSEC1_TX_EN/GPIO[31] | 19-2/19-8, 17-1/17-3 |
| TSEC_1588_ALARM_OUT2 | 1588 timer alarm-out 2 | eTSEC | 1 | O | 19-2/19-8 | TSEC1_TX_ER | 19-2/19-8 |
| TSEC_1588_PULSE_OUT1 | 1588 timer pulse-out 1 | eTSEC | 1 | I/O | 19-2/19-8 | TSEC1_TXD[2]/GPIO[29] | 19-2/19-8, 24-1/24-2 |
| TSEC_1588_PULSE_OUT2 | 1588 timer pulse-out 2 | eTSEC | 1 | I/O | 19-2/19-8 | TSEC1_TXD[1]/GPIO[30] | 19-2/19-8, 24-1/24-2 |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|----------------------|---------------------------|------------------|----------------|-----|-------------|-----------------------------------|----------------------|
| TSEC_1588_PULSE_OUT3 | 1588 timer pulse-out 3 | eTSEC | 1 | O | 19-2/19-8 | TSEC1_TXD[0]/ USBDR_STP | 19-2/19-8, 17-1/17-3 |
| TSEC1_COL | eTSEC1 collision detect | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[0]/ GPIO[24] | 17-1/17-3, 24-1/24-2 |
| TSEC1_CRS | eTSEC11 carrier sense | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[1]/ GPIO[25] | 17-1/17-3, 24-1/24-2 |
| TSEC1_GTX_CLK | eTSEC1 transmit clock out | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[2] | 17-1/17-3 |
| TSEC1_RX_CLK | eTSEC1 receive clock | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[3] | 17-1/17-3 |
| TSEC1_RX_DV | eTSEC1 receive data valid | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[4] | 17-1/17-3 |
| TSEC1_RXD[3:2] | eTSEC1 receive data 3–2 | eTSEC1 | 2 | I/O | 19-2/19-8 | USBDR_TXDRXD [5:6] | 17-1/17-3 |
| TSEC1_RXD[1] | eTSEC1 receive data 1 | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_CLK_IN/ USBDR_TXDRXD[7] | 19-2/19-8, 17-1/17-3 |
| TSEC1_RXD[0] | eTSEC1 receive data 0 | eTSEC1 | 1 | I | 19-2/19-8 | TSEC_1588_TRIG_IN1/USBDR_NXT | 19-2/19-8, 17-1/17-3 |
| TSEC1_RX_ER | eTSEC1 receiver error | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_TRIG_IN2/USBDR_DIR | 19-2/19-8, 17-1/17-3 |
| TSEC1_TX_ER | eTSEC1 transmit error | eTSEC1 | 1 | O | 19-2/19-8 | TSEC_1588_ALARM_OUT1 | 19-2/19-8 |
| TSEC1_TX_CLK | eTSEC1 transmit clock in | eTSEC1 | 1 | I | 19-2/19-8 | USBDR_CLK | 17-1/17-3 |
| TSEC1_TXD[3] | eTSEC1 transmit data 3 | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_CLK_OUT/GPIO[28] | 19-2/19-8, 24-1/24-2 |
| TSEC1_TXD[2] | eTSEC1 transmit data 2 | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_PULSE_OUT1/GPIO[29] | 19-2/19-8, 24-1/24-2 |
| TSEC1_TXD[1] | eTSEC1 transmit data 1 | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_PULSE_OUT2/GPIO[30] | 19-2/19-8, 24-1/24-2 |
| TSEC1_TXD[0] | eTSEC1 transmit data 0 | eTSEC1 | 1 | O | 19-2/19-8 | TSEC_1588_PULSE_OUT3/USBDR_STP | 19-2/19-8, 17-1/17-3 |
| TSEC1_TX_EN | eTSEC1 transmit enable | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_ALARM_OUT1/GPIO[31] | 19-2/19-8, 17-1/17-3 |
| TSEC2_COL | eTSEC2 collision detect | eTSEC2 | 1 | I/O | 19-2/19-8 | GPIO[26] | 24-1/24-2 |
| TSEC2_CRS | eTSEC2 carrier sense | eTSEC2 | 1 | I/O | 19-2/19-8 | GPIO[27] | 24-1/24-2 |
| TSEC2_GTX_CLK | eTSEC2 transmit clock out | eTSEC2 | 1 | O | 19-2/19-8 | — | — |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|-------------------------|---|------------------------|----------------|-----|-------------|----------------------------------|-------------|
| TSEC2_RX_CLK | eTSEC2 receive clock | eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC2_RX_DV | eTSEC2 receive data valid | eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC2_RXD[3:0] | eTSEC2 receive data 3–0 | eTSEC2 | 4 | I | 19-2/19-8 | — | — |
| TSEC2_RX_ER | eTSEC2 receiver error | eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC2_TXD[3:0] | eTSEC2 transmit data 3–0 | eTSEC2 | 4 | I/O | 19-2/19-8 | CFG_RESET_SOURCE[0:3] | 4-1/4-1 |
| TSEC2_TX_ER | eTSEC2 transmit error | eTSEC2 | 1 | O | 19-2/19-8 | — | — |
| TSEC2_TX_EN | eTSEC2 transmit enable | eTSEC2 | 1 | O | 19-2/19-8 | — | — |
| TSEC2_TX_CLK | eTSEC2 transmit clock in | eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC_GTX_CLK125 | Gigabit reference clock | eTSEC1/ eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC_MDC | Ethernet management data clock | Ethernet management | 1 | O | 19-2/19-8 | LB_POR_CFG_BOOT_ECC ¹ | — |
| TSEC_MDIO | Ethernet management data in/out | Ethernet management | 1 | I/O | 19-2/19-8 | — | — |
| RXA | Serial receiver, lane A, positive data | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| $\overline{\text{RXA}}$ | Serial receiver, lane A, negative data (complement) | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| RXB | Serial receiver, lane B, positive data | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| $\overline{\text{RXB}}$ | Serial receiver, lane B, negative data (complement) | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| SDAVDD | Analog supply for SerDes PLL | SGMII PHY ² | 1 | I | 16-1/16-3 | — | — |
| SDAVSS | Analog ground for SerDes PLL | SGMII PHY | 1 | I | 16-1/16-3 | — | — |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|----------------------------------|--|------------------------|----------------|-----|---------------------------|-----------------------|-------------|
| SD_IMP_CAL_RX | Receiver impedance control signal | SGMII PHY ³ | 1 | I | 16-1/16-3 | — | — |
| SD_IMP_CAL_TX | Transmitter impedance control signal | SGMII PHY ³ | 1 | I | 16-1/16-3 | — | — |
| SD_PLL_TPA_ANA | Analog test point for SerDes PLL testing | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| SD_PLL_TPD | Digital test point for SerDes PLL testing | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| SD_REF_CLK | SerDes PLL reference clock | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| $\overline{\text{SD_REF_CLK}}$ | SerDes PLL reference clock (complement) | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| TXA | Serial transmitter, lane A, positive data | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| $\overline{\text{TXA}}$ | Serial transmitter, lane A, negative data (complement) | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| TXB | Serial transmitter, lane B, positive data | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| $\overline{\text{TXB}}$ | Serial transmitter, lane B, negative data (complement) | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| XCOREVDD[0:3] | SerDes transceiver core supply | SGMII PHY ² | 4 | PWR | 16-1/16-3 | — | — |
| XCOREVSS[0:3] | SerDes transceiver core ground | SGMII PHY | 4 | GND | 16-1/16-3 | — | — |
| XPADVDD[0:2] | SerDes transceiver pad supply | SGMII PHY ² | 3 | PWR | 16-1/16-3 | — | — |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|-------------------------------|--|------------------|----------------|-----|-------------|-----------------------------------|----------------------|
| XPADVSS[0:2] | SerDes transceiver pad ground | SGMII PHY | 3 | GND | 16-1/16-3 | — | — |
| LAD[0:15] | LBC address/data | eLBC | 16 | I/O | 10-1/10-4 | — | — |
| LA[16:25] | LBC port address | eLBC | 10 | O | 10-1/10-4 | — | — |
| $\overline{\text{LCS}}$ [0:3] | LBC chip select 0–3 | eLBC | 4 | O | 10-1/10-4 | — | — |
| $\overline{\text{LWE}}$ [0:1] | LBC write enable | eLBC | 2 | O | 10-1/10-4 | — | — |
| LBCTL | LBC data buffer control | eLBC | 1 | O | 10-1/10-4 | — | — |
| LALE | LBC address latch enable | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL0 | LBC UPM general purpose line 0/ | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL1/LFALE | LBC GP line 1/Flash address latch enable | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL2 | LBC output enable | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL3 | LBC GP line 3 | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL4 | LBC GP line 4 | eLBC | 1 | I/O | 10-1/10-4 | — | — |
| LGPL5 | LBC GP line 5 | eLBC | 1 | O | 10-1/10-4 | — | — |
| LCLK[0:1] | LBC clocks 0–1 | eLBC | 2 | O | 10-1/10-4 | — | — |
| USBDR_CLK | Clocking signal for ULPI PHY interface | USB | 1 | I | 17-1/17-3 | TSEC1_TX_CLK | 19-2/19-8 |
| USBDR_DIR | Direction of data bus | USB | 1 | I/O | 17-1/17-3 | TSEC_1588_TRIG_IN2/TSEC1_RX_ER | 19-2/19-8 |
| USBDR_STP | End of a transfer on the bus | USB | 1 | O | 17-1/17-3 | TSEC1_TXD[0]/TSEC_1588_PULSE_OUT3 | 19-2/19-8 |
| USBDR_NXT | Next data | USB | 1 | I | 17-1/17-3 | TSEC_1588_TRIG_IN1/TSEC1_RXD[0] | 19-2/19-8 |
| USBDR_TXDRXD[0] | Data bit 0 | USB | 1 | I/O | 17-1/17-3 | TSEC1_COL/GPIO[24] | 19-2/19-8, 24-1/24-2 |
| USBDR_TXDRXD[1] | Data bit 1 | USB | 1 | I/O | 17-1/17-3 | TSEC1_CRS/GPIO[25] | 19-2/19-8, 24-1/24-2 |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|-------------------|--|------------------|----------------|-----|-------------|---|-------------------------|
| USBDR_TXDRXD[2] | Data bit 2 | USB | 1 | I/O | 17-1/17-3 | TSEC1_GTX_CLK | 19-2/19-8 |
| USBDR_TXDRXD[3] | Data bit 3 | USB | 1 | I/O | 17-1/17-3 | TSEC1_RX_CLK | 19-2/19-8 |
| USBDR_TXDRXD[4] | Data bit 4 | USB | 1 | I/O | 17-1/17-3 | TSEC1_RX_DV | 19-2/19-8 |
| USBDR_TXDRXD[5:6] | Data bit 5–6 | USB | 2 | I/O | 17-1/17-3 | TSEC1_RXD[3:2] | 19-2/19-8 |
| USBDR_TXDRXD[7] | Data bit 7 | USB | 1 | I/O | 17-1/17-3 | TSEC1_RXD[1]/ TSEC_1588_CLK_IN | 19-2/19-8 |
| USBDR_PCTL0 | Port control 0 | USB | 1 | I/O | 17-1/17-3 | $\overline{\text{GTM1_TOUT2}}$ / $\overline{\text{GTM2_TOUT1}}$ / GPIO[10] | 5-58/5-56, 24-1/24-2 |
| USBDR_PCTL1 | Port control 1 | USB | 1 | I/O | 17-1/17-3 | $\overline{\text{GTM1_TOUT4}}$ / $\overline{\text{GTM2_TOUT3}}$ / GPIO[11] | 5-58/5-56, 24-1/24-2 |
| USBDR_DRIVE_VBUS | USB VBus power enable | USB | 1 | I/O | 17-1/17-3 | GTM1_TIN1/ GTM2_TIN2/ GPIO[8]/ | 5-58/5-56, 24-1/24-2 |
| USBDR_PWRFAULT | USB VBus power fault | USB | 1 | I/O | 17-1/17-3 | $\overline{\text{GTM1_TGATE1}}$ / $\overline{\text{GTM2_TGATE2}}$ / GPIO[0] | 5-58/5-56, 24-1/24-2 |
| USBDR_PCTL0 | USB status LED 0 control | USB | 1 | I/O | 17-1/17-3 | $\overline{\text{GTM1_TOUT2}}$ / $\overline{\text{GTM2_TOUT1}}$ / GPIO[10] | 5-58/5-56, 24-1/24-2 |
| USBDR_PCTL1 | USB status LED 0 control | USB | 1 | I/O | 17-1/17-3 | $\overline{\text{GTM1_TOUT4}}$ / $\overline{\text{GTM2_TOUT3}}$ / GPIO[11] | 5-58/5-56, 24-1/24-2 |
| USB_DP | USB 2.0 D+ line | USB PHY | 1 | I/O | 17-1/17-3 | — | — |
| USB_DM | USB 2.0 D– line | USB PHY | 1 | I/O | 17-1/17-3 | — | — |
| USB_RBIAIS | USB bias input ⁴ | USB_PHY | 1 | I | 17-1/17-3 | — | — |
| USB_VBUS | USB 2.0 VBUS line | USB PHY | 1 | I/O | 17-1/17-3 | — | — |
| USB_TPA | Dedicated analog test signal | USB PHY | 1 | I/O | 17-1/17-3 | — | — |
| USB_PLL_PWR1 | Dedicated 1.0 V analog power for USB PLL | USB PHY | 1 | I | 17-1/17-3 | — | — |
| SATA_ANAVIZ | | SATA PHY | 1 | O | — | — | — |
| SATA_CLK_IN | Clock in | SATA PHY | 1 | I | — | — | — |
| SATA_VDD | Voltage supply | SATA PHY | 2 | I | — | — | — |
| SATA_VSS | Ground | SATA PHY | 2 | I | — | — | — |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|--------------------------------|-------------------------------|------------------|----------------|-----|-------------|---------------------------------|--------------------|
| PINRXPLUSA | Receive plus A | SATA PHY | 1 | I | — | — | — |
| PINRXPLUSB | Receive plus B | SATA PHY | 1 | I | — | — | — |
| PINTXPLUSA | Transmit plus A | SATA PHY | 1 | O | — | — | — |
| PINTXPLUSB | Transmit plus A | SATA PHY | 1 | O | — | — | — |
| PINRXMINUSA | Receive minus A | SATA PHY | 1 | I | — | — | — |
| PINRXMINUSB | Receive minus B | SATA PHY | 1 | I | — | — | — |
| PINTXMINUSA | Transmit minus A | SATA PHY | 1 | O | — | — | — |
| PINTXMINUSB | Transmit minus B | SATA PHY | 1 | O | — | — | — |
| VDD33ANA | Analog 3.3 V supply | SATA PHY | 1 | I | — | — | — |
| VDD33PLL | PLL 3.3 V supply | SATA PHY | 1 | I | — | — | — |
| VSSRESREF | Reset reference ground | SATA PHY | 1 | I | — | — | — |
| RESREF | Reset reference | SATA PHY | 1 | I | — | — | — |
| IIC_SDA | I ² C serial data | I ² C | 1 | I/O | — | $\overline{\text{CKSTOP_OUT}}$ | 4-3/4-5 |
| IIC_SCL | I ² C serial clock | I ² C | 1 | I/O | — | $\overline{\text{CKSTOP_IN}}$ | 4-3/4-5 |
| UART_SIN1 | DUART serial data in | DUART | 1 | I/O | 21-1/21-3 | MSRCID1/LSRCID1 | 9-1/9-3, 10-1/10-4 |
| UART_SIN2 | DUART serial data in | DUART | 1 | I/O | 21-1/21-3 | MDVAL/LDVAL | 9-1/9-3, 10-1/10-4 |
| UART_SOUT1 | DUART serial data out | DUART | 1 | O | 21-1/21-3 | MSRCID0/LSRCID0 | 9-1/9-3, 10-1/10-4 |
| UART_SOUT2 | DUART serial data out | DUART | 1 | O | 21-1/21-3 | MSRCID4/LSRCID4 | 9-1/9-3, 10-1/10-4 |
| $\overline{\text{UART_CTS1}}$ | DUART clear to send | DUART | 1 | I/O | 21-1/21-3 | MSRCID2/LSRCID2 | 9-1/9-3, 10-1/10-4 |
| $\overline{\text{UART_CTS2}}$ | DUART clear to send | DUART | 1 | I/O | 21-1/21-3 | — | — |
| $\overline{\text{UART_RTS1}}$ | DUART ready to send | DUART | 1 | I/O | 21-1/21-3 | MSRCID3/LSRCID3 | 9-1/9-3, 10-1/10-4 |
| $\overline{\text{UART_RTS2}}$ | DUART ready to send | DUART | 1 | I/O | 21-1/21-3 | — | — |
| SPIMOSI | SPI master-out slave-in | SPI | 1 | I/O | 22-1/22-7 | GPIO[15] | 24-1/24-2 |
| SPIMISO | SPI master-in slave-out | SPI | 1 | I/O | 22-1/22-7 | GPIO[16] | 24-1/24-2 |
| SPICLK | SPI clock | SPI | 1 | I/O | 22-1/22-7 | — | — |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|-----------------------------|--------------------------------|------------------|----------------|-----|-------------|------------------------------|-------------------------|
| $\overline{\text{SPISEL}}$ | SPI slave select | SPI | 1 | I/O | 22-1/22-7 | GPIO[17] | 24-1/24-2 |
| TDM_RCK | TDM receive clock | TDM | 1 | I/O | 25-1/25-2 | GPIO[18] | 24-1/24-2 |
| TDM_TCK | TDM transmit clock | TDM | 1 | I/O | 25-1/25-2 | GPIO[21] | 24-1/24-2 |
| TDM_RFS | TDM receive frame sync | TDM | 1 | I/O | 25-1/25-2 | GPIO[19] | 24-1/24-2 |
| TDM_TFS | TDM transmit frame sync | TDM | 1 | I/O | 25-1/25-2 | GPIO[22] | 24-1/24-2 |
| TDM_RD | TDM receive data | TDM | 1 | I/O | 25-1/25-2 | GPIO[20] | 24-1/24-2 |
| TDM_TD | TDM transmit data | TDM | 1 | I/O | 25-1/25-2 | GPIO[23] | 24-1/24-2 |
| GTM1_TIN1/ GTM2_TIN2 | Timer in 1/2 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[8]/USBDR_ DRIVE_VBUS | 24-1/24-2, 17-1/17-3 |
| GTM1_TGATE1/ GTM2_TGATE2 | Timer gate 1/2 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[9]/USBDR_ PWRFAULT | 24-1/24-2, 17-1/17-3 |
| GTM1_TOUT1 | Timer out 1 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[0]/ DMA_DREQ1 | 24-1/24-2, 12-1/12-2 |
| GTM1_TIN2/ GTM2_TIN1 | Timer in 2/1 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[1]/ DMA_DACK1 | 24-1/24-2, 12-1/12-2 |
| GTM1_TGATE2/ GTM2_TGATE1 | Timer gate 2/1 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[2]/ DMA_DDONE1 | 24-1/24-2, 12-1/12-2 |
| GTM1_TOUT2/ GTM2_TOUT1 | Timer out 2/1 | Global Timers | 1 | I/O | 5-58/5-56 | USBDR_PCTL0/ GPIO[10] | 17-1/17-3, 24-1/24-2 |
| GTM1_TIN3/ GTM2_TIN4 | Timer in 3/4 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[3] | 24-1/24-2 |
| GTM1_TGATE3/ GTM2_TGATE4 | Timer gate 3/4 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[4] | 24-1/24-2 |
| GTM1_TOUT3/ GTM2_TOUT1 | Timer out 3/1 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[5] | 24-1/24-2 |
| GTM1_TIN4/ GTM2_TIN3 | Timer in 4/3 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[6] | 24-1/24-2 |
| GTM1_TGATE4/ GTM2_TGATE3 | Timer gate 4/3 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[7] | 24-1/24-2 |
| GTM1_TOUT4/ GTM2_TOUT3 | Timer out 4/3 | Global Timers | 1 | I/O | 5-58/5-56 | USBDR_PCTL1/ GPIO[11] | 17-1/17-3, 24-1/24-2 |
| MCP_OUT | Machine check interrupt output | IPIC | 1 | O | 8-1/8-5 | — | — |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|--|------------------------------|------------------|----------------|-----|----------------|---|-------------------------|
| $\overline{\text{IRQ}}[0]/\overline{\text{MCP_IN}}$ | External interrupt 0 | IPIC | 1 | I | 8-1/8-5 | — | — |
| $\overline{\text{IRQ}}[1:4]$ | External interrupt 1–4 | IPIC | 4 | I | 8-1/8-5 | — | — |
| $\overline{\text{IRQ}}[5]$ | External interrupt 5 | IPIC | 1 | I | 8-1/8-5 | $\overline{\text{CORE_SRESET_IN}}$ | 4-1/4-1 |
| $\overline{\text{IRQ}}[6]$ | External interrupt 6 | IPIC | 1 | I/O | 8-1/8-5 | $\overline{\text{CKSTOP_OUT}}$ | 4-3/4-5 |
| $\overline{\text{IRQ}}[7]$ | External interrupt 7 | IPIC | 1 | I | 8-1/8-5 | $\overline{\text{CKSTOP_IN}}$ | 4-3/4-5 |
| TCK | Test clock | JTAG | 1 | I | 23-1/23-2 | — | — |
| TDI | Test data in | JTAG | 1 | I | 23-1/23-2 | — | — |
| TDO | Test data out | JTAG | 1 | O | 23-1/23-2 | — | — |
| TMS | Test mode select | JTAG | 1 | I | 23-1/23-2 | — | — |
| $\overline{\text{TRST}}$ | Test reset | JTAG | 1 | I | 23-1/23-2 | — | — |
| GPIO[0] | General-purpose I/O signal 0 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{GTM1_TOUT1/}}$ $\overline{\text{DMA_DREQ1}}$ | 5-58/5-56, 12-1/12-2 |
| GPIO[1] | General-purpose I/O signal 1 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{GTM1_TIN2/}}$ $\overline{\text{GTM2_TIN1/}}$ $\overline{\text{DMA_DACK1}}$ | 5-58/5-56, 12-1/12-2 |
| GPIO[2] | General-purpose I/O signal 2 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{GTM1_TGATE2/}}$ $\overline{\text{GTM2_TGATE1/}}$ $\overline{\text{DMA_DDONE1}}$ | 5-58/5-56, 12-1/12-2 |
| GPIO[3] | General-purpose I/O signal 3 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{GTM1_TIN3/}}$ $\overline{\text{GTM2_TIN4}}$ | 5-58/5-56 |
| GPIO[4] | General-purpose I/O signal 4 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{GTM1_TGATE3/}}$ $\overline{\text{GTM2_TGATE4}}$ | 5-58/5-56 |
| GPIO[5] | General-purpose I/O signal 5 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{GTM1_TOUT3/}}$ $\overline{\text{GTM2_TOUT1}}$ | 5-58/5-56 |
| GPIO[6] | General-purpose I/O signal 6 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{GTM1_TIN4/}}$ $\overline{\text{GTM2_TIN3}}$ | 5-58/5-56 |
| GPIO[7] | General-purpose I/O signal 7 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{GTM1_TGATE4/}}$ $\overline{\text{GTM2_TGATE3}}$ | 5-58/5-56 |
| GPIO[8] | General-purpose I/O signal 8 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{USBDR_DRIVE_}}$ $\overline{\text{VBUS/GTM1_TIN1/}}$ $\overline{\text{GTM2_TIN2}}$ | 17-1/17-3, 5-58/5-56 |
| GPIO[9] | General-purpose I/O signal 9 | GPIO | 1 | I/O | 24-1/24-2 | $\overline{\text{USBDR_PWRFAULT/}}$ $\overline{\text{GTM1_TGATE1/}}$ $\overline{\text{GTM2_TGATE2}}$ | 17-1/17-3, 5-58/5-56 |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|----------|-------------------------------|------------------|----------------|-----|---------------------------|--|--|
| GPIO[10] | General-purpose I/O signal 10 | GPIO | 1 | I/O | 24-1/24-2 | USBD _R _PCTL0/ GTM1_TOUT2/ GTM2_TOUT1 | 17-1/17-3 , 5-58/5-56 |
| GPIO[11] | General-purpose I/O signal 11 | GPIO | 1 | I/O | 24-1/24-2 | USBD _R _PCTL1/ GTM1_TOUT4/ GTM2_TOUT3 | 17-1/17-3 , 5-58/5-56 |
| GPIO[12] | General-purpose I/O signal 12 | GPIO | 1 | I/O | 24-1/24-2 | DMA_DREQ0 | 12-1/12-2 |
| GPIO[13] | General-purpose I/O signal 13 | GPIO | 1 | I/O | 24-1/24-2 | DMA_DACK0 | 12-1/12-2 |
| GPIO[14] | General-purpose I/O signal 14 | GPIO | 1 | I/O | 24-1/24-2 | DMA_DDONE0 | 12-1/12-2 |
| GPIO[15] | General-purpose I/O signal 15 | GPIO | 1 | I/O | 24-1/24-2 | SPI MOSI | 22-2/22-7 |
| GPIO[16] | General-purpose I/O signal 16 | GPIO | 1 | I/O | 24-1/24-2 | SPI MISO | 22-2/22-7 |
| GPIO[17] | General-purpose I/O signal 17 | GPIO | 1 | I/O | 24-1/24-2 | SPI SEL | 22-2/22-7 |
| GPIO[18] | General-purpose I/O signal 18 | GPIO | 1 | I/O | 24-1/24-2 | TDM_RCK | 25-1/25-2 |
| GPIO[19] | General-purpose I/O signal 19 | GPIO | 1 | I/O | 24-1/24-2 | TDM_RFS | 25-1/25-2 |
| GPIO[20] | General-purpose I/O signal 20 | GPIO | 1 | I/O | 24-1/24-2 | TDM_RD | 25-1/25-2 |
| GPIO[21] | General-purpose I/O signal 21 | GPIO | 1 | I/O | 24-1/24-2 | TDM_TCK | 25-1/25-2 |
| GPIO[22] | General-purpose I/O signal 22 | GPIO | 1 | I/O | 24-1/24-2 | TDM_TFS | 25-1/25-2 |
| GPIO[23] | General-purpose I/O signal 23 | GPIO | 1 | I/O | 24-1/24-2 | TDM_TD | 25-1/25-2 |
| GPIO[24] | General-purpose I/O signal 24 | GPIO | 1 | I/O | 24-1/24-2 | USBD _R _TXDRXD[0]/ TSEC1_COL | 17-1/17-3 , 19-2/19-8 |
| GPIO[25] | General-purpose I/O signal 25 | GPIO | 1 | I/O | 24-1/24-2 | USBD _R _TXDRXD[1]/ TSEC1_CRS | 17-1/17-3 , 19-2/19-8 |
| GPIO[26] | General-purpose I/O signal 26 | GPIO | 1 | I/O | 24-1/24-2 | TSEC2_COL | 19-2/19-8 |
| GPIO[27] | General-purpose I/O signal 27 | GPIO | 1 | I/O | 24-1/24-2 | TSEC2_CRS | 19-2/19-8 |
| GPIO[28] | General-purpose I/O signal 28 | GPIO | 1 | I/O | 24-1/24-2 | TSEC1_TXD[3]/ TSEC_1588_CLK_ OUT | 19-2/19-8 |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|--------------------------------------|---|------------------|----------------|-----|-------------|--|-------------------------|
| GPIO[29] | General-purpose I/O signal 29 | GPIO | 1 | I/O | 24-1/24-2 | TSEC_1588_PULSE_OUT1/ TSEC1_TXD[2] | 19-2/19-8 |
| GPIO[30] | General-purpose I/O signal 30 | GPIO | 1 | I/O | 24-1/24-2 | TSEC_1588_PULSE_OUT2/ TSEC1_TXD[1] | 19-2/19-8 |
| GPIO[31] | General-purpose I/O signal 31 | GPIO | 1 | I/O | 24-1/24-2 | TSEC1_TX_EN/ TSEC_1588_ALARM_OUT1 | 19-2/19-8 |
| $\overline{\text{DMA_DREQ0}}$ | DMA request 0 | DMA | 1 | I/O | 12-1/12-2 | GPIO[12] | 24-1/24-2 |
| $\overline{\text{DMA_DACK0}}$ | DMA acknowledge 0 | DMA | 1 | I/O | 12-1/12-2 | GPIO[13] | 24-1/24-2 |
| $\overline{\text{DMA_DDONE0}}$ | DMA done 0 | DMA | 1 | I/O | 12-1/12-2 | GPIO[14] | 24-1/24-2 |
| $\overline{\text{DMA_DREQ1}}$ | DMA request 1 | DMA | 1 | I/O | 12-1/12-2 | GPIO[0]/ $\overline{\text{GTM1_TOUT1}}$ | 24-1/24-2, 5-58/5-56 |
| $\overline{\text{DMA_DACK1}}$ | DMA acknowledge 1 | DMA | 1 | I/O | 12-1/12-2 | GPIO[1]/ GTM1_TIN2/ GTM2_TIN1 | 24-1/24-2, 5-58/5-56 |
| $\overline{\text{DMA_DDONE1}}$ | DMA done 1 | DMA | 1 | I/O | 12-1/12-2 | GPIO[2]/ $\overline{\text{GTM1_TGATE2}}$ / $\overline{\text{GTM2_TGATE1}}$ | 24-1/24-2, 5-58/5-56 |
| EXT_PWR_CTRL | External power control | PMC | 1 | O | 5-70/5-70 | — | — |
| PMC_PWR_OK | Stable power | PMC | 1 | I | 5-70/5-70 | — | — |
| $\overline{\text{QUIESCE}}$ | Quiesce state | PMC | 1 | O | 5-70/5-70 | — | — |
| $\overline{\text{CORE_SRESET_IN}}$ | Soft reset to the e300 core | System control | 1 | I | 4-1/4-1 | — | — |
| $\overline{\text{PORESET}}$ | Power on reset | System control | 1 | I | 4-1/4-1 | — | — |
| $\overline{\text{HRESET}}$ | Hard reset | System control | 1 | I/O | 4-1/4-1 | — | — |
| $\overline{\text{CFG_CLKIN_DIV}}$ | Configuration clock in division selection | Reset and clock | 1 | I | 4-1/4-1 | — | — |
| CFG_RESET_SOURCE[0:3] | Reset configuration word source selection | Reset and clock | 4 | I/O | 4-1/4-1 | TSEC2_TXD[3:0] | 19-2/19-8 |
| $\overline{\text{CKSTOP_IN}}$ | | Reset and clock | 1 | I/O | 4-2/4-3 | IIC_SCL | 20-1/20-3 |

Table 2-1. MPC8315E Signal Reference by Functional Block (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|---------------------------|--------------------------------|------------------|----------------|-----|-----------------------|-----------------------|-------------|
| CKSTOP_OUT | | Reset and clock | 1 | I/O | 4-2/4-3 | IIC_SDA | 20-1/20-3 |
| SYS_XTAL_IN | Crystal clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| SYS_XTAL_OUT | Crystal clock output | Clocks | 1 | O | 4-2/4-3 | — | — |
| SYS_CLK_IN | Clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| USB_XTAL_IN | USB crystal clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| USB_XTAL_OUT | USB crystal clock output | Clocks | 1 | O | 4-2/4-3 | — | — |
| USB_CLK_IN | USB clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| PCI_CLOCK/ PCI_SYNC_IN | PCI clock/PCI clock sync input | Clocks | 1 | I | 4-2/4-3 | — | — |
| PCI_SYNC_OUT | PCI clock sync output | Clocks | 1 | O | 4-2/4-3 | — | — |
| SATA_CLK_IN | SATA clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| RTC_CLK | Timer clock / Real time clock | PIT/RTC | 1 | I | 4-2/4-3 | — | — |
| MSRCID0/LSRCID0 | Memory debug source ID | Debug | 1 | O | 9-1/9-3/ 10-1/10-4 | UART_SOUT1 | 21-1/21-3 |
| MSRCID1/LSRCID1 | Memory debug source ID | Debug | 1 | I/O | 9-1/9-3/ 10-1/10-4 | UART_SIN1 | 21-1/21-3 |
| MSRCID2/LSRCID2 | Memory debug source ID | Debug | 1 | I/O | 9-1/9-3/ 10-1/10-4 | UART_CTS1 | 21-1/21-3 |
| MSRCID3/LSRCID3 | Memory debug source ID | Debug | 1 | I/O | 9-1/9-3/ 10-1/10-4 | UART_RTS1 | 21-1/21-3 |
| MSRCID4/LSRCID4 | Memory debug source ID | Debug | 1 | O | 9-1/9-3/ 10-1/10-4 | UART_SOUT2 | 21-1/21-3 |
| MDVAL/LDVAL | Memory debug data valid | Debug | 1 | I/O | 9-1/9-3 | UART_SIN2 | 21-1/21-3 |

¹ In silicon rev 1.0, NAND boot ECC checking is enabled. In silicon rev 1.1, NAND boot ECC checking is configurable; by default it is enabled (pulled down).

² See Chapter 4, “Reset, Clocking, and Initialization” for proper connection to power.

³ See hardware specification for resistor values.

⁴ Must be connected to a 10K ±1% precision resistor if using the integrated USB PHY through the UTMI.

Table 2-2 lists the signals in alphabetical order.

Table 2-2. MPC8315E Signal Reference by Alphabetical Order

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|-----------------------|---|------------------|----------------|-----|-------------|---|-------------------------|
| CFG_CLKIN_DIV | Configuration clock in division selection | Reset and clock | 1 | I | 4-1/4-1 | — | — |
| CFG_RESET_SOURCE[0:3] | Reset configuration word source selection | Reset and clock | 4 | I/O | 4-1/4-1 | TSEC2_TXD[3:0] | 19-2/19-8 |
| CKSTOP_IN | | Reset and clock | 1 | I/O | 4-2/4-3 | IIC_SCL | 20-1/20-3 |
| CKSTOP_OUT | | Reset and clock | 1 | I/O | 4-2/4-3 | IIC_SDA | 20-1/20-3 |
| CORE_SRESET_IN | Soft reset to the e300 core | System control | 1 | I | 4-1/4-1 | | — |
| CPCI_HS_ENUM | CompactPCI hot swap enumerator | PCI | 1 | O | 13-2/13-4 | PCI_GNT2 | 13-2/13-4 |
| CPCI_HS_ES | CompactPCI hot swap ejector switch | PCI | 1 | I | 13-2/13-4 | PCI_REQ1 | 13-2/13-4 |
| CPCI_HS_LED | CompactPCI hot swap LED | PCI | 1 | O | 13-2/13-4 | PCI_GNT1 | 13-2/13-4 |
| DMA_DACK0 | DMA acknowledge 0 | DMA | 1 | I/O | 12-1/12-2 | GPIO[13] | 24-1/24-2 |
| DMA_DACK1 | DMA acknowledge 1 | DMA | 1 | I/O | 12-1/12-2 | GPIO[1]/ GTM1_TIN2/ GTM2_TIN1 | 24-1/24-2, 5-58/5-56 |
| DMA_DDONE0 | DMA done 0 | DMA | 1 | I/O | 12-1/12-2 | GPIO[14] | 24-1/24-2 |
| DMA_DDONE1 | DMA done 1 | DMA | 1 | I/O | 12-1/12-2 | GPIO[2]/ GTM1_TGATE2/ GTM2_TGATE1 | 24-1/24-2, 5-58/5-56 |
| DMA_DREQ0 | DMA request 0 | DMA | 1 | I/O | 12-1/12-2 | GPIO[12] | 24-1/24-2 |
| DMA_DREQ1 | DMA request 1 | DMA | 1 | I/O | 12-1/12-2 | GPIO[0]/ GTM1_TOUT1 | 24-1/24-2, 5-58/5-56 |
| EXT_PWR_CTRL | External power control | PMC | 1 | O | 5-70/5-70 | — | — |
| GPIO[0] | General-purpose I/O signal 0 | GPIO | 1 | I/O | 24-1/24-2 | GTM1_TOUT1/ DMA_DREQ1 | 5-58/5-56, 12-1/12-2 |
| GPIO[1] | General-purpose I/O signal 1 | GPIO | 1 | I/O | 24-1/24-2 | GTM1_TIN2/ GTM2_TIN1/ DMA_DACK1 | 5-58/5-56, 12-1/12-2 |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|----------|-------------------------------|------------------|----------------|-----|---------------------------|--|---|
| GPIO[10] | General-purpose I/O signal 10 | GPIO | 1 | I/O | 24-1/24-2 | USBD _R _PCTL0/ GTM1_TOUT2/ GTM2_TOUT1 | 17-1/17-3, 5-58/5-56 |
| GPIO[11] | General-purpose I/O signal 11 | GPIO | 1 | I/O | 24-1/24-2 | USBD _R _PCTL1/ GTM1_TOUT4/ GTM2_TOUT3 | 17-1/17-3, 5-58/5-56 |
| GPIO[12] | General-purpose I/O signal 12 | GPIO | 1 | I/O | 24-1/24-2 | DMA_DREQ0 | 12-1/12-2 |
| GPIO[13] | General-purpose I/O signal 13 | GPIO | 1 | I/O | 24-1/24-2 | DMA_DACK0 | 12-1/12-2 |
| GPIO[14] | General-purpose I/O signal 14 | GPIO | 1 | I/O | 24-1/24-2 | DMA_DDONE0 | 12-1/12-2 |
| GPIO[15] | General-purpose I/O signal 15 | GPIO | 1 | I/O | 24-1/24-2 | SPI_MOSI | 22-2/22-7 |
| GPIO[16] | General-purpose I/O signal 16 | GPIO | 1 | I/O | 24-1/24-2 | SPI_MISO | 22-2/22-7 |
| GPIO[17] | General-purpose I/O signal 17 | GPIO | 1 | I/O | 24-1/24-2 | SPI_SEL | 22-2/22-7 |
| GPIO[18] | General-purpose I/O signal 18 | GPIO | 1 | I/O | 24-1/24-2 | TDM_RCK | 25-1/25-2 |
| GPIO[19] | General-purpose I/O signal 19 | GPIO | 1 | I/O | 24-1/24-2 | TDM_RFS | 25-1/25-2 |
| GPIO[2] | General-purpose I/O signal 2 | GPIO | 1 | I/O | 24-1/24-2 | GTM1_TGATE2/ GTM2_TGATE1/ DMA_DDONE1 | 5-58/5-56, 12-1/12-2 |
| GPIO[20] | General-purpose I/O signal 20 | GPIO | 1 | I/O | 24-1/24-2 | TDM_RD | 25-1/25-2 |
| GPIO[21] | General-purpose I/O signal 21 | GPIO | 1 | I/O | 24-1/24-2 | TDM_TCK | 25-1/25-2 |
| GPIO[22] | General-purpose I/O signal 22 | GPIO | 1 | I/O | 24-1/24-2 | TDM_TFS | 25-1/25-2 |
| GPIO[23] | General-purpose I/O signal 23 | GPIO | 1 | I/O | 24-1/24-2 | TDM_TD | 25-1/25-2 |
| GPIO[24] | General-purpose I/O signal 24 | GPIO | 1 | I/O | 24-1/24-2 | USBD _R _TXDRXD[0]/ TSEC1_COL | 17-1/17-3, 19-2/19-8 |
| GPIO[25] | General-purpose I/O signal 25 | GPIO | 1 | I/O | 24-1/24-2 | USBD _R _TXDRXD[1]/ TSEC1_CRS | 17-1/17-3, 19-2/19-8 |
| GPIO[26] | General-purpose I/O signal 26 | GPIO | 1 | I/O | 24-1/24-2 | TSEC2_COL | 19-2/19-8 |
| GPIO[27] | General-purpose I/O signal 27 | GPIO | 1 | I/O | 24-1/24-2 | TSEC2_CRS | 19-2/19-8 |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|-----------------------------|-------------------------------|------------------|----------------|-----|-------------|--|-------------------------|
| GPIO[28] | General-purpose I/O signal 28 | GPIO | 1 | I/O | 24-1/24-2 | TSEC1_TXD[3]/ TSEC_1588_CLK_OUT | 19-2/19-8 |
| GPIO[29] | General-purpose I/O signal 29 | GPIO | 1 | I/O | 24-1/24-2 | TSEC_1588_PULSE_OUT1/ TSEC1_TXD[2] | 19-2/19-8 |
| GPIO[3] | General-purpose I/O signal 3 | GPIO | 1 | I/O | 24-1/24-2 | GTM1_TIN3/ GTM2_TIN4 | 5-58/5-56 |
| GPIO[30] | General-purpose I/O signal 30 | GPIO | 1 | I/O | 24-1/24-2 | TSEC_1588_PULSE_OUT2/ TSEC1_TXD[1] | 19-2/19-8 |
| GPIO[31] | General-purpose I/O signal 31 | GPIO | 1 | I/O | 24-1/24-2 | TSEC1_TX_EN/ TSEC_1588_ALARM_OUT1 | 19-2/19-8 |
| GPIO[4] | General-purpose I/O signal 4 | GPIO | 1 | I/O | 24-1/24-2 | GTM1_TGATE3/ GTM2_TGATE4 | 5-58/5-56 |
| GPIO[5] | General-purpose I/O signal 5 | GPIO | 1 | I/O | 24-1/24-2 | GTM1_TOUT3/ GTM2_TOUT1 | 5-58/5-56 |
| GPIO[6] | General-purpose I/O signal 6 | GPIO | 1 | I/O | 24-1/24-2 | GTM1_TIN4/ GTM2_TIN3 | 5-58/5-56 |
| GPIO[7] | General-purpose I/O signal 7 | GPIO | 1 | I/O | 24-1/24-2 | GTM1_TGATE4/ GTM2_TGATE3 | 5-58/5-56 |
| GPIO[8] | General-purpose I/O signal 8 | GPIO | 1 | I/O | 24-1/24-2 | USBDR_DRIVE_VBUS/ GTM1_TIN1/ GTM2_TIN2 | 17-1/17-3, 5-58/5-56 |
| GPIO[9] | General-purpose I/O signal 9 | GPIO | 1 | I/O | 24-1/24-2 | USBDR_PWRFAULT/ GTM1_TGATE1/ GTM2_TGATE2 | 17-1/17-3, 5-58/5-56 |
| GTM1_TGATE1/ GTM2_TGATE2 | Timer gate 1/2 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[9]/USBDR_PWRFAULT | 24-1/24-2, 17-1/17-3 |
| GTM1_TGATE2/ GTM2_TGATE1 | Timer gate 2/1 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[2]/ DMA_DDONE1 | 24-1/24-2, 12-1/12-2 |
| GTM1_TGATE3/ GTM2_TGATE4 | Timer gate 3/4 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[4] | 24-1/24-2 |
| GTM1_TGATE4/ GTM2_TGATE3 | Timer gate 4/3 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[7] | 24-1/24-2 |
| GTM1_TIN1/ GTM2_TIN2 | Timer in 1/2 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[8]/USBDR_DRIVE_VBUS | 24-1/24-2, 17-1/17-3 |
| GTM1_TIN2/ GTM2_TIN1 | Timer in 2/1 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[1]/ DMA_DACK1 | 24-1/24-2, 12-1/12-2 |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|---------------------------|-------------------------------|------------------|----------------|-----|-------------|--------------------------|-------------------------|
| GTM1_TIN3/ GTM2_TIN4 | Timer in 3/4 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[3] | 24-1/24-2 |
| GTM1_TIN4/ GTM2_TIN3 | Timer in 4/3 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[6] | 24-1/24-2 |
| GTM1_TOUT1 | Timer out 1 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[0]/ DMA_DREQ1 | 24-1/24-2, 12-1/12-2 |
| GTM1_TOUT2/ GTM2_TOUT1 | Timer out 2/1 | Global Timers | 1 | I/O | 5-58/5-56 | USBDR_PCTL0/ GPIO[10] | 17-1/17-3, 24-1/24-2 |
| GTM1_TOUT3/ GTM2_TOUT1 | Timer out 3/1 | Global Timers | 1 | I/O | 5-58/5-56 | GPIO[5] | 24-1/24-2 |
| GTM1_TOUT4/ GTM2_TOUT3 | Timer out 4/3 | Global Timers | 1 | I/O | 5-58/5-56 | USBDR_PCTL1/ GPIO[11] | 17-1/17-3, 24-1/24-2 |
| HRESET | Hard reset | System control | 1 | I/O | 4-1/4-1 | — | — |
| IIC_SCL | I ² C serial clock | I ² C | 1 | I/O | — | CKSTOP_IN | 4-3/4-5 |
| IIC_SDA | I ² C serial data | I ² C | 1 | I/O | — | CKSTOP_OUT | 4-3/4-5 |
| IRQ[0]/MCP_IN | External interrupt 0 | IPIC | 1 | I | 8-1/8-5 | — | — |
| IRQ[1:4] | External interrupt 1–4 | IPIC | 4 | I | 8-1/8-5 | — | — |
| IRQ[5] | External interrupt 5 | IPIC | 1 | I | 8-1/8-5 | CORE_SRESET_IN | 4-1/4-1 |
| IRQ[6] | External interrupt 6 | IPIC | 1 | I/O | 8-1/8-5 | CKSTOP_OUT | 4-3/4-5 |
| IRQ[7] | External interrupt 7 | IPIC | 1 | I | 8-1/8-5 | CKSTOP_IN | 4-3/4-5 |
| LA[16:25] | LBC port address | eLBC | 10 | O | 10-1/10-4 | — | — |
| LAD[0:15] | LBC address/data | eLBC | 16 | I/O | 10-1/10-4 | — | — |
| LALE | LBC address latch enable | eLBC | 1 | O | 10-1/10-4 | — | — |
| LBCTL | LBC data buffer control | eLBC | 1 | O | 10-1/10-4 | — | — |
| LCLK[0:1] | LBC clocks 0–1 | eLBC | 2 | O | 10-1/10-4 | — | — |
| LCS[0:3] | LBC chip select 0–3 | eLBC | 4 | O | 10-1/10-4 | — | — |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|--|--|------------------|----------------|-----|-----------------------|-----------------------|-------------|
| LGPL0 | LBC UPM general purpose line 0/ | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL1/LFALE | LBC GP line 1/Flash address latch enable | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL2 | LBC output enable | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL3 | LBC GP line 3 | eLBC | 1 | O | 10-1/10-4 | — | — |
| LGPL4 | LBC GP line 4 | eLBC | 1 | I/O | 10-1/10-4 | — | — |
| LGPL5 | LBC GP line 5 | eLBC | 1 | O | 10-1/10-4 | — | — |
| $\overline{\text{LWE}}[0:1]$ | LBC write enable | eLBC | 2 | O | 10-1/10-4 | — | — |
| M66EN | 66-MHz system configuration | PCI | 1 | I | 13-2/13-4 | — | — |
| MA[14:0] | DDR address | DDR | 15 | O | 9-1/9-3 | — | — |
| MBA[0:2] | DDR bank select | DDR | 3 | O | 9-1/9-3 | — | — |
| $\overline{\text{MCAS}}$ | DDR column address strobe | DDR | 1 | O | 9-1/9-3 | — | — |
| MCK[0:1], $\overline{\text{MCK}}[0:1]$ | DDR differential clocks | DDR | 4 | O | 9-1/9-3 | — | — |
| MCKE | DDR clock enable | DDR | 1 | O | 9-1/9-3 | — | — |
| $\overline{\text{MCP_OUT}}$ | Machine check interrupt output | IPIC | 1 | O | 8-1/8-5 | — | — |
| $\overline{\text{MCS}}[0:1]$ | DDR chip select (2/DIMM) | DDR | 2 | O | 9-1/9-3 | — | — |
| MDM[0:3] | DDR data mask | DDR | 4 | O | 9-1/9-3 | — | — |
| MDQ[0:31] | DDR data | DDR | 32 | I/O | 9-1/9-3 | — | — |
| MDQS[0:3] | DDR data strobe | DDR | 4 | I/O | 9-1/9-3 | — | — |
| MDVAL/LDVAL | Memory debug data valid | Debug | 1 | I/O | 9-1/9-3 | UART_SIN2 | 21-1/21-3 |
| MODT[0:1] | DRAM on-die termination | DDR | 2 | O | 9-1/9-3 | — | — |
| $\overline{\text{MRAS}}$ | DDR row address strobe | DDR | 1 | O | 9-1/9-3 | — | — |
| MSRCID0/LSRCID0 | Memory debug source ID | Debug | 1 | O | 9-1/9-3/ 10-1/10-4 | UART_SOUT1 | 21-1/21-3 |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|-------------------------------------|--------------------------------|------------------|----------------|-----|-----------------------|-----------------------|-------------|
| MSRCID1/LSRCID1 | Memory debug source ID | Debug | 1 | I/O | 9-1/9-3 | UART_SIN1 | 21-1/21-3 |
| MSRCID2/LSRCID2 | Memory debug source ID | Debug | 1 | I/O | 9-1/9-3/ 10-1/10-4 | UART_CTS1 | 21-1/21-3 |
| MSRCID3/LSRCID3 | Memory debug source ID | Debug | 1 | I/O | 9-1/9-3/ 10-1/10-4 | UART_RTS1 | 21-1/21-3 |
| MSRCID4/LSRCID4 | Memory debug source ID | Debug | 1 | O | 9-1/9-3/ 10-1/10-4 | UART_SOUT2 | 21-1/21-3 |
| $\overline{\text{MWE}}$ | DDR write enable | DDR | 1 | O | 9-1/9-3 | — | — |
| PCI_AD[14] | PCI address/data | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_AD[31:0] | PCI address/data | PCI | 14 | I/O | 13-2/13-4 | — | — |
| PCI_AD[31:15] | PCI address/data | PCI | 17 | I/O | 13-2/13-4 | — | — |
| PCI_C/ $\overline{\text{BE}}$ [3:0] | PCI command/byte enable | PCI | 4 | I/O | 13-2/13-4 | — | — |
| PCI_CLK0 | PCI clock 0 | PCI | 1 | O | 4-2/4-3 | — | — |
| PCI_CLK1 | PCI clock 1 | PCI | 1 | O | 4-2/4-3 | — | — |
| PCI_CLK2 | PCI clock 2 | PCI | 1 | O | 4-2/4-3 | — | — |
| PCI_CLOCK/ PCI_SYNC_IN | PCI clock/PCI clock sync input | Clocks | 1 | I | 4-2/4-3 | — | — |
| $\overline{\text{PCI_DEVSEL}}$ | PCI device select | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_FRAME}}$ | PCI frame | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_GNT0}}$ | PCI grant 0 | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_GNT1}}$ | PCI grant 1 | PCI | 1 | O | 13-2/13-4 | CPCI_HS_LED | 13-2/13-4 |
| $\overline{\text{PCI_GNT2}}$ | PCI grant 2 | PCI | 1 | O | 13-2/13-4 | CPCI_HS_ENUM | 13-2/13-4 |
| $\overline{\text{PCI_IDSEL}}$ | PCI initial device select | PCI | 1 | I | 13-2/13-4 | — | — |
| $\overline{\text{PCI_INTA}}$ | PCI interrupt output | PCI | 1 | O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_IRDY}}$ | PCI initiator ready | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_PAR | PCI parity | PCI | 1 | I/O | 13-2/13-4 | — | — |
| $\overline{\text{PCI_PERR}}$ | PCI parity error | PCI | 1 | I/O | 13-2/13-4 | — | — |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|---------------|---|------------------|----------------|-----|-------------|-----------------------|-------------|
| PCI_PME | PCI PME assertion request | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_REQ0 | PCI request 0 | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_REQ1 | PCI request 1 | PCI | 1 | I | 13-2/13-4 | CPCI_HS_ES | 13-2/13-4 |
| PCI_REQ2 | PCI request 2 | PCI | 1 | I | 13-2/13-4 | — | — |
| PCI_RESET_OUT | PCI reset | PCI | 1 | O | 13-2/13-4 | — | — |
| PCI_SERR | PCI system error | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_STOP | PCI stop | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PCI_SYNC_OUT | PCI clock sync output | Clocks | 1 | O | 4-2/4-3 | — | — |
| PCI_TRDY | PCI target ready | PCI | 1 | I/O | 13-2/13-4 | — | — |
| PINRXMINUSA | Receive minus A | SATA PHY | 1 | I | — | — | — |
| PINRXMINUSB | Receive minus B | SATA PHY | 1 | I | — | — | — |
| PINRXPLUSA | Receive plus A | SATA PHY | 1 | I | — | — | — |
| PINRXPLUSB | Receive plus B | SATA PHY | 1 | I | — | — | — |
| PINTXMINUSA | Transmit minus A | SATA PHY | 1 | O | — | — | — |
| PINTXMINUSB | Transmit minus B | SATA PHY | 1 | O | — | — | — |
| PINTXPLUSA | Transmit plus A | SATA PHY | 1 | O | — | — | — |
| PINTXPLUSB | Transmit plus A | SATA PHY | 1 | O | — | — | — |
| PMC_PWR_OK | Stable power | PMC | 1 | I | 5-70/5-70 | — | — |
| PORESET | Power on reset | System control | 1 | I | 4-1/4-1 | — | — |
| QUIESCE | Quiesce state | PMC | 1 | O | 5-70/5-70 | — | — |
| RESREF | Reset reference | SATA PHY | 1 | I | — | — | — |
| RTC_CLK | Timer clock / Real time clock | PIT/RTC | 1 | I | 4-2/4-3 | — | — |
| RXA | Serial receiver, lane A, positive data | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| RXA | Serial receiver, lane A, negative data (complement) | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| RXB | Serial receiver, lane B, positive data | SGMII PHY | 1 | I | 16-1/16-3 | — | — |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|----------------------------------|---|------------------------|----------------|-----|---------------------------|-----------------------|---------------------------|
| $\overline{\text{RXB}}$ | Serial receiver, lane B, negative data (complement) | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| SATA_ANAVIZ | | SATA PHY | 1 | O | — | — | — |
| SATA_CLK_IN | Clock in | SATA PHY | 1 | I | — | — | — |
| SATA_CLK_IN | SATA clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| SATA_VDD | Voltage supply | SATA PHY | 2 | I | — | — | — |
| SATA_VSS | Ground | SATA PHY | 2 | I | — | — | — |
| SD_IMP_CAL_RX | Receiver impedance control signal | SGMII PHY ¹ | 1 | I | 16-1/16-3 | — | — |
| SD_IMP_CAL_TX | Transmitter impedance control signal | SGMII PHY ³ | 1 | I | 16-1/16-3 | — | — |
| SD_PLL_TPA_ANA | Analog test point for SerDes PLL testing | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| SD_PLL_TPD | Digital test point for SerDes PLL testing | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| SD_REF_CLK | SerDes PLL reference clock | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| $\overline{\text{SD_REF_CLK}}$ | SerDes PLL reference clock (complement) | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| SDAVDD | Analog supply for SerDes PLL | SGMII PHY ² | 1 | I | 16-1/16-3 | — | — |
| SDAVSS | Analog ground for SerDes PLL | SGMII PHY | 1 | I | 16-1/16-3 | — | — |
| SPICLK | SPI clock | SPI | 1 | I/O | 22-1/22-7 | — | — |
| SPIMISO | SPI master-in slave-out | SPI | 1 | I/O | 22-1/22-7 | GPIO[16] | 24-1/24-2 |
| SPIMOSI | SPI master-out slave-in | SPI | 1 | I/O | 22-1/22-7 | GPIO[15] | 24-1/24-2 |
| $\overline{\text{SPISEL}}$ | SPI slave select | SPI | 1 | I/O | 22-1/22-7 | GPIO[17] | 24-1/24-2 |
| SYS_CLK_IN | Clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| SYS_XTAL_IN | Crystal clock input | Clocks | 1 | I | 4-2/4-3 | — | — |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|--------------------------|-------------------------|-------------------|----------------|-----|-------------|----------------------------------|-------------------------|
| SYS_XTAL_OUT | Crystal clock output | Clocks | 1 | O | 4-2/4-3 | — | — |
| TCK | Test clock | JTAG | 1 | I | 23-1/23-2 | — | — |
| TDI | Test data in | JTAG | 1 | I | 23-1/23-2 | — | — |
| TDM_RCK | TDM receive clock | TDM | 1 | I/O | 25-1/25-2 | GPIO[18] | 24-1/24-2 |
| TDM_RD | TDM receive data | TDM | 1 | I/O | 25-1/25-2 | GPIO[20] | 24-1/24-2 |
| TDM_RFS | TDM receive frame sync | TDM | 1 | I/O | 25-1/25-2 | GPIO[19] | 24-1/24-2 |
| TDM_TCK | TDM transmit clock | TDM | 1 | I/O | 25-1/25-2 | GPIO[21] | 24-1/24-2 |
| TDM_TD | TDM transmit data | TDM | 1 | I/O | 25-1/25-2 | GPIO[23] | 24-1/24-2 |
| TDM_TFS | TDM transmit frame sync | TDM | 1 | I/O | 25-1/25-2 | GPIO[22] | 24-1/24-2 |
| TDO | Test data out | JTAG | 1 | O | 23-1/23-2 | — | — |
| TMS | Test mode select | JTAG | 1 | I | 23-1/23-2 | — | — |
| $\overline{\text{TRST}}$ | Test reset | JTAG | 1 | I | 23-1/23-2 | — | — |
| TSEC_1588_ALARM_OUT1 | 1588 timer alarm-out 1 | eTSEC | 1 | I/O | 19-2/19-8 | TSEC1_TX_EN/ GPIO[31] | 19-2/19-8, 17-1/17-3 |
| TSEC_1588_ALARM_OUT2 | 1588 timer alarm-out 2 | eTSEC | 1 | O | 19-2/19-8 | TSEC1_TX_ER | 19-2/19-8 |
| TSEC_1588_CLK_IN | 1588 clock-in | eTSEC | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD [7]/TSEC1_RXD[1] | 17-1/17-3, 19-2/19-8 |
| TSEC_1588_CLK_OUT | 1588 clock-out | eTSEC | 1 | I/O | 19-2/19-8 | TSEC1_TXD[3]/ GPIO[28] | 19-2/19-8, 24-1/24-2 |
| TSEC_1588_PULSE_OUT1 | 1588 timer pulse-out 1 | eTSEC | 1 | I/O | 19-2/19-8 | TSEC1_TXD[2]/ GPIO[29] | 19-2/19-8, 24-1/24-2 |
| TSEC_1588_PULSE_OUT2 | 1588 timer pulse-out 2 | eTSEC | 1 | I/O | 19-2/19-8 | TSEC1_TXD[1]/ GPIO[30] | 19-2/19-8, 24-1/24-2 |
| TSEC_1588_PULSE_OUT3 | 1588 timer pulse-out 3 | eTSEC | 1 | O | 19-2/19-8 | TSEC1_TXD[0]/ USBDR_STP | 19-2/19-8, 17-1/17-3 |
| TSEC_1588_TRIG_IN1 | 1588 trigger-in 1 | eTSEC | 1 | I | 19-2/19-8 | USBDR_NXT/ TSEC1_RXD[0] | 17-1/17-3, 19-2/19-8 |
| TSEC_1588_TRIG_IN2 | 1588 trigger-in 2 | eTSEC | 1 | I/O | 19-2/19-8 | USBDR_DIR/ TSEC1_RX_ER | 17-1/17-3, 19-2/19-8 |
| TSEC_GTX_CLK125 | Gigabit reference clock | eTSEC1/ eTSEC2 | 1 | I | 19-2/19-8 | — | — |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|----------------|---------------------------------|---------------------|----------------|-----|----------------|----------------------------------|----------------------|
| TSEC_MDC | Ethernet management data clock | Ethernet management | 1 | O | 19-2/19-8 | LB_POR_CFG_BOOT_ECC ³ | — |
| TSEC_MDIO | Ethernet management data in/out | Ethernet management | 1 | I/O | 19-2/19-8 | — | — |
| TSEC1_COL | eTSEC1 collision detect | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[0]/GPIO[24] | 17-1/17-3, 24-1/24-2 |
| TSEC1_CRS | eTSEC1 carrier sense | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[1]/GPIO[25] | 17-1/17-3, 24-1/24-2 |
| TSEC1_GTX_CLK | eTSEC1 transmit clock out | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[2] | 17-1/17-3 |
| TSEC1_RX_CLK | eTSEC1 receive clock | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[3] | 17-1/17-3 |
| TSEC1_RX_DV | eTSEC1 receive data valid | eTSEC1 | 1 | I/O | 19-2/19-8 | USBDR_TXDRXD[4] | 17-1/17-3 |
| TSEC1_RX_ER | eTSEC1 receiver error | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_TRIG_IN2/USBDR_DIR | 19-2/19-8, 17-1/17-3 |
| TSEC1_RXD[0] | eTSEC1 receive data 0 | eTSEC1 | 1 | I | 19-2/19-8 | TSEC_1588_TRIG_IN1/USBDR_NXT | 19-2/19-8, 17-1/17-3 |
| TSEC1_RXD[1] | eTSEC1 receive data 1 | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_CLK_IN/USBDR_TXDRXD[7] | 19-2/19-8, 17-1/17-3 |
| TSEC1_RXD[3:2] | eTSEC1 receive data 3–2 | eTSEC1 | 2 | I/O | 19-2/19-8 | USBDR_TXDRXD [5:6] | 17-1/17-3 |
| TSEC1_TX_CLK | eTSEC1 transmit clock in | eTSEC1 | 1 | I | 19-2/19-8 | USBDR_CLK | 17-1/17-3 |
| TSEC1_TX_EN | eTSEC1 transmit enable | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_ALARM_OUT1/GPIO[31] | 19-2/19-8, 17-1/17-3 |
| TSEC1_TX_ER | eTSEC1 transmit error | eTSEC1 | 1 | O | 19-2/19-8 | TSEC_1588_ALARM_OUT1 | 19-2/19-8 |
| TSEC1_TXD[0] | eTSEC1 transmit data 0 | eTSEC1 | 1 | O | 19-2/19-8 | TSEC_1588_PULSE_OUT3/USBDR_STP | 19-2/19-8, 17-1/17-3 |
| TSEC1_TXD[1] | eTSEC1 transmit data 1 | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_PULSE_OUT2/GPIO[30] | 19-2/19-8, 24-1/24-2 |
| TSEC1_TXD[2] | eTSEC1 transmit data 2 | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_PULSE_OUT1/GPIO[29] | 19-2/19-8, 24-1/24-2 |
| TSEC1_TXD[3] | eTSEC1 transmit data 3 | eTSEC1 | 1 | I/O | 19-2/19-8 | TSEC_1588_CLK_OUT/GPIO[28] | 19-2/19-8, 24-1/24-2 |
| TSEC2_COL | eTSEC2 collision detect | eTSEC2 | 1 | I/O | 19-2/19-8 | GPIO[26] | 24-1/24-2 |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|--------------------------------|--|------------------|----------------|-----|-------------|-----------------------|-----------------------|
| TSEC2_CRS | eTSEC2 carrier sense | eTSEC2 | 1 | I/O | 19-2/19-8 | GPIO[27] | 24-1/24-2 |
| TSEC2_GTX_CLK | eTSEC2 transmit clock out | eTSEC2 | 1 | O | 19-2/19-8 | — | — |
| TSEC2_RX_CLK | eTSEC2 receive clock | eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC2_RX_DV | eTSEC2 receive data valid | eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC2_RX_ER | eTSEC2 receiver error | eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC2_RXD[3:0] | eTSEC2 receive data 3–0 | eTSEC2 | 4 | I | 19-2/19-8 | — | — |
| TSEC2_TX_CLK | eTSEC2 transmit clock in | eTSEC2 | 1 | I | 19-2/19-8 | — | — |
| TSEC2_TX_EN | eTSEC2 transmit enable | eTSEC2 | 1 | O | 19-2/19-8 | — | — |
| TSEC2_TX_ER | eTSEC2 transmit error | eTSEC2 | 1 | O | 19-2/19-8 | — | — |
| TSEC2_TXD[3:0] | eTSEC2 transmit data 3–0 | eTSEC2 | 4 | I/O | 19-2/19-8 | CFG_RESET_SOURCE[0:3] | 4-1/4-1 |
| TXA | Serial transmitter, lane A, positive data | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| $\overline{\text{TXA}}$ | Serial transmitter, lane A, negative data (complement) | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| TXB | Serial transmitter, lane B, positive data | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| $\overline{\text{TXB}}$ | Serial transmitter, lane B, negative data (complement) | SGMII PHY | 1 | O | 16-1/16-3 | — | — |
| $\overline{\text{UART_CTS1}}$ | DUART clear to send | DUART | 1 | I/O | 21-1/21-3 | MSRCID2/ LSRCID2 | 9-1/9-3, 10-1/10-4 |
| $\overline{\text{UART_CTS2}}$ | DUART clear to send | DUART | 1 | I/O | 21-1/21-3 | — | — |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|------------------|--|------------------|----------------|-----|-------------|--------------------------------------|-------------------------|
| UART_RTS1 | DUART ready to send | DUART | 1 | I/O | 21-1/21-3 | MSRCID3/ LSRCID3 | 9-1/9-3, 10-1/10-4 |
| UART_RTS2 | DUART ready to send | DUART | 1 | I/O | 21-1/21-3 | — | — |
| UART_SIN1 | DUART serial data in | DUART | 1 | I/O | 21-1/21-3 | MSRCID1/LSRCID1 | 9-1/9-3, 10-1/10-4 |
| UART_SIN2 | DUART serial data in | DUART | 1 | I/O | 21-1/21-3 | MDVAL/LDVAL | 9-1/9-3, 10-1/10-4 |
| UART_SOUT1 | DUART serial data out | DUART | 1 | O | 21-1/21-3 | MSRCID0/LSRCID0 | 9-1/9-3, 10-1/10-4 |
| UART_SOUT2 | DUART serial data out | DUART | 1 | O | 21-1/21-3 | MSRCID4/LSRCID4 | 9-1/9-3, 10-1/10-4 |
| USB_CLK_IN | USB clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| USB_DM | USB 2.0 D– line | USB PHY | 1 | I/O | 17-1/17-3 | — | — |
| USB_DP | USB 2.0 D+ line | USB PHY | 1 | I/O | 17-1/17-3 | — | — |
| USB_PLL_PWR1 | Dedicated 1.0 V analog power for USB PLL | USB PHY | 1 | I | 17-1/17-3 | — | — |
| USB_RBIA5 | USB bias input ⁴ | USB_PHY | 1 | I | 17-1/17-3 | — | — |
| USB_TPA | Dedicated analog test signal | USB PHY | 1 | I/O | 17-1/17-3 | — | — |
| USB_VBUS | USB 2.0 VBUS line | USB PHY | 1 | I/O | 17-1/17-3 | — | — |
| USB_XTAL_IN | USB crystal clock input | Clocks | 1 | I | 4-2/4-3 | — | — |
| USB_XTAL_OUT | USB crystal clock output | Clocks | 1 | O | 4-2/4-3 | — | — |
| USBDR_CLK | Clocking signal for ULPI PHY interface | USB | 1 | I | 17-1/17-3 | TSEC1_TX_CLK | 19-2/19-8 |
| USBDR_DIR | Direction of data bus | USB | 1 | I/O | 17-1/17-3 | TSEC_1588_TRIG_IN2/TSEC1_RX_ER | 19-2/19-8 |
| USBDR_DRIVE_VBUS | USB VBus power enable | USB | 1 | I/O | 17-1/17-3 | GTM1_TIN1/ GTM2_TIN2/ GPIO[8]/ | 5-58/5-56, 24-1/24-2 |
| USBDR_NXT | Next data | USB | 1 | I | 17-1/17-3 | TSEC_1588_TRIG_IN1/TSEC1_RXD[0] | 19-2/19-8 |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|-------------------|--------------------------------|------------------------|----------------|-----|-------------|---|-------------------------|
| USBDR_PCTL0 | Port control 0 | USB | 1 | I/O | 17-1/17-3 | GTM1_TOUT2/ GTM2_TOUT1/ GPIO[10] | 5-58/5-56, 24-1/24-2 |
| USBDR_PCTL0 | USB status LED 0 control | USB | 1 | I/O | 17-1/17-3 | GTM1_TOUT2/ GTM2_TOUT1/ GPIO[10] | 5-58/5-56, 24-1/24-2 |
| USBDR_PCTL1 | Port control 1 | USB | 1 | I/O | 17-1/17-3 | GTM1_TOUT4/ GTM2_TOUT3/ GPIO[11] | 5-58/5-56, 24-1/24-2 |
| USBDR_PCTL1 | USB status LED 0 control | USB | 1 | I/O | 17-1/17-3 | GTM1_TOUT4/ GTM2_TOUT3/ GPIO[11] | 5-58/5-56, 24-1/24-2 |
| USBDR_PWRFAULT | USB VBus power fault | USB | 1 | I/O | 17-1/17-3 | GTM1_TGATE1/ GTM2_TGATE2/ GPIO[0] | 5-58/5-56, 24-1/24-2 |
| USBDR_STP | End of a transfer on the bus | USB | 1 | O | 17-1/17-3 | TSEC1_TXD[0]/ TSEC_1588_PULSE_OUT3 | 19-2/19-8 |
| USBDR_TXDRXD[0] | Data bit 0 | USB | 1 | I/O | 17-1/17-3 | TSEC1_COL/ GPIO[24] | 19-2/19-8, 24-1/24-2 |
| USBDR_TXDRXD[1] | Data bit 1 | USB | 1 | I/O | 17-1/17-3 | TSEC1_CRS/ GPIO[25] | 19-2/19-8, 24-1/24-2 |
| USBDR_TXDRXD[2] | Data bit 2 | USB | 1 | I/O | 17-1/17-3 | TSEC1_GTX_CLK | 19-2/19-8 |
| USBDR_TXDRXD[3] | Data bit 3 | USB | 1 | I/O | 17-1/17-3 | TSEC1_RX_CLK | 19-2/19-8 |
| USBDR_TXDRXD[4] | Data bit 4 | USB | 1 | I/O | 17-1/17-3 | TSEC1_RX_DV | 19-2/19-8 |
| USBDR_TXDRXD[5:6] | Data bit 5–6 | USB | 2 | I/O | 17-1/17-3 | TSEC1_RXD[3:2] | 19-2/19-8 |
| USBDR_TXDRXD[7] | Data bit 7 | USB | 1 | I/O | 17-1/17-3 | TSEC1_RXD[1]/ TSEC_1588_CLK_IN | 19-2/19-8 |
| VDD33ANA | Analog 3.3 V supply | SATA PHY | 1 | I | — | — | — |
| VDD33PLL | PLL 3.3 V supply | SATA PHY | 1 | I | — | — | — |
| VSSRESREF | Reset reference ground | SATA PHY | 1 | I | — | — | — |
| XCOREVDD[0:3] | SerDes transceiver core supply | SGMII PHY ² | 4 | PWR | 16-1/16-3 | — | — |
| XCOREVSS[0:3] | SerDes transceiver core ground | SGMII PHY | 4 | GND | 16-1/16-3 | — | — |

Table 2-2. MPC8315E Signal Reference by Alphabetical Order (continued)

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|--------------|-------------------------------|------------------------|----------------|-----|-------------|-----------------------|-------------|
| XPADVDD[0:2] | SerDes transceiver pad supply | SGMII PHY ² | 3 | PWR | 16-1/16-3 | — | — |
| XPADVSS[0:2] | SerDes transceiver pad ground | SGMII PHY | 3 | GND | 16-1/16-3 | — | — |

¹ See hardware specification for resistor values.

² See [Chapter 4, “Reset, Clocking, and Initialization”](#) for proper connection to power.

³ In silicon rev 1.0, NAND boot ECC checking is enabled. In silicon rev 1.1, NAND boot ECC checking is configurable; by default it is enabled (pulled down).

⁴ Must be connected to a 10K ±1% precision resistor if using the integrated USB PHY through the UTMI.

2.2 Output Signal States During Reset

When a system reset is recognized ($\overline{\text{PORESET}}$ or $\overline{\text{HRESET}}$ are asserted), the device aborts all current internal and external transactions and releases all bidirectional I/O signals to a high-impedance state. See [Chapter 4, “Reset, Clocking, and Initialization,”](#) for a complete description of the reset functionality.

During reset, the device ignores most input signals (except for the reset configuration signals) and drives most of the output-only signals to an inactive state. [Table 2-3](#) shows the states of the output-only signals.

Table 2-3. Output Signal States During System Reset

| Interface | Signal | State During Reset |
|-------------------------------------|---------------------------|--------------------|
| MDM[0:3] | DDR data mask | High-Z |
| MBA[0:2] | DDR bank select | High-Z |
| MA[14:0] | DDR address | High-Z |
| $\overline{\text{MWE}}$ | DDR write enable | High-Z |
| $\overline{\text{MRAS}}$ | DDR row address strobe | High-Z |
| $\overline{\text{MCAS}}$ | DDR column address strobe | High-Z |
| $\overline{\text{MCS}}[0:1]$ | DDR chip select (2/DIMM) | High-Z |
| MCKE | DDR clock enable | Driven |
| MCK[0:1] | DDR differential clocks | Driven Toggling |
| $\overline{\text{MCK}}[0:1]$ | DDR differential clocks | Driven Toggling |
| MODT[0:1] | DRAM on-die termination | Driven Low |
| $\overline{\text{PCI_INTA}}$ | PCI interrupt output | High-Z |
| $\overline{\text{PCI_RESET_OUT}}$ | PCI reset output | High-Z |
| $\overline{\text{PCI_GNT}}[0:2]$ | PCI grant | High-Z |

Table 2-3. Output Signal States During System Reset (continued)

| Interface | Signal | State During Reset |
|--------------------------------------|--------------------------------|--|
| UART_SOUT[0:1] | DUART serial data out | High-Z |
| $\overline{\text{UART_RTS}}[0:1]$ | DUART ready to send | High-Z |
| TSEC1_GTX_CLK | TSEC1 transmit clock out | High-Z |
| TSEC1_TX_EN | TSEC1 transmit enable | Driven Low |
| TSEC2_GTX_CLK | TSEC2 transmit clock out | High-Z |
| LA[16:25] | LBC port address | Active—used to load reset configuration word |
| $\overline{\text{LCS}}[0]$ | LBC chip select 0 | Active—used to load reset configuration word |
| $\overline{\text{LCS}}[1:3]$ | LBC chip select | High-Z |
| $\overline{\text{LWE}}[0:1]$ | LBC write enable | High-Z |
| LBCTL | LBC data buffer control | Active—used to load reset configuration word |
| LALE | LBC address latch enable | Active—used to load reset configuration word |
| $\overline{\text{LOE}}/\text{LGPL2}$ | LBC output enable/GP line 2 | Active—used to load reset configuration word |
| LCLK[0:1] | LBC clock | High-Z |
| $\overline{\text{MCP_OUT}}$ | Machine check interrupt output | High-Z |
| TDO | Test data out | High-Z |
| $\overline{\text{QUIESCE}}$ | Quiesce state | High-Z |
| PCI_SYNC_OUT | PCI sync output | Driven Low |
| USBDR_STP_SUSPEND ¹ | USB host 0 data stop/suspend | High-Z |

¹ This pin should be pulled high during a hard reset for proper functionality of the device since it has a weak internal pull up. No external pull-down resistors are allowed to be mounted on this net.

Table 2-4 provides the list of registers that control the device’s signal multiplexing.

Table 2-4. Signals for Multiplexing

| Signal Group | Multiplexing is Controlled By | Table/Page |
|--------------|-------------------------------|---|
| PCI/CPCI | RCWH[PCIARB] | 4.3.2.2/4-15 |
| All others | SICRL/SICRH | 5.3.2.5/5-23/5.3.2.6/5-26 |

Chapter 3

Memory Map

This chapter describes the MPC8315E memory map. The internal memory-mapped registers are described, including a complete listing of all memory-mapped registers with cross references to the sections detailing descriptions of each.

3.1 Internal Memory-Mapped Registers

All of the memory-mapped registers in the device are contained within a 1-Mbyte address region. To allow for flexibility, the base address of the memory-mapped registers is relocatable in the local address space. The local address map location of this register block is controlled by the internal memory-mapped registers base address register (IMMRBAR); see [Appendix A 5.2.4.1, “Internal Memory Map Registers Base Address Register \(IMMRBAR\)”](#) for more information. The default value for IMMRBAR is 0xFF40_0000.

3.2 Accessing IMMR Memory From the Local Processor

When the local e300 processor is used to configure IMMR space, the IMMR memory space should typically be marked as cache-inhibited and guarded.

In addition, many configuration registers affect accesses to other memory regions; therefore, writes to these registers must be guaranteed to have taken effect before accesses are made to the associated memory regions.

To guarantee that the results of any sequence of writes to configuration registers are in effect, the final configuration register write should be followed immediately by a read of the same register, and that should be followed by a sync instruction. Then accesses can safely be made to memory regions affected by the configuration register write.

3.3 IMMR Address Map

[Table 3-1](#) lists the location of the functional block base addresses for the entire IMMRBAR space. Unless stated otherwise in a particular block, all accesses to and from the memory-mapped registers must be made with 32-bit accesses. There is no support for accesses of sizes other than 32 bits.

Reading from address locations which appear as reserved in the memory map table is not guaranteed to return predictable data. Writing to address locations which appear as reserved in the memory map table is not allowed and could lead to unpredictable behavior of the device. Reserved bits in non-reserved registers will be read as zero unless the reset value of those bits is different due to internal logic considerations.

Memory Map

When writing to registers with reserved bits, those reserved bits should be cleared. By doing so, existing software would be able to run on a future modified device in which some reserved bits were allocated for enhanced modes. This would allow for maintaining the legacy functionality when set to zero.

In certain specific cases, reserved bits should not be cleared but should keep their reset value. Thus, the software should perform a ‘read-modify-write’ and make sure that it does not change the reset value of those bits. The description of the specific bits will indicate when this is needed.

Cross-references are provided to the IMMRBAR maps for each individual block. A complete listing of all registers is provided in [Appendix A, “Complete List of Configuration, Control, and Status Registers.”](#)

Table 3-1. IMMR Memory Map

| Block Base Address | Block | Section/Page | Comments |
|--------------------|--|-----------------------------|--|
| 0x0_0000 | Local access window | 5.2.3.1/5-5 | — |
| | System configuration | 5.3.1/5-19 | — |
| 0x0_0200 | Watchdog timer (WDT) | 5.4.4/5-34 | — |
| 0x0_0300 | Real time clock (RTC) | 5.5.5/5-41 | — |
| 0x0_0400 | Periodic interval timer (PIT) | 5.6.5/5-48 | — |
| 0x0_0500 | General purpose (global) timers (GTMs) | 5.7.5/5-57 | Global timers module 1: 0x0_0500 Global timers module 2: 0x0_0600 |
| 0x0_0700 | Integrated programable interrupt controller (IPIC) | 8.5/8-6 | — |
| 0x0_0800 | System arbiter | 6.2/6-2 | — |
| 0x0_0900 | Reset configuration | 4.5.1/4-33 | — |
| 0x0_0A00 | Clock configuration | 4.5.2/4-38 | — |
| 0x0_0B00 | Power management controller (PMC) | 5.8.2/5-71 | — |
| 0x0_0C00 | General purpose I/O (GPIO) | 24.3/24-2 | — |
| 0x0_0D00 | Reserved | — | — |
| 0x0_2000 | DDR memory controller | 9.4/9-8 | — |
| 0x0_3000 | I ² C controller | 20.3/20-4 | — |
| 0x0_3100 | Reserved | — | — |
| 0x0_4500 | Dual UART (DUART) | 21.3/21-4 | UART1: 0x0_4500 UART2: 0x0_4600 |
| 0x0_5000 | Enhanced local bus controller (eLBC) | 10.3/10-8 | — |
| 0x0_6000 | Reserved | — | — |
| 0x0_7000 | Serial peripheral interface (SPI) | 22.4/22-8 | — |
| 0x0_8000 | DMA controller | 12.3/12-3 | — |
| 0x0_8300 | PCI configuration access | 13.3/13-11 | — |
| 0x0_8380 | Reserved | — | — |
| 0x0_8400 | I/O sequencer (IOS) | 11.3/11-2 | — |

Table 3-1. IMMR Memory Map (continued)

| Block Base Address | Block | Section/Page | Comments |
|--------------------|--|---|--|
| 0x0_8500 | PCI controller | 13.3/13-11 | — |
| 0x0_8600 | Reserved | — | — |
| 0x0_9000 | PCI Express controller | 14.3.1/14-5 | PCI Express 1: 0x0_9000 PCI Express 2: 0x0_A000 |
| 0x0_B000 | Reserved | — | — |
| 0x1_6000 | Time division multiplexing (TDM) interface | 25.4.2/25-7 , 25.4.4/25-28 , 25.9.3/25-52 | TDM configuration: 0x1_6000 TDM data: 0x1_6100 TDM clock control: 0x1_6180 |
| 0x1_6200 | Reserved | — | — |
| 0x1_8000 | Serial ATA (SATA) controller | 15.3.1/15-4 | SATA1: 0x1_8000 SATA2: 0x1_9000 |
| 0x1_A000 | Reserved | — | — |
| 0x2_3000 | USB dual-role (DR) controller | 17.3/17-6 | — |
| 0x2_4000 | Enhanced three-speed Ethernet controller (eTSEC) | 19.5/19-11 | eTSEC 1: 0x2_4000 eTSEC 2: 0x2_5000 |
| 0x2_6000 | Reserved | — | — |
| 0x2_C000 | TDM DMA controller (DMAC) | 25.10.2/25-55 | — |
| 0x2_E000 | Reserved | — | — |
| 0x3_0000 | Security engine controller (SEC) | 18.2/18-10 | — |
| 0x4_0000 | Reserved | — | — |
| 0xE_3000 | SerDes PHY (PCI Express/SGMII) | 16.3/16-4 | — |
| 0xE_3100 | Reserved | — | — |

Chapter 4

Reset, Clocking, and Initialization

The reset, clocking, and control signals offer many options for operating the device. Various modes and features can be configured during hard reset or power-on reset. Most configurable features are loaded to the device through a reset configuration word, and a few device signals are used as reset configuration inputs during the reset sequence.

4.1 External Signals

The following sections describe the reset and clock signals in detail.

4.1.1 Reset Signals

Table 4-1 describes the reset signals of the device. Section 4.3.2, “Reset Configuration Words,” describes the signals that also function as reset configuration signals.

Table 4-1. System Control Signals

| Signal | I/O | Description | | |
|---|---|---|---|---|
| $\overline{\text{PORESET}}$ | I | Power-on reset. Initiates the power-on reset flow that resets the device and configures various attributes of the device, including its clock modes. | | |
| | | <table border="1"> <tr> <td>State Meaning</td> <td>Asserted—An external agent has triggered a power-on reset sequence. Negated—No power-on reset.</td> </tr> </table> | State Meaning | Asserted—An external agent has triggered a power-on reset sequence. Negated—No power-on reset. |
| | | State Meaning | Asserted—An external agent has triggered a power-on reset sequence. Negated—No power-on reset. | |
| | | <table border="1"> <tr> <td>Timing</td> <td>See the hardware specifications for timing information.</td> </tr> </table> | Timing | See the hardware specifications for timing information. |
| Timing | See the hardware specifications for timing information. | | | |
| <table border="1"> <tr> <td>Reset State</td> <td>Always input.</td> </tr> </table> | Reset State | Always input. | | |
| Reset State | Always input. | | | |
| $\overline{\text{HRESET}}$ | I/O | Hard reset. Causes the device to abort all current internal and external transactions and set most registers to their default values. $\overline{\text{HRESET}}$ can be asserted completely asynchronously with respect to all other signals. The device can detect an external assertion of $\overline{\text{HRESET}}$ while the device is not asserting hard reset. $\overline{\text{HRESET}}$ is an open-drain signal. | | |
| | | <table border="1"> <tr> <td>State Meaning</td> <td>Asserted—An external agent or internal hardware has triggered a hard reset. The internal hardware drives $\overline{\text{HRESET}}$ until the sequence completes. Negated—No hard reset.</td> </tr> </table> | State Meaning | Asserted—An external agent or internal hardware has triggered a hard reset. The internal hardware drives $\overline{\text{HRESET}}$ until the sequence completes. Negated—No hard reset. |
| | | State Meaning | Asserted—An external agent or internal hardware has triggered a hard reset. The internal hardware drives $\overline{\text{HRESET}}$ until the sequence completes. Negated—No hard reset. | |
| | | <table border="1"> <tr> <td>Timing</td> <td>Assertion—Occur at any time, asynchronously to any clock. Negation—Must be asserted for at least 32 SYS_CLK_IN (PCI host mode) or PCI_CLK (PCI agent mode) cycles.</td> </tr> </table> | Timing | Assertion—Occur at any time, asynchronously to any clock. Negation—Must be asserted for at least 32 SYS_CLK_IN (PCI host mode) or PCI_CLK (PCI agent mode) cycles. |
| | | Timing | Assertion—Occur at any time, asynchronously to any clock. Negation—Must be asserted for at least 32 SYS_CLK_IN (PCI host mode) or PCI_CLK (PCI agent mode) cycles. | |
| <table border="1"> <tr> <td>Requirements</td> <td>An open-drain signal. An external pull-up is required.</td> </tr> </table> | Requirements | An open-drain signal. An external pull-up is required. | | |
| Requirements | An open-drain signal. An external pull-up is required. | | | |
| <table border="1"> <tr> <td>Reset State</td> <td>Output, driven low during power-on and hard reset flows. High impedance after reset flow completes.</td> </tr> </table> | Reset State | Output, driven low during power-on and hard reset flows. High impedance after reset flow completes. | | |
| Reset State | Output, driven low during power-on and hard reset flows. High impedance after reset flow completes. | | | |

Table 4-1. System Control Signals (continued)

| Signal | I/O | Description | | |
|---|--|--|---|---|
| CORE_SRESET_IN | I | The CORE_SRESET_IN input is connected directly to the soft reset input of the e300c3 core. The assertion of the e300 soft reset input, CORE_SRESET_IN, causes a high priority interrupt to the e300 core as described in Section 4.2.1, “Reset Operations.” It does not reset the e300 core or any other portion of the device. The CORE_SRESET_IN input is not registered in the reset status register (RSR). | | |
| | | <table border="1"> <tr> <td>State Meaning</td> <td>Asserted—Indicates that the processor must initiate a system reset interrupt. Negated—Indicates that the interrupt is not being requested.</td> </tr> </table> | State Meaning | Asserted—Indicates that the processor must initiate a system reset interrupt. Negated—Indicates that the interrupt is not being requested. |
| | | State Meaning | Asserted—Indicates that the processor must initiate a system reset interrupt. Negated—Indicates that the interrupt is not being requested. | |
| | | <table border="1"> <tr> <td>Timing</td> <td>Assertion—Occurs at any time, asynchronously to any clock. Negation—Occurs after being serviced. (PCI host mode) or PCI_CLK (PCI agent mode)</td> </tr> </table> | Timing | Assertion—Occurs at any time, asynchronously to any clock. Negation—Occurs after being serviced. (PCI host mode) or PCI_CLK (PCI agent mode) |
| | | Timing | Assertion—Occurs at any time, asynchronously to any clock. Negation—Occurs after being serviced. (PCI host mode) or PCI_CLK (PCI agent mode) | |
| <table border="1"> <tr> <td>Requirements</td> <td>An open-drain signal. An external pull-up is required.</td> </tr> </table> | Requirements | An open-drain signal. An external pull-up is required. | | |
| Requirements | An open-drain signal. An external pull-up is required. | | | |
| <table border="1"> <tr> <td>Reset State</td> <td>Always input.</td> </tr> </table> | Reset State | Always input. | | |
| Reset State | Always input. | | | |
| CFG_RESET_SOURCE[0:3] | I | Reset configuration word source selection. These signals are on device pins that have other functions when the device is not in reset. They are sampled during the assertion of $\overline{\text{PORESET}}$ to determine the interface from which the device loads the reset configuration words. | | |
| | | <table border="1"> <tr> <td>State Meaning</td> <td>See Section 4.3.1.1, “Reset Configuration Word Source.”</td> </tr> </table> | State Meaning | See Section 4.3.1.1, “Reset Configuration Word Source.” |
| | | State Meaning | See Section 4.3.1.1, “Reset Configuration Word Source.” | |
| | | <table border="1"> <tr> <td>Timing</td> <td>These signals are sampled during the assertion of $\overline{\text{PORESET}}$ after a stable clock is supplied ($\overline{\text{PORESET}}$ flow) and must be pulled high or low by external resistors as long as HRESET is asserted.</td> </tr> </table> | Timing | These signals are sampled during the assertion of $\overline{\text{PORESET}}$ after a stable clock is supplied ($\overline{\text{PORESET}}$ flow) and must be pulled high or low by external resistors as long as HRESET is asserted. |
| | | Timing | These signals are sampled during the assertion of $\overline{\text{PORESET}}$ after a stable clock is supplied ($\overline{\text{PORESET}}$ flow) and must be pulled high or low by external resistors as long as HRESET is asserted. | |
| <table border="1"> <tr> <td>Requirements</td> <td>During $\overline{\text{PORESET}}$ and $\overline{\text{HRESET}}$ flows, all other signal drivers connected to these signals must be in the high-impedance state. Refer to the hardware specifications for proper resistor values to pull reset configuration signals high or low.</td> </tr> </table> | Requirements | During $\overline{\text{PORESET}}$ and $\overline{\text{HRESET}}$ flows, all other signal drivers connected to these signals must be in the high-impedance state. Refer to the hardware specifications for proper resistor values to pull reset configuration signals high or low. | | |
| Requirements | During $\overline{\text{PORESET}}$ and $\overline{\text{HRESET}}$ flows, all other signal drivers connected to these signals must be in the high-impedance state. Refer to the hardware specifications for proper resistor values to pull reset configuration signals high or low. | | | |
| <table border="1"> <tr> <td>Reset State</td> <td>Input during power-on and hard reset flows. Functional signal after reset flow completes.</td> </tr> </table> | Reset State | Input during power-on and hard reset flows. Functional signal after reset flow completes. | | |
| Reset State | Input during power-on and hard reset flows. Functional signal after reset flow completes. | | | |
| CFG_CLKIN_DIV | I | Clock in division selection. This signal is located on a device pin that has another function when the device is not in reset state. This signal is sampled during the assertion of $\overline{\text{PORESET}}$ to determine whether SYS_CLK_IN is divided by two. | | |
| | | <table border="1"> <tr> <td>State Meaning</td> <td>See Section 4.3.1.2, “SYS_CLK_IN Division.”</td> </tr> </table> | State Meaning | See Section 4.3.1.2, “SYS_CLK_IN Division.” |
| | | State Meaning | See Section 4.3.1.2, “SYS_CLK_IN Division.” | |
| | | <table border="1"> <tr> <td>Timing</td> <td>This signal is sampled during the assertion of $\overline{\text{PORESET}}$ after a stable clock is supplied ($\overline{\text{PORESET}}$ flow), and it must be pulled high or low by external resistors as long as HRESET is asserted.</td> </tr> </table> | Timing | This signal is sampled during the assertion of $\overline{\text{PORESET}}$ after a stable clock is supplied ($\overline{\text{PORESET}}$ flow), and it must be pulled high or low by external resistors as long as HRESET is asserted. |
| | | Timing | This signal is sampled during the assertion of $\overline{\text{PORESET}}$ after a stable clock is supplied ($\overline{\text{PORESET}}$ flow), and it must be pulled high or low by external resistors as long as HRESET is asserted. | |
| <table border="1"> <tr> <td>Requirements</td> <td>During $\overline{\text{PORESET}}$ and $\overline{\text{HRESET}}$ flows, all other signal drivers connected to this signal must be in the high-impedance state. Refer to the hardware specifications for proper resistors values to pull reset configuration signals high or low.</td> </tr> </table> | Requirements | During $\overline{\text{PORESET}}$ and $\overline{\text{HRESET}}$ flows, all other signal drivers connected to this signal must be in the high-impedance state. Refer to the hardware specifications for proper resistors values to pull reset configuration signals high or low. | | |
| Requirements | During $\overline{\text{PORESET}}$ and $\overline{\text{HRESET}}$ flows, all other signal drivers connected to this signal must be in the high-impedance state. Refer to the hardware specifications for proper resistors values to pull reset configuration signals high or low. | | | |
| <table border="1"> <tr> <td>Reset State</td> <td>Always input</td> </tr> </table> | Reset State | Always input | | |
| Reset State | Always input | | | |

4.1.2 Clock Signals

In Table 4-2, some clock signals are specific to blocks within the device. Although some of their functionality is described in Section 4.4, “Clocking,” they are defined in detail in their respective chapters. See Figure 4-7 for the internal distribution of clocks in the device.

Table 4-2. External Clock Signals

| Signal | I/O | Description | |
|----------------|-----|--|---|
| SYS_CLK_IN | I | System clock. In PCI host mode, SYS_CLK_IN is the primary input clock. SYS_CLK_IN directly feeds the PCI output clock dividers and is driven out on the PCI_SYNC_OUT signal for de-skewing external PCI clocks routing. If the device is used as PCI agent, PCI_CLK can be provided from a PCI source. In this case, SYS_CLK_IN may no longer be needed and should be tied low. If SYS_CLK_IN is the clock source, SYS_CR_CLK_IN should be tied low and SYS_CR_CLK_OUT should be left unconnected. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information. |
| | | Requirements | Should be tied low in PCI agent mode. |
| | | Reset State | Always input. |
| SYS_CR_CLK_IN | I | Crystal input. SYS_CR_CLK_IN/SYS_CR_CLK_OUT allows the system clock to be provided using an external crystal oscillator. If a crystal source is used, SYS_CLK_IN should be tied low. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information. |
| | | Requirements | Should be tied low if unused, for example in PCI agent mode or when the clock is provided through SYS_CLK_IN. |
| | | Reset State | Always input. |
| SYS_CR_CLK_OUT | O | Crystal output. SYS_CR_CLK_IN/SYS_CR_CLK_OUT allows the system clock to be provided through an external crystal oscillator. If a crystal source is used, SYS_CLK_IN should be tied low. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information. |
| | | Requirements | Should be left unconnected if unused, for example in PCI agent mode or when the clock is provided through SYS_CLK_IN. |
| | | Reset State | Always output. |
| USB_CLK_IN | I | USB PHY clock. USB_CLK_IN feeds into the USB PHY PLL. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information. |
| | | Requirements | Should be tied low if unused, for example when the clock is provided through USB_CR_CLK_IN or when derived from the system clock. |
| | | Reset State | Always input. |
| USB_CR_CLK_IN | I | USB crystal input. USB_CR_CLK_IN/USB_CR_CLK_OUT allows the USB clock to be provided using an external crystal oscillator. If a crystal source is used, USB_CLK_IN should be tied low. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information. |
| | | Requirements | Should be tied low if unused, for example when the clock is provided through USB_CR_CLK_IN or when derived from the system clock. |
| | | Reset State | Always input. |

Table 4-2. External Clock Signals (continued)

| Signal | I/O | Description | |
|-------------------------|-----|--|---|
| USB_CR_CLK_OUT | O | USB crystal output. USB_CR_CLK_IN/USB_CR_CLK_OUT allows the USB clock to be provided through an external crystal oscillator. If a crystal source is used, USB_CLK_IN should be tied low. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information |
| | | Requirements | Should be left unconnected if unused, for example when the clock is provided through USB_CR_CLK_IN or when derived from the system clock. |
| | | Reset State | Always output. |
| PCI_CLK/ PCI_SYNC_IN | I | PCI clock/ PCI synchronization clock (PCI_CLK/PCI_SYNC_IN). PCI_CLK is the primary clock input to the device, the reference clock for the system APLL. When the device is in PCI host mode SYS_CLK_IN or SYS_XTAL_IN is used as the clock source. In this case the PCI_CLK_OUT[0:2] signals are driven and PCI_SYNC_IN should be tied to PCI_SYNC_OUT. When the device is in PCI agent mode PCI_CLK will be tied directly to a PCI system clock source. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information |
| | | Reset State | Always input. |
| PCI_SYNC_OUT | O | Reference PCI output synchronization clock (PCI_SYNC_OUT). In PCI host mode with the PCI_CLK_OUT[0:2] signals driven, PCI_SYNC_OUT is connected externally to PCI_SYNC_IN signal for de-skewing external PCI clocks routing. PCI_SYNC_OUT has the same frequency as CLKIN or CLKIN/2 depending on the state of CFG_CLKIN_DIV at reset. See Section 4.3.1.2, “SYS_CLK_IN Division.” In PCI agent mode, this signal is typically not used. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information. |
| | | Reset State | Always output, toggling in PCI host mode. |
| PCI_CLK_OUT[0:2] | O | PCI output clocks bank. In PCI host mode, the device provides three separate clock output signals for feeding PCI agent devices. | |
| | | Timing | Assertion/Negation—See the hardware specifications for timing information. |
| | | Reset State | Always output. Drive ‘0’ and after power-on reset flow. Enabled by a memory-mapped register. |

4.2 Functional Description

This section describes the various ways to reset the device, the power-on reset configurations, and clocking.

4.2.1 Reset Operations

The device has several inputs to the reset logic:

- Power-on reset ($\overline{\text{PORESET}}$)
- External hard reset ($\overline{\text{HRESET}}$)
- External soft reset (CORE_SRESET_IN)
- Software watchdog reset
- System bus monitor reset
- Checkstop reset

- JTAG reset
- Software hard reset

All of these reset sources are fed into the reset controller and, depending on the source of the reset, different actions are taken. The reset status register, described in [Section 4.5.1.3, “Reset Status Register \(RSR\),”](#) indicates the last sources to cause a reset.

4.2.1.1 Reset Causes

[Table 4-3](#) describes reset causes.

Table 4-3. Reset Causes

| Name | Description |
|---|---|
| Power-on reset (PORESET) | Input signal. Asserting this signal initiates the power-on reset flow that resets the entire device and configures various attributes of the device including its clock modes. |
| Hard reset ($\overline{\text{HRESET}}$) | A bidirectional I/O signal. The device can detect an external assertion of $\overline{\text{HRESET}}$ only while it is not asserting hard reset. $\overline{\text{HRESET}}$ is an open-drain signal. |
| Soft reset (CORE_SRESET_IN) | Input signal. Connected to the \overline{sreset} input of the e300 core, and when asserted, causes a high priority interrupt to the e300 core. |
| Software watchdog reset | After the device watchdog counts to zero, a software watchdog reset is signaled. The enabled software watchdog event then generates an internal hard reset sequence. |
| System bus monitor reset | After the device CSB bus monitor reaches a timeout condition, a bus monitor reset is asserted. The enabled bus monitor event then generates an internal hard reset sequence. |
| Checkstop reset | If the core enters checkstop state and the checkstop reset is enabled ($\text{RMR}[\text{CSRE}] = 1$), the checkstop reset is asserted. The enabled checkstop event then generates an internal hard reset sequence. |
| JTAG reset | When JTAG logic asserts the JTAG soft reset signal, an internal soft reset sequence is generated. |
| Software hard reset | A hard reset sequence can be initialized by writing to a memory-mapped register (RCR). |

4.2.1.2 Reset Actions

The reset control logic determines the cause of reset, synchronizes it if necessary, and resets the appropriate internal hardware. The CORE_SRESET_IN input is not routed to the reset control logic, but directly to the \overline{sreset} input on the e300 core. Each reset flow has a different impact on the device logic:

- Power-on reset has the greatest impact, resetting the entire device, including clock logic and error capture registers.
- Hard reset resets the entire device excluding clock logic and error capture registers.
- Soft reset initializes the internal logic while maintaining the system configuration.

The memory controller, system protection logic, interrupt controller, and I/O signals are initialized only on hard reset. A soft reset causes a reset exception to the e300 core but does not reset other device logic. [Table 4-4](#) identifies the reset actions for each reset source.

Table 4-4. Reset Actions

| Action | Reset Source | | |
|---|----------------|---|------------|
| | Power-On Reset | External Hard Reset Software Watchdog Bus Monitor Checkstop Software Hard Reset | JTAG Reset |
| Resets: PLLs, clocks, RTC unit, and error capture registers | Yes | No | No |
| Resets: DDR, LBC, I/O multiplexers, GTM, PIT, GPIO, system configuration, and local access windows | Yes | Yes | No |
| Resets other internal logic | Yes | Yes | Yes |
| Reset configuration words loaded | Yes | Yes | No |
| $\overline{\text{HRESET}}$ driven | Yes | Yes | No |
| Hard reset to e300 core | Yes | Yes | No |
| High priority interrupt to the e300 core | No | No | No |

4.2.2 Power-On Reset Flow

Assertion of the $\overline{\text{PORESET}}$ external signal initiates the power-on reset flow. $\overline{\text{PORESET}}$ should be asserted externally for at least 32 input clock cycles after stable external power to the device is applied.

Directly after the negation of $\overline{\text{PORESET}}$, the device starts the configuration process. The device asserts $\overline{\text{HRESET}}$ throughout the power-on reset process, including configuration. Configuration time varies according to the configuration source and SYS_CLK_IN (PCI host mode) or PCI_CLK (PCI agent mode) frequency. Initially, the reset configuration inputs are sampled to determine the configuration source and the input clock division mode. Next, the device starts loading the reset configuration words. The system PLL begins to lock according to the clock mode values in the reset configuration word low. When the system PLL is locked, the clock unit starts distributing clock signals in the device. At this stage, the core PLL begins to lock. When it is locked and the reset configuration words are loaded, $\overline{\text{HRESET}}$ is released.

The detailed power-on reset (POR) flow for the device is as follows:

1. Power is applied to meet the specifications in the device data sheet.
2. The system asserts $\overline{\text{PORESET}}$ and $\overline{\text{TRST}}$, causing all registers to be initialized to their default states and most I/O drivers to be released to high-impedance.
Some clock, clock enabled, and system control signals remain active.
3. The system applies a stable SYS_CLK_IN (PCI host mode) or PCI_CLK (PCI agent mode) signal and stable reset configuration inputs (CFG_RESET_SOURCE , $\overline{\text{CFG_CLKIN_DIV}}$).
4. The system negates $\overline{\text{PORESET}}$ after at least 32 stable SYS_CLK_IN (PCI host mode) or PCI_CLK (PCI agent mode) clock cycles.

5. The device samples the reset configuration input signals to determine the clock division and the reset configuration words source.
6. The device starts loading the reset configuration words.
Loading time depends on the reset configuration word source.
7. When the reset configuration word low is loaded, the system PLL begins to lock.
When the system PLL is locked, *csb_clk* is supplied to the core PLL.
8. The core PLL begins to lock.
9. The device drives $\overline{\text{HRESET}}$ asserted until the e300 PLL is locked and the reset configuration words are loaded.
10. The user optionally negates $\overline{\text{HRESET}}$ if it was not negated earlier.
JTAG logic must always be initialized by asserting $\overline{\text{TRST}}$. If the JTAG signals are not used, $\overline{\text{TRST}}$ should be connected directly to $\overline{\text{PORESET}}$. $\overline{\text{TRST}}$ must not remain asserted after the negation of $\overline{\text{PORESET}}$.
11. The internal reset to the core and the rest of the logic is negated. I/O drivers are enabled. The PCI interface can assert $\overline{\text{DEVSEL}}$ in response to configuration cycles.
12. The device stops driving $\overline{\text{HRESET}}$. The reset to the e300 core is negated and the core is enabled. The boot sequencer, if enabled, is released, causing it to load configuration data from serial ROMs, as described in [Section 20.4.5, “Boot Sequencer Mode.”](#)
13. Before the boot sequencer finishes, it can enable the PCI interface to accept external requests, if required, by clearing the `CFG_LOCK` bit in the PCI function configuration register as described in [Table 13-40](#). If the e300 core is required to proceed, the boot sequencer should enable boot vector fetch by clearing `ACR[COREDIS]` as described in [Section 6.2.1, “Arbiter Configuration Register \(ACR\).”](#)
14. The PCI interface can now accept external requests, if enabled, and the boot vector fetch by the core can proceed, if enabled.

The device is now in its ready state.

Figure 4-1 shows a timing diagram of the power-on reset flow.




Figure 4-1. Power-On Reset Flow

4.2.3 Hard Reset Flow

The $\overline{\text{HRESET}}$ signal is initiated externally by asserting $\overline{\text{HRESET}}$ or internally when the device detects a reason to generate an internal hard reset sequence. In both cases, the device continues asserting $\overline{\text{HRESET}}$ throughout the $\overline{\text{HRESET}}$ state. The hard reset sequence time varies according to the configuration source and SYS_CLK_IN (PCI host mode) or PCI_CLK (PCI agent mode) frequency. The reset configuration input signals (CFG_RESET_SOURCE and CFG_CLKIN_DIV) are not sampled by hard reset (only by power-on reset), so the device immediately starts loading the reset configuration words and configures the device as explained in Section 4.3.3, “Loading the Reset Configuration Words.” After the configuration sequence completes, the device releases the $\overline{\text{HRESET}}$ signal and exits the $\overline{\text{HRESET}}$ state. An external pull-up resistor should negate the signals. After negation is detected, a 16-cycle period is taken before testing for the presence of an external (hard) reset.

NOTE

Because the device does not sample the reset configuration input signals (CFG_RESET_SOURCE , CFG_CLKIN_DIV) during a hard reset flow, setting a new value on those signals (other than that set during power-on reset) has no effect.

Figure 4-2 shows a timing diagram of the hard reset flow.




Figure 4-2. Hard Reset Flow

4.2.4 Soft Reset Flow

The CORE_SRESET_IN signal can be initiated externally by asserting CORE_SRESET_IN. CORE_SRESET_IN assertion asserts the *sreset* input to the e300 core. No other hardware is reset.

4.3 Reset Configuration

The device is initialized using two complementary methods, latching CFG_RESET_SOURCE and loading the reset configuration words. Initially, a few input signals are sampled during the assertion of the PORESET signal. These signals determine whether a reset configuration word is required and the device source interface from which it is loaded. According to the value on these signals, the device continues loading the reset configuration word.

4.3.1 Reset Configuration Signals

Reset configuration input signals are on device pins that have other functions when the device is not in reset state. These input signals are sampled into registers during the assertion of PORESET, after a stable clock is supplied (PCI_CLK), and must be pulled high or low by external resistors as long as HRESET is asserted. While the PORESET and HRESET signals are asserted, all other signal drivers connected to these signals must be in the high-impedance state. Refer to the hardware specifications for proper resistor values for pulling reset configuration signals high or low.

This section describes the modes configured by the reset configuration signals. Note that the reset configuration inputs sampled values are accessible to software through memory-mapped registers described in [Section 4.5.1.3, “Reset Status Register \(RSR\),”](#) and [Section 4.5.2.1, “System PLL Mode Register \(SPMR\).”](#)

NOTE

Implement one of the following methods to control the selection between the reset and non-reset function of these pins.

- Resistors. Use pullup or pulldown resistors to set the desired value on the reset configuration input signals. During the power-on and hard reset sequences, these signals are inputs to the device.
- Active driving device. Use $\overline{\text{HRESET}}$ to control the driving device. When $\overline{\text{HRESET}}$ is asserted, drive reset configuration values on the pins; when $\overline{\text{HRESET}}$ is negated, stop driving the reset configuration input signals.

4.3.1.1 Reset Configuration Word Source

The reset configuration word source options, shown in [Table 4-5](#), select whether the device loads a reset configuration word from NOR Flash, NAND Flash, or an I²C EEPROM or uses hard-coded default options. The value of these signals also affects the duration of power-on and hard reset sequences. In any case, the reset sequence does not exceed 1 ms.

Table 4-5. Reset Configuration Words Source

| CFG_RESET_SOURCE[0:3] | Meaning |
|-----------------------|---|
| 0000 | Reset configuration word is loaded from NOR Flash |
| 0001 | Reset configuration word is loaded from NAND Flash memory (8-bit small page). |
| 0010 | Reserved |
| 0011 | Reserved |
| 0100 | Reset configuration word is loaded from an I ² C EEPROM. PCI_CLK/PCI_SYNC_IN is valid for any PCI frequency up to 66.666 MHz (range of 24–66.666 MHz). |
| 0101 | Reset configuration word is loaded from NAND Flash memory (8-bit large page). |
| 0110 | Reserved |
| 0111 | Reserved |
| 1000 | Hard-coded option 0. Reset configuration word is not loaded. |
| 1001 | Hard-coded option 1. Reset configuration word is not loaded. |
| 1010 | Hard-coded option 2. Reset configuration word is not loaded. |
| 1011 | Hard-coded option 3. Reset configuration word is not loaded. |
| 1100 | Hard-coded option 4. Reset configuration word is not loaded. |
| 1101 | Reserved |
| 1110 | Reserved |
| 1111 | Reserved |

4.3.1.2 SYS_CLK_IN Division

When the device is configured as a PCI host, the $\overline{\text{CFG_CLKIN_DIV}}$ configuration input selects the relationship between SYS_CLK_IN and PCI_SYNC_OUT as shown in Table 4-6. As a PCI host, the device supports three output signals. The frequency of the output clocks will be equal to the PCI_SYNC_OUT frequency.

When the device is configured as a PCI agent, the $\overline{\text{CFG_CLKIN_DIV}}$ configuration input can be used to double the internal clock frequencies, if sampled as '0' during power-on reset assertion. This feature is useful if a fixed internal frequency is desired regardless of whether the PCI clock is running at 33 or 66 MHz. PCI specifications require the PCI clock frequency information to be provided by the M66EN signal.

When the device is configured as PCI host, there are two scenarios for connecting the $\overline{\text{CFG_CLKIN_DIV}}$ configuration input. If the frequency of SYS_CLK_IN is 33 MHz (that is, the PCI system is running on a 33-MHz clock), $\overline{\text{CFG_CLKIN_DIV}}$ should be connected high. If the frequency of SYS_CLK_IN is 66 MHz (that is, the PCI system can run at 33- or 66-MHz clock signaled by M66EN), $\overline{\text{CFG_CLKIN_DIV}}$ should be connected to the M66EN signal.

Table 4-6. SYS_CLK_IN Division

| CFG_CLKIN_DIV | Description |
|---------------|--|
| 1 | In PCI host mode, SYS_CLK_IN: PCI_SYNC_OUT = 1:1 and all PCI_CLK_OUT[0:2] clocks are running at a frequency which is equal to the SYS_CLK_IN frequency. |
| 0 | In PCI agent mode, SYS_CLK_IN: PCI_SYNC_OUT = 2:1 and all PCI_CLK_OUT[0:2] clocks are running at a frequency which is equal to the PCI_SYNC_OUT frequency. |

4.3.1.3 Selecting Reset Configuration Input Signals

The example described in Table 4-7 shows how the user should pull down or pull up the reset configuration input signals (CFG_RESET_SOURCE, $\overline{\text{CFG_CLKIN_DIV}}$). The reset sequence duration is measured from the negation of $\overline{\text{PORESET}}$ to the negation of $\overline{\text{HRESET}}$. Note that the duration mentioned in this table is typical and does not represent cases in which the process of loading the reset configuration word had to be retried due to errors.

Table 4-7. Selecting Reset Configuration Input Signals

| I ² C EEPROM Configuration Words | SYS_CLK_IN Frequency (Host Mode) | $\overline{\text{CFG_CLKIN_DIV}}$ (Host Mode) | PCI_CLK Frequency (Agent Mode) | CFG_RESET_SOURCE[0:3] | Reset Sequence Duration in SYS_CLK_IN/ PCI_CLK Cycles | Duration |
|---|----------------------------------|---|--------------------------------|----------------------------------|---|-------------|
| No | 33 MHz | 1 | 33 MHz | 0000 (RCW loaded from NOR Flash) | 15210 | 456 μ s |
| No | 66 MHz | 1 | 66 MHz | 1000–1100 (use hard coded RCW) | 15380 | 231 μ s |

Table 4-7. Selecting Reset Configuration Input Signals (continued)

| I ² C EEPROM Configuration Words | SYS_CLK_IN Frequency (Host Mode) | CFG_CLKIN_DIV (Host Mode) | PCI_CLK Frequency (Agent Mode) | CFG_RESET_SOURCE[0:3] | Reset Sequence Duration in SYS_CLK_IN/PCI_CLK Cycles | Duration |
|---|----------------------------------|---------------------------|--------------------------------|--|--|----------|
| No | 66 MHz | 0 | 33 MHz | 0000 (RCW loaded from NOR Flash) | 30420/15210 | 456 μs |
| Yes | 33 MHz | 1 | 33 MHz | 0100 (I ² C EEPROM) | 106534 | 196 μs |
| Yes | 66 MHz | 1 | 66 MHz | 0100 (I ² C EEPROM) | 106534 | 1598 μs |
| Yes | 66 MHz | 0 | 33 MHz | 0100 (I ² C EEPROM) | 213068/106534 | 3196 μs |
| No | 66 MHz | 1 | 66 MHz | 0001 (RCW loaded from 8-bit small page NAND Flash) | 23024 | 345 μs |
| No | 66 MHz | 1 | 66 MHz | 0101 (RCW loaded from 8-bit large page NAND Flash) | 45284 | 679 μs |

4.3.2 Reset Configuration Words

The reset configuration words control the clock ratios and other basic device functions such as PCI host or agent mode, boot location, eTSEC modes, and endian mode. The reset configuration words are loaded from NOR Flash, NAND Flash, or the I²C interfaces or from hard-coded values during the power-on or hard reset flows. See [Section 4.3.1, “Reset Configuration Signals,”](#) for information on the reset configuration word source. Although the configuration reset words are loaded during hard reset flows, the clocks and PLL modes are reset only when PORESET is asserted during a power-on reset flow. See [Section 4.2.1.2, “Reset Actions.”](#) The values of fields in the reset configuration words registers (RCWLR and RCWHR) reflect only their state during the reset flow. Some of these parameters and modes can be modified by changing their values in the memory-mapped registers of other units, which does not affect RCWLR and RCWHR.

The reset configuration settings are accessible to software through the following read-only memory-mapped registers:

- Reset configuration word low register (RCWLR)
- Reset configuration word high register (RCWHR)
- Reset status register (RSR)
- System PLL mode register (SPMR)

See [Section 4.5, “Memory Map/Register Definitions.”](#)

4.3.2.1 Reset Configuration Word Low Register (RCWLR)

RCWLR is shown in [Figure 4-3](#). This read-only register gets its values according to the reset configuration word low loaded during the reset flow.




Figure 4-3. Reset Configuration Word Low Register (RCWLR)

[Table 4-8](#) defines the RCWLR bit fields.

Table 4-8. RCWLR Bit Settings

| Bits | Name | Description |
|-------|---------|--|
| 0 | LBCM | Local bus memory controller clock mode. Selects the local bus controller clock ratio. The local bus memory controller operates with a frequency equal to the frequency of <i>csb_clk</i> . This bit should be cleared. |
| 1 | DDRCM | DDR SDRAM memory controller clock mode. Selects the DDR SDRAM memory controller clock ratio. The DDR SDRAM memory controller operates at twice the frequency of the <i>csb_clk</i> . 0 <i>csb_clk</i> ratio is 1:1 1 <i>csb_clk</i> ratio is 2:1 |
| 2–3 | SVCOD | System PLL VCO division. See Section 4.3.2.1.1, “System PLL VCO Division.” |
| 4–7 | SPMF | System PLL multiplication factor. |
| 8 | — | Reserved, should be cleared. |
| 9–15 | COREPLL | Core PLL configuration. COREPLL sets the ratio between the e300 core clock and the internal <i>csb_clk</i> of the device. The encodings for COREPLL are given in the hardware specifications for this device. |
| 16–31 | — | Reserved, should be cleared. |

4.3.2.1.1 System PLL VCO Division

The RCWLR field SVCOD (system PLL VCO division), shown in [Table 4-9](#), establishes the internal ratio between the system PLL VCO frequency and the PLL output clock frequency. The PLL output clock frequency equals *csb_clk* frequency if RCWLR[LBCM] and RCWLR[DDRCM] are both cleared or twice the *csb_clk* frequency if RCWLR[LBCM] or RCWLR[DDRCM] or both of them are set.

Table 4-9 describes the setting of SVCOD bits.

Table 4-9. System PLL VCO Division

| Reset Configuration Word Low Register (RCWLR) Bits | Field Name | Value (Binary) | VCO Division Factor |
|--|------------|----------------|---------------------|
| 2–3 | SVCOD | 00 | 2 |
| | | 01 | 4 |
| | | 10 | 8 |
| | | 11 | 1 |

4.3.2.1.2 System PLL Configuration

The system PLL ratio reset, shown in Table 4-10, establishes the clock ratio between the SYS_CLK_IN signal and the internal *csb_clk* of the device. *csb_clk* drives internal units and feeds the e300 core PLL.

Table 4-10. System PLL Ratio

| RCWLR Bits | Field Name | Value (Binary) | <i>csb_clk</i> : CLKIN (PCI Host Mode) <i>csb_clk</i> : $(\text{PCI_CLK} \times (1 + \text{sampled_cfg_clk_div}))$ (PCI Agent Mode) |
|------------|------------|----------------|--|
| 4–7 | SPMF | 0000 | Reserved |
| | | 0001 | Reserved |
| | | 0010 | 2 : 1 |
| | | 0011 | 3 : 1 |
| | | 0100 | 4 : 1 |
| | | 0101 | 5 : 1 |
| | | 0110 | 6 : 1 |
| | | 0111–1111 | Reserved, should not be set |

NOTE

In PCI host mode, the SPMF field described in Table 4-10 always selects the *csb_clk*:SYS_CLK_IN ratio regardless of the $\overline{\text{CFG_CLKIN_DIV}}$ reset configuration input value during reset flow.

The SPMF field maximum allowed value is dependent on the value sampled on $\overline{\text{CFG_CLKIN_DIV}}$ during power-on resetIU. Table 4-11 defines the upper limit of SPMF with respect to these values. Values for SPMF are as follows:

Table 4-11. SPMF Maximum Values

| SYS_CLK_IN (MHz) | $\overline{\text{CFG_CLKIN_DIV}}$ | LBCM | DDRCM | Maximum SPMF Value (Decimal) |
|------------------|-------------------------------------|------|-------|------------------------------|
| 25 | 1 | 0 | 1 | 5 |
| 33 | 1 | 0 | 1 | 4 |

Table 4-11. SPMF Maximum Values (continued)

| SYS_CLK_IN (MHz) | CFG_CLKIN_DIV | LBCM | DDRCM | Maximum SPMF Value (Decimal) |
|------------------|---------------|------|-------|------------------------------|
| 50 | 0 | 0 | 1 | 2 |
| 66 | 0 | 0 | 1 | 2 |

4.3.2.2 Reset Configuration Word High Register (RCWHR)

RCWHR is shown in [Figure 4-4](#). This read-only register gets its values according to the reset configuration word high loaded during the reset flow.

Offset 0x0_0904

Access: Read/Write

| | | | | | | | | | | | | | | | |
|-------|---------|---|--------|----|---------|-----|---------|------|--------|----|-------|------|----|----|----|
| Field | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 11 | 12 | 13 | 14 | 15 |
| | PCIHOST | — | PCIARB | — | COREDIS | BMS | BOOTSEQ | SWEN | ROMLOC | | RLEXT | | — | — | |
| Field | 16 | | 18 | 19 | 21 | | 22 | 23 | 24 | 27 | | 28 | 29 | 30 | 31 |
| | TSEC1M | | TSEC2M | | — | | — | | — | | TLE | LALE | — | | |

Figure 4-4. Reset Configuration Word High Register (RCWHR)

[Table 4-12](#) defines the reset configuration word high bit fields.

Table 4-12. Reset Configuration Word High Bit Settings

| Bits | Name | Description | | | | | | | | |
|------------------------------|---------------------------------|--|------------------------------|------------------------------|------------|---------------------------------|-------------|---------------------------------|--------------|---------------------------------|
| 0 | PCIHOST | PCI host mode. See Section 4.3.2.2.1, “PCI Host/Agent Configuration,” for more information. | | | | | | | | |
| 1 | — | Reserved, should be cleared. | | | | | | | | |
| 2 | PCIARB | PCI internal arbiter mode. Enables the on-chip PCI arbiter. 0 On-chip PCI arbiter is disabled. External arbitration is required. 1 On-chip PCI arbiter is enabled. The value of PCIARB also defines the function of the PCI arbitration signals that are multiplexed with CompactPCI signals, as follows: <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Pin Function When PCIARB = 0</th> <th>Pin Function When PCIARB = 1</th> </tr> </thead> <tbody> <tr> <td>CPCI_HS_ES</td> <td>$\overline{\text{PCI_REQ}}[1]$</td> </tr> <tr> <td>CPCI_HS_LED</td> <td>$\overline{\text{PCI_GNT}}[1]$</td> </tr> <tr> <td>CPCI_HS_ENUM</td> <td>$\overline{\text{PCI_GNT}}[2]$</td> </tr> </tbody> </table> | Pin Function When PCIARB = 0 | Pin Function When PCIARB = 1 | CPCI_HS_ES | $\overline{\text{PCI_REQ}}[1]$ | CPCI_HS_LED | $\overline{\text{PCI_GNT}}[1]$ | CPCI_HS_ENUM | $\overline{\text{PCI_GNT}}[2]$ |
| Pin Function When PCIARB = 0 | Pin Function When PCIARB = 1 | | | | | | | | | |
| CPCI_HS_ES | $\overline{\text{PCI_REQ}}[1]$ | | | | | | | | | |
| CPCI_HS_LED | $\overline{\text{PCI_GNT}}[1]$ | | | | | | | | | |
| CPCI_HS_ENUM | $\overline{\text{PCI_GNT}}[2]$ | | | | | | | | | |
| 3 | — | Reserved, should be cleared. | | | | | | | | |

Table 4-12. Reset Configuration Word High Bit Settings (continued)

| Bits | Name | Description |
|-------|----------|--|
| 4 | COREDISE | <p>Core disable mode. Specifies the e300 core mode out of reset. If COREDIS is set, the core cannot fetch boot code until it is configured by an external master. The external master frees the core to boot by clearing the COREDIS bit in the arbiter configuration register as described in Section 6.2.1, “Arbiter Configuration Register (ACR).”</p> <p>This bit must be set when the boot sequencer is enabled to initiate the device (BOOTSEQ is not 0b00). Otherwise, unpredictable behavior occurs.</p> <p>0 The core can boot without waiting for configuration by an external master. 1 Core boot holdoff mode. The core is prevented from booting until it is configured by an external master.</p> |
| 5 | BMS | <p>Boot memory space. See Section 4.3.2.2.2, “Boot Memory Space (BMS),” for more information.</p> |
| 6–7 | BOOTSEQ | <p>Boot sequencer configuration. See Section 4.3.2.2.3, “Boot Sequencer Configuration,” for more information.</p> |
| 8 | SWEN | <p>Software watchdog enable. Selects whether the software watchdog is enabled to start counting down immediately when coming out of reset. The user can override this value by writing to the system watchdog control register (SWCRR[SWEN]) during system initialization.</p> <p>0 Disabled 1 Enabled</p> |
| 9–11 | ROMLOC | <p>Boot ROM interface location. This bit combined with bit RLEXT determines where the device boots from. See Section 4.3.2.2.4, “Boot ROM Location,” for more information.</p> |
| 12–13 | RLEXT | <p>Boot ROM location extension. This bit combined with bit ROMLOC determines where the device boots from. See Section 4.3.2.2.4, “Boot ROM Location,” for more information.</p> <p>00 Legacy mode—allows for booting from on-chip peripherals. Refer to Table 4-16 for more information. 01 NAND Flash mode—allows for booting from NAND flash devices. Refer to Table 4-16 for more information. 10 Reserved 11 Reserved</p> |
| 14–15 | — | Reserved, should be cleared. |
| 16–18 | TSEC1M | <p>TSEC1 mode. See Section 4.3.2.2.5, “eTSEC1 Mode,” for more information.</p> |
| 19–21 | TSEC2M | <p>TSEC2 mode. See Section 4.3.2.2.6, “eTSEC2 Mode,” for more information.</p> |
| 22–27 | — | Reserved, should be cleared. |
| 28 | TLE | True little-endian. See Section 4.3.2.2.7, “e300 Core True Little-Endian,” for more information. |
| 29 | LALE | Reserved, should be cleared. |
| 30–31 | — | Reserved, should be cleared. |

4.3.2.2.1 PCI Host/Agent Configuration

The PCIHOST configuration parameter, shown in [Table 4-13](#), configures the device to act as a PCI host or as a PCI agent device. In host mode, the device can immediately master transactions to the PCI interface. If the device is a PCI agent device, the device is disabled from mastering PCI transactions until the external

host enables it to do so. The external host does this by setting the control registers of the device's interfaces appropriately. See details in the PCI programming model described in [Section 13.3, "PCI Memory Map/Register Definitions."](#)

Table 4-13. PCI Host/Agent Configuration

| RCWHR Bit | Field Name | Value (Binary) | Meaning |
|-----------|------------|----------------|--|
| 0 | PCIHOST | 0 | The device acts as a PCI agent device. |
| | | 1 | The device acts as the host processor (default). |

NOTE

If the device is a PCI agent, and the e300 core is not in holdoff mode (as described in [Section 4.3.2.2, "Reset Configuration Word High Register \(RCWHR\)"](#)), the boot ROM should not be located on the PCI interface because the device is not enabled to master reads onto the PCI bus.

4.3.2.2.2 Boot Memory Space (BMS)

BMS defines the initial value of the e300 core MSR[IP] bit, which specifies the location of the interrupt vectors (including the hard reset exception vector). The device defines the default boot ROM memory space to be 8 Mbytes at addresses All zeros to 0x007F_FFFF or 0xFF80_0000 to 0xFFFF_FFFF. When the core comes out of reset, if it is enabled to boot, it begins fetching boot code from one of two addresses: 0x0000_0100 or 0xFFFF0_0100, and exceptions are vectored to physical addresses 0x000n_nnnn or 0xFFFFn_nnnn appropriately. This bit specifies whether an interrupt vector offset is prepended with 0xFFF or 0x000. In the description below, *n_nnnn* is the offset of the exception vector.

The boot memory space reset configuration word field, shown in [Table 4-14](#), specifies both the device boot ROM address window and the initial e300 core boot address.

Table 4-14. Boot Memory Space

| RCWHR Bit | Field Name | Value (Binary) | Meaning |
|-----------|------------|----------------|--|
| 5 | BMS | 0 | Boot memory space is 8 Mbytes at All zeros to 0x007F_FFFF. e300 core register MSR[IP] initial value is 0b0. The core, if enabled to boot, begins fetching boot code from address 0x0000_0100 and exceptions are vectored to the physical address of 0x000n_nnnn. |
| | | 1 | Boot memory space is 8 Mbytes at 0xFF80_0000 to 0xFFFF_FFFF. e300 core register MSR[IP] initial value is 0b1. The core, if enabled to boot, begins fetching boot code from address 0xFFFF0_0100 and exceptions are vectored to the physical address of 0xFFFFn_nnnn. |

4.3.2.2.3 Boot Sequencer Configuration

The boot sequencer configuration options, shown in [Table 4-15](#), allow the boot sequencer to load configuration data from the serial ROM located on the I²C port before the host tries to configure the device.

These options also specify normal or extended I²C addressing modes. See [Section 20.4.5, “Boot Sequencer Mode.”](#)

Table 4-15. Boot Sequencer Configuration

| RCWHR Bits | Field Name | Value (Binary) | Meaning |
|------------|------------|----------------|---|
| 6–7 | BOOTSEQ | 00 | Boot sequencer is disabled. No I ² C ROM is accessed. |
| | | 01 | Normal I ² C addressing mode is used. Boot sequencer is enabled and loads configuration information from a ROM on the I ² C interface. A valid ROM must be present. |
| | | 10 | Extended I ² C addressing mode is used. Boot sequencer is enabled and loads configuration information from a ROM on the I ² C interface. A valid ROM must be present. |
| | | 11 | Reserved, should be cleared. |

NOTE

When the boot sequencer is enabled, the e300 core must be prevented from fetching boot code, by setting the core disable reset configuration word field (COREDIS) as described in [Section 4.3.2.2, “Reset Configuration Word High Register \(RCWHR\).”](#) If the e300 core is required to proceed, the boot sequencer should enable boot vector fetch by clearing ACR[COREDIS] as described in [Section 6.2.1, “Arbiter Configuration Register \(ACR\).”](#)

4.3.2.2.4 Boot ROM Location

The device defines the default boot ROM address range to be 8 Mbytes at addresses All zeros to 0x007F_FFFF or 0xFF80_0000 to 0xFFFF_FFFF (selected by the BMS reset configuration word field). However, the on-chip peripheral that manages these boot ROM accesses can be selected at power up.

The boot ROM location reset configuration word field, shown in [Table 4-16](#), establishes the location of boot ROM. The exact boot ROM location table to be used is defined by the setting of RCWHR[RLEXT]

bits, as shown in [Table 4-12](#). Accesses to the boot vector and the default boot ROM region of the local address map are directed to the interface specified by this field.

Table 4-16. Boot ROM Location

| RCWHR Bits | Field Name | Value (Binary) | Meaning | |
|------------|------------|----------------|------------------------------|---|
| | | | Legacy Mode (RLEXT = 00) | NAND Flash Mode (RLEXT = 01) |
| 9–11 | ROMLOC | 000 | DDR SDRAM | Reserved |
| | | 001 | PCI | Local bus NAND Flash—8-bit small page ROM |
| | | 010 | Reserved, should be cleared. | Reserved |
| | | 011 | Reserved, should be cleared. | Reserved |
| | | 100 | Reserved | Reserved |
| | | 101 | Local bus GPCM—8-bit ROM | Local bus NAND Flash—8-bit large page ROM |
| | | 110 | Local bus GPCM—16-bit ROM | Reserved |
| | | 111 | Reserved | Reserved |

The local access window of the selected boot ROM interface is enabled and initialized with the proper base address and size, as described in [Section 5.2, “Local Memory Map Overview and Example.”](#)

4.3.2.2.5 eTSEC1 Mode

The TSEC1 mode reset configuration word field, shown in [Table 4-17](#), selects the protocol used by the eTSEC1 controller (enhanced three-speed Ethernet controller interface).

Table 4-17. eTSEC1 Mode Configuration

| Reset Configuration Word High Register (RCWHR) Bits | Field Name | Value (Binary) | Meaning |
|---|------------|----------------|--|
| 16–18 | TSEC1M | 000 | The eTSEC1 controller operates in the MII protocol, using only four transmit data signals and four receive data signals. |
| | | 001 | The eTSEC1 controller operates in the RMII protocol, using only two transmit data signals and two receive data signals. |
| | | 010 | Reserved |
| | | 011 | The eTSEC1 controller operates in the RGMII protocol, using four transmit data signals and four receive data signals. |
| | | 100 | Reserved |
| | | 101 | The eTSEC1 controller operates in the RTBI protocol, using only four transmit data signals and four receive data signals. |
| | | 110 | The eTSEC1 controller operates in the SGMII protocol, using the on-chip PHY. For other modes, SerDes will be selected as the PCI Express mode. |
| | | 111 | Reserved |

NOTE

The reset value of the system I/O configuration register high (SICRH) depends on the reset configuration word high TSEC1M field setting. This is used to avoid contention in systems not using the RTBI modes. In non-TBI modes, device signals which have additional functions are set to be in a non-TSEC function, thus not driven during and after reset. The function of these signals can be changed by writing to this register during system initialization. See [Section 5.3.2.6, “System I/O Configuration Register High \(SICRH\).”](#)

4.3.2.2.6 eTSEC2 Mode

The eTSEC2 mode reset configuration word field, shown in [Table 4-18](#), selects the protocol used by the eTSEC2 controller (enhanced three-speed Ethernet controller interface).

Table 4-18. eTSEC2 Mode Configuration

| Reset Configuration Word High Register (RCWHR) Bits | Field Name | Value (Binary) | Meaning |
|---|------------|----------------|--|
| 19–21 | TSEC2M | 000 | The eTSEC2 controller operates in the MII protocol, using only four transmit data signals and four receive data signals. |
| | | 001 | The eTSEC2 controller operates in the RMII protocol, using only two transmit data signals and two receive data signals. |
| | | 010 | Reserved |
| | | 011 | The eTSEC2 controller operates in the RGMII protocol, using four transmit data signals and four receive data signals. |
| | | 100 | Reserved |
| | | 101 | The eTSEC2 controller operates in the RTBI protocol, using only four transmit data signals and four receive data signals. |
| | | 110 | The eTSEC2 controller operates in the SGMII protocol, using the on-chip PHY. For other modes, SerDes will be selected as the PCI Express mode. |
| | | 111 | Reserved |

NOTE

The reset value of the system I/O configuration register high (SICRH) depends on the setting of TSEC2M. This is used to avoid contention in systems not using TBI or RTBI modes. In non-TBI modes, device signals which have additional functions are set to be in a non-TSEC function, thus not driven during and after reset. The function of these signals can be changed by writing to this register during system initialization. See [Section 5.3.2.6, “System I/O Configuration Register High \(SICRH\).”](#)

4.3.2.2.7 e300 Core True Little-Endian

The true little endian reset configuration word field, shown in [Table 4-19](#), selects whether the e300 core operates in big-endian mode or true little-endian mode at reset.

Table 4-19. e300 Core True Little-Endian

| Reset Configuration Word High Register (RCWHR) Bit | Field Name | Value (Binary) | Meaning |
|--|------------|----------------|-------------------------|
| 28 | TLE | 0 | Big-endian mode |
| | | 1 | True little-endian mode |

4.3.2.2.8 LALE Configuration

The LALE reset configuration word field, shown in [Table 4-20](#), configures the timing of the local bus LALE signal. Refer to the hardware specifications for specific timing information.

Table 4-20. LALE Configuration

| Reset Configuration Word High Register (RCWHR) Bit | Field Name | Value (Binary) | Meaning |
|--|------------|----------------|--|
| 29 | LALE | 0 | Normal LALE timing |
| | | 1 | LALE is negated 1/2 a LBC_controller_clk earlier |

4.3.3 Loading the Reset Configuration Words

The device loads the reset configuration words from a local bus EEPROM, a local bus NAND Flash, or an I²C serial EEPROM, or uses hard-coded configuration, as selected by the reset configuration inputs described in [Section 4.3.1, “Reset Configuration Signals.”](#) The following sections describe each of these options.

4.3.3.1 Loading from Local Bus

The reset configuration words are assumed to reside in an EEPROM or NOR Flash or NAND Flash device connected to $\overline{\text{LCS0}}$ of the device local bus. Because the port size of this EEPROM is unknown, the device reads all configuration words byte-by-byte only from locations that are independent of port size. $+\text{LCS0}$ is the default for GPCM, so GPCM controlled is used to read the reset configuration word from EEPROM. $\overline{\text{LGTA}}$ should be high to avoid unintended early termination of the read cycle.

[Table 4-21](#) shows addresses that should be used to contain the reset configuration words. Byte addresses that do not appear in this table have no effect on the configuration of the device. The values of the bytes in [Table 4-21](#) are always read on byte lane LAD[0:7] regardless of the port size.

Table 4-21. Local Bus Configuration EEPROM Addresses

| Reset Configuration Word | Bits [0:7] Address | Bits [8:15] Address | Bits [16:23] Address | Bits [24:31] Address |
|--------------------------|--------------------|---------------------|----------------------|----------------------|
| Low | 0x00 | 0x08 | 0x10 | 0x18 |
| High | 0x20 | 0x28 | 0x30 | 0x38 |

[Table 4-22](#) shows the data structure of the local bus device containing the reset configuration words (RCWL and RCWH).

Table 4-22. Local Bus Reset Configuration Words Data Structure

| EEPROM Address | EEPROM Data Bits | | | |
|----------------|------------------|--------|---------|---------|
| | [0:7] | [8:15] | [16:23] | [24:31] |
| 0x00 | RCWL[0:7] | | | |
| 0x04 | | | | |
| 0x08 | RCWL[8:15] | | | |

Table 4-22. Local Bus Reset Configuration Words Data Structure (continued)

| EEPROM Address | EEPROM Data Bits | | | |
|----------------|------------------|--------|---------|---------|
| | [0:7] | [8:15] | [16:23] | [24:31] |
| 0x0C | | | | |
| 0x10 | RCWL[16:23] | | | |
| 0x14 | | | | |
| 0x18 | RCWL[24:31] | | | |
| 0x1C | | | | |
| 0x20 | RCWH[0:7] | | | |
| 0x24 | | | | |
| 0x28 | RCWH[8:15] | | | |
| 0x2C | | | | |
| 0x30 | RCWH[16:23] | | | |
| 0x34 | | | | |
| 0x38 | RCWH[24:31] | | | |
| 0x3C | | | | |

4.3.3.1.1 Local Bus Controller Setting

The device will use GPCM to load the reset configuration from EEPROM or NOR Flash. The device will read 64 bytes in this case. The local bus controller's registers setting will be set according to [Table 4-23](#).

The device will use FCM to load the reset configuration from NAND Flash. The device will read 512 bytes if small size NAND Flash is used, or 2048 bytes if large page NAND Flash is used. The local bus controller's registers setting will be set according to [Table 4-23](#).

The device will use PCI_SYNC_IN clock to generate the internal LCLK, which will run at half the frequency of PCI_SYNC_IN.

Table 4-23. Local Bus Controller Setting when Loading RCW

| CFG_RESET_SOURCE | Meaning | BR0[PS] | BR0[MSEL] | OR0[SCY] | OR0[PGS] |
|------------------|-------------------------------|---------|-----------|----------|----------|
| 0000 | NOR Flash | 11 | 000 | 1111 | NA |
| 0001 | NAND Flash, 8 bit, small page | 01 | 001 | 0010 | 0 |
| 0101 | NAND Flash, 8 bit, large page | 01 | 001 | 0010 | 1 |

4.3.3.2 Loading from I²C EEPROM

The device is capable of loading the reset configuration word from the I²C interface. If the device is configured to load the reset configuration word from the I²C interface, according to the reset configuration

input signals, it uses the I²C unit boot sequencer in a special mode. In this mode, the I²C boot sequencer is activated while the rest of the device is still in reset state ($\overline{\text{HRESET}}$ asserted) to load the reset configuration words from an I²C serial EEPROM.

Note that this does not prevent using the I²C boot sequencer to initiate the device in the normal functional mode after reset state has completed. The only restriction is that the first two EEPROM data structures contain dedicated reset information.

4.3.3.2.1 Using the Boot Sequencer Reset Configuration

For a detailed description about the I²C interface and the boot sequencer refer to [Section 20.4.5, “Boot Sequencer Mode.”](#)

NOTE

When reset configuration words are loaded from an I²C EEPROM, an I²C serial EEPROM of extended addressing type must be used.

If the I²C interface is used for loading the reset configuration words, the I²C module addresses the EEPROM and reads the first two data structures (after reading the preamble). Upon being read, the reset configuration words are latched inside the device and the I²C module enters its reset state until $\overline{\text{HRESET}}$ is negated. There should be no other I²C traffic when the boot sequencer is active.

After $\overline{\text{HRESET}}$ is negated, the functional boot sequencer, in extended I²C addressing mode, may be activated if the BOOTSEQ field of the reset configuration word high is set to 0b10.

4.3.3.2.2 EEPROM Calling Address

The device uses 0b101_0000 for the EEPROM calling address. The EEPROM to be addressed must contain the reset configuration information and be programmed to respond to this address. No additional EEPROMs are accessed by the boot sequencer in reset configuration mode.

4.3.3.2.3 EEPROM Data Format in Reset Configuration Mode

The I²C module expects that a particular data format be used for data in the EEPROM. A preamble should be the first 3 bytes programmed into the EEPROM. It should have a value of 0xAA55AA. The I²C module checks to ensure that this preamble is correctly detected before proceeding further. Following the preamble, there should be the two reset configuration words, programmed according to a particular format, as shown in [Figure 4-5](#).

The first 3 bytes hold the attributes and address offset. The addresses of the two reset configuration words must be programmed to the offset of the reset configuration word low register (RCWLR) and reset configuration word high register (RCWHR) respectively (see [Section 4.5.1.1, “Reset Configuration Word Low Register \(RCWLR\),”](#) and [Section 4.5.1.2, “Reset Configuration Word High Register \(RCWHR\)”](#)). The attributes should be programmed as follows: alternate configuration space (ACS) should be cleared (0b0), byte enables should be all ones, and continue (CONT) should be set.

After the first 3 bytes, 4 bytes of data should hold the desired value of the reset configuration word. The boot sequencer assumes that a big-endian address is stored in the EEPROM.

IMMRBAR value is prepended to the EEPROM address to generate the complete memory-mapped register's address.

When the I²C operates in reset configuration mode, the cyclic redundancy check (CRC) is ignored, as well as any registers following the first two reset configuration words.

| 0 | 1 | 4 | 5 | 6 | 7 |
|---------------------------------------|-------------------|---|---|-------------|----------------------|
| ACS (0) | BYTE_EN (1111) | | | CONT (1) | RCWLR ADDR[12–13] |
| RCWLR ADDR[14:21] | | | | | |
| RCWLR ADDR[22:29] | | | | | |
| Reset configuration word low [0–7] | | | | | |
| Reset configuration word low [8–15] | | | | | |
| Reset configuration word low [16–23] | | | | | |
| Reset configuration word low [24–31] | | | | | |
| ACS (0) | BYTE_EN (1111) | | | CONT (1) | RCWHR ADDR[12–13] |
| RCWHR ADDR[14–21] | | | | | |
| RCWHR ADDR[22–29] | | | | | |
| Reset configuration word high [0–7] | | | | | |
| Reset configuration word high [8–15] | | | | | |
| Reset configuration word high [16–23] | | | | | |
| Reset configuration word high [24–31] | | | | | |

Figure 4-5. EEPROM Data Format for Reset Configuration Words Preload Command

Figure 4-6 shows an example of the EEPROM contents, including the preamble, reset configuration words and additional initialization data, and CRC. In this example, it is assumed that the EEPROM contains

information additional to the reset configuration words, which should be loaded in the functional state after the device completes its reset flow.

| | | | | | | | | | |
|---------------------------------------|---------|---|---|---|---|-------------------|-------------------------|---|------------------------------------|
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | Preamble | |
| 1 | 0 | 1 | 0 | 1 | 0 | 1 | 0 | | |
| 0 | 1 | 0 | 1 | 0 | 1 | 0 | 1 | | |
| 1 | 0 | 1 | 0 | 1 | 0 | 1 | 0 | | |
| 0 | 1 | 1 | 1 | 1 | 1 | RCWLR ADDR[12:13] | | Reset configuration word low preload command | |
| RCWLR ADDR[14-21] | | | | | | | | | |
| RCWLR ADDR[22-29] | | | | | | | | | |
| Reset configuration word low [0-7] | | | | | | | | | |
| Reset configuration word low [8-15] | | | | | | | | | |
| Reset configuration word low [16-23] | | | | | | | | | |
| Reset configuration word low [24-31] | | | | | | | | | |
| 0 | 1 | 1 | 1 | 1 | 1 | RCWHR ADDR[12:13] | | Reset configuration word high preload command | |
| RCWHR ADDR[14-21] | | | | | | | | | |
| RCWHR ADDR[22-29] | | | | | | | | | |
| Reset configuration word high [0-7] | | | | | | | | | |
| Reset configuration word high [8-15] | | | | | | | | | |
| Reset configuration word high [16-23] | | | | | | | | | |
| Reset configuration word high [24-31] | | | | | | | | | |
| * | | | | | | | | | |
| ACS | BYTE_EN | | | | 1 | ADDR[12-13] | | | Last configuration preload command |
| ADDR[14-21] | | | | | | | | | |
| ADDR[22-29] | | | | | | | | | |
| DATA[0-7] | | | | | | | | | |
| DATA[8-15] | | | | | | | | | |
| DATA[16-23] | | | | | | | | | |
| DATA[24-31] | | | | | | | | | |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | End command | |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | | |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | | |
| CRC[0-7] | | | | | | | Cyclic redundancy check | | |
| CRC[8-15] | | | | | | | | | |
| CRC[16-23] | | | | | | | | | |
| CRC[24-31] | | | | | | | | | |

Figure 4-6. EEPROM Contents

4.3.3.2.4 Reset Configuration Load Fail

Failure of reset configuration load by the I²C boot sequencer can be caused by an incorrect EEPROM data structure or I²C bus problem. If a reset configuration load failure occurs, due to preamble fail or any other I²C bus error detection, the device will continuously attempt to reload the hard reset configuration words from the I²C bus. The device does not negate $\overline{\text{HRESET}}$ and remains in hard reset state until the HRCWs are successfully loaded or the PORESET flow is restarted.

4.3.3.3 Default Reset Configuration Words

If the device is configured not to load the reset configuration words from NOR Flash, NAND Flash, or an I²C EEPROM, it can also be initialized with one of five hard-coded default options, selected by the reset configuration input signals, CFG_RESET_SOURCE[0:3]. In this mode, the device is assumed to be a PCI agent, and therefore only clock modes differ among the four options.

The reset configuration words are driven internally with the values shown in [Table 4-24](#) and [Table 4-25](#).

NOTE

In this mode the device is also configured to accept PCI configuration cycles when completing its reset sequence (In PCI function configuration register, the CFG_LOCK bit is cleared). In addition, the inbound window size of the PCI inbound window attribute registers (PIWAR_n[IWS]) is set to 0b010011, defining 1-Mbyte ($2^{(19+1)}$) memory windows. See [Section 13.3.3.24, “PCI Function Configuration Register.”](#)

Table 4-24. Hard Coded Reset Configuration Word Low Fields Values

| RCWL Bits: | 0 | 1 | 2–3 | 4–7 | 8 | 9–15 | 16–31 |
|------------------------|--------------------------------------|--------------------------------------|-----|---|-----|-------------------------------------|-------|
| Field: | LBCM | DDRCM | Res | SPMF | Res | COREPLL | Res |
| Meaning: | LBC controller clock: <i>csb_clk</i> | DDR controller clock: <i>csb_clk</i> | — | <i>csb_clk</i> : PCI_CLK ratio SPMF:1 | — | Core clock: <i>csb_clk</i> ratio | — |
| CFG_RESET_SOURCE Value | 0 1:1 1 2:1 | 0 1:1 1 2:1 | | | | | |
| 1000 | 0 | 1 | 00 | 0100 | 0 | 0000100 | 16'b0 |
| 1001 | 0 | 1 | 00 | 0010 | 0 | 0000101 | 16'b0 |
| 1010 | 0 | 1 | 00 | 0100 | 0 | 0000101 | 16'b0 |
| 1011 | 0 | 1 | 00 | 0010 | 0 | 0000110 | 16'b0 |
| 1100 | 0 | 1 | 00 | 0100 | 0 | 0000110 | 16'b0 |

Table 4-25 defines the hard-coded reset configuration word high fields values. These values select hard-coded reset configuration words options, as described in Section 4.3.1.1, “Reset Configuration Word Source.”

Table 4-25. Hard-Coded Reset Configuration Word High Field Values

| Bits | Name | Field Values when CFG_RESET_SOURCE[0:3] = 1000–1100 | | | | | Meaning |
|-------|----------|--|------|------|------|------|---|
| | | 1000 | 1001 | 1010 | 1011 | 1100 | |
| 0 | PCIHOST | 0 | | | | | PCI agent mode |
| 1 | Reserved | 1 | | | | | — |
| 2 | PCIARB | 0 | | | | | External arbiter is used |
| 3 | Reserved | 0 | | | | | — |
| 4 | COREDIS | 1 | | | | | e300 core is disabled (boot holdoff) |
| 5 | BMS | 1 | | | | | Boot memory space is 0xFF80_0000–0xFFFF_FFFF. MSR[IP] initial value is 0b1. |
| 6–7 | BOOTSEQ | 00 | | | | | Boot sequencer is disabled. |
| 8 | SWEN | 0 | | | | | Software watchdog disabled. |
| 9–11 | ROMLOC | 000 | | | | | Boot ROM interface location. |
| 12–13 | RLEXT | 00 | | | | | Legacy mode. |
| 14–15 | Reserved | 00 | | | | | — |
| 16–18 | TSEC1M | 101 | 011 | 000 | 001 | 000 | 000 = MII mode 001 = RMII 011 = RGMII mode 101 = RTBI mode |
| 19–21 | TSEC2M | 000 | 101 | 001 | 101 | 011 | |
| 22–27 | Reserved | 000000 | | | | | — |
| 28 | TLE | 0 | | | | | Big-endian mode |
| 29 | LALE | 0 | | | | | Normal timing |
| 30–31 | Reserved | 00 | | | | | — |

4.3.3.3.1 Examples for Hard-Coded Reset Configuration Words Usage

Examples for various clock modes are listed in Table 4-26.

Table 4-26. Examples For Hard-Coded Reset Configuration Words Usage

| CFG_RESET_SOURCE[0:3] | 1000 | 1001 | 1010 | 1011 | 1100 |
|----------------------------|------|------|------|------|------|
| PCI_CLK (MHz) | 33 | 66 | 33 | 66 | 33 |
| <i>csb_clk</i> (MHz) | 133 | 133 | 133 | 133 | 133 |
| DDR Controller clock (MHz) | 266 | 266 | 266 | 266 | 266 |
| Core clock (MHz) | 266 | 333 | 333 | 400 | 333 |

4.4 Clocking

The following external clock sources are utilized on the MPC8315E:

- System clock (SYS_CLK_IN or PCI_CLK)
- Ethernet clock (GTX_CLK125 for eTSEC n or SD_REF_CLK/ $\overline{\text{SD_REF_CLK}}$ for SGMII PHY interface)
- USB clock (USB_CLK_IN for USB PHY)
- Real-time clock (RTC_CLK)

All clock inputs can be supplied using an external canned oscillator, a clock generation chip, or some other source that provides a standard CMOS square wave input.

The system and USB clock sources can alternatively be provided using an external crystal (on-chip crystal oscillator provided). Separate crystal (CR) inputs/outputs are provided for this case. The crystal oscillator dissipates 50 mW of current, so this may not be the best choice for very low power applications. When a crystal is not used, a square wave clock input must be supplied.

If SYS_CLK_IN is used as the system clock, the SYS_XTAL_IN (crystal input) must be tied to quiet ground; otherwise, SYS_CLK_IN must be tied to quiet ground. Similarly, if USB_CLK_IN is used as the USB clock, USB_XTAL_IN (crystal input) must be tied to quiet ground; otherwise, USB_CLK_IN must be tied to quiet ground.

Figure 4-7 shows the internal distribution of clocks within the device.




Figure 4-7. Clock Subsystem Block Diagram

The primary clock input to this device is PCI_CLK. This clock is the reference to the system APLL.

4.4.1 Clocking in PCI Host Mode

When the device is configured as a PCI host device (RCWH[PCIHOST] = 1), SYS_CLK_IN is the primary input clock. SYS_CLK_IN feeds the PCI clock divider ($\div 2$) and the PCI_SYNC_OUT and PCI_CLK_OUT multiplexers. The CFG_CLKIN_DIV configuration input selects whether SYS_CLK_IN or SYS_CLK_IN/2 is driven out on the PCI_SYNC_OUT and PCI_CLK_OUT signals.

PCI_SYNC_OUT is connected externally to PCI_SYNC_IN to allow the internal clock subsystem to synchronize to the system PCI clocks. PCI_SYNC_OUT must be connected properly to PCI_SYNC_IN, with equal delay to all PCI agent devices in the system.

4.4.1.1 PCI Clock Outputs (PCI_CLK_OUT[0:2])

When the device is configured as a PCI host, it provides three clock output signals, PCI_CLK_OUT[0:2], for external PCI agents.

When the device comes out of reset, the PCI clock outputs are disabled and are actively driven to a steady low state. Each of the individual clock outputs can be enabled (enable toggling of the clock) by setting its corresponding OCCR[PCICOEn] bit. All output clocks are phase aligned to each other and to PCI_SYNC_OUT.

4.4.2 Clocking In PCI Agent Mode

When the device is configured as a PCI agent, PCI_CLK is the primary input clock. In agent mode, the SYS_CLK_IN signal may not be used. If it is unused, SYS_CLK_IN should be tied to GND, and the clock output signals, PCI_CLK_OUT n and PCI_SYNC_OUT, are not used.

In agent mode, the $\overline{\text{CFG_CLKIN_DIV}}$ configuration input can be used to double the internal clock frequencies, if sampled as 1 during PORESET assertion. This feature is useful if a fixed internal frequency is desired regardless of whether the PCI clock is running at 33 or 66 MHz. PCI specifications require that the signal M66EN provides the PCI clock frequency information.

4.4.3 System Clock Domains

As shown in [Figure 4-7](#), the primary clock input (PCI_CLK/PCI_SYNC_IN) frequency is multiplied up by the system phase-locked loop (PLL) and the clock unit to create three major clock domains:

- The coherent system bus clock (*csb_clk*)
- The internal clock for the DDR controller (*ddr_clk*)
- The internal clock for the local bus interface unit (*lbc_clk*)

The *csb_clk* frequency is derived from a complex set of factors that can be simplified into the following equation:

$$csb_clk = [\text{PCI_SYNC_IN} \times (1 + \overline{\text{CFG_CLKIN_DIV}})] \times \text{SPMF}$$

In PCI host mode, $\text{PCI_SYNC_IN} \times (1 + \overline{\text{CFG_CLKIN_DIV}})$ is the SYS_CLK_IN frequency.

The *csb_clk* serves as the clock input to the e300 core. A second PLL inside the core multiplies up the *csb_clk* frequency to create the internal clock for the core (*core_clk*). The system and core PLL multipliers are selected by the SPMF and COREPLL fields in the reset configuration word low (RCWL), which is loaded at power-on reset or by one of the hard-coded reset options. See [Section 4.3, “Reset Configuration.”](#)

The DDR SDRAM memory controller will operate with a frequency equal to twice the frequency of *csb_clk*. Note that *ddr_clk* is not the external memory bus frequency; *ddr_clk* passes through the DDR

clock divider ($\div 2$) to create the differential DDR memory bus clock outputs (MCK and $\overline{\text{MCK}}$). However, the data rate is the same frequency as *ddr_clk*.

The local bus memory controller will operate with a frequency equal to the frequency of *csb_clk*. Note that *lbc_clk* is not the external local bus frequency; *lbc_clk* passes through the LBC clock divider to create the external local bus clock output (LCLK[]). The LBC clock divider ratio is controlled by LCCR[CLKDIV]. See [Section 10.1.3.1, “eLBC Bus Clock and Clock Ratios,”](#) for more information.

In addition, some of the internal units may be required to be shut off or operate at lower frequency than the *csb_clk* frequency. These units have a default clock ratio that can be configured by a memory-mapped register after the device comes out of reset. [Table 4-27](#) specifies which units have a configurable clock frequency. Refer to [Section 4.5.2.3, “System Clock Control Register \(SCCR\).”](#)

Table 4-27. Configurable Clock Units

| Unit | Default Frequency | Options |
|---------------------------------|-------------------|---|
| eTSEC1 and eTSEC2 | <i>csb_clk</i> | Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i> |
| Security core, I ² C | <i>csb_clk</i> | Off, <i>csb_clk/2</i> , <i>csb_clk/3</i> |
| USB DR | <i>csb_clk</i> | Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i> |
| PCI and DMA complex | <i>csb_clk</i> | Off, <i>csb_clk</i> |
| PCI Express 1 and PCI Express 2 | <i>csb_clk</i> | Off, <i>csb_clk</i> |
| SATA1/SATA2 | <i>csb_clk</i> | Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i> |
| TDM | <i>csb_clk</i> | Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i> |

NOTE

The clock ratios of these units must be set before they are accessed.

eTSEC1 and eTSEC2 share the same clock control, and therefore the same clock ratio. They can be independently switched off.

NOTE

A portion of the eTSEC_n and USB DR controller run at the line rate on a clock provided by the SGMII and USB PHY blocks, respectively. Synchronization between the line rate clock and the *csb_clk* is provided in these controllers.

4.4.4 USB Clocking

If the on-chip USB PHY is utilized, the reference input can be provided externally using a separate clock source, either a crystal or an external oscillator. The PHY supplies the clock to the USB DR controller in UTMI mode (when the on-chip PHY is used). Synchronization between the PHY clock domain and the CSB clock domain occurs in the USB controller.

An option is provided to supply the USB reference clock from the SYS_CLK_IN or SYS_XTAL_IN inputs. This allows for a single crystal or clock input to supply both system and USB references. The USB reference clock can be provided with a divide by 1 or 2 from these inputs (see [Figure 4-7](#)). When using the

single crystal option, the frequency for SYS_CLK_IN must be chosen such that the USB reference will be 24 or 48 MHz when utilizing the divide by 1 or 2 option, that is, the SYS_CLK_IN must be 24 or 48 MHz.

If this option is used in PCI agent mode, the PCI clock supplied to device must still be chosen at the appropriate frequencies mentioned above. In this case the PCI source clock would be tied to both SYS_CLK_IN and PCI_CLK inputs. The clock source will need to meet the input clock specifications for both SYS_CLK_IN and USB_CLK_IN.

For more details refer to [Chapter 17, “Universal Serial Bus Interface.”](#)

4.4.5 Ethernet Clocking

The device contains an integrated high speed serial/deserializing (SerDes) PHY block that provides an SGMII interface to an external PHY. The SGMII PHY interface has its own PLL and requires a reference clock which is 1 V signal that can either be single ended or differential. The SerDes PHY generates its own transmit and receive sampling clock. The receive clock is re-created from the receive data.

When running in RGMII or RTBI modes (not using the SerDes) the Gigabit reference clock is the GTX_CLK125 input on the eTSEC1 interface which is common to both eTSECs. This can either be a 2.5- or 3.3-V signal.

For more information refer to [Chapter 19, “Enhanced Three-Speed Ethernet Controllers.”](#)

4.4.6 Real-Time Clock (RTC)

The source for the RTC can come from the internal system clock (csb_clk) through a divider circuit or from an off-chip source. For more information refer to [Section 5.5, “Real Time Clock Module \(RTC\).”](#)

4.5 Memory Map/Register Definitions

This section presents the memory maps and register descriptions for both reset and clocking.

4.5.1 Reset Configuration Register Descriptions

The reset configuration and status registers are shown in [Table 4-28](#).

Table 4-28. Reset Configuration and Status Registers Memory Map

| Address | Register | Access | Reset | Section/Page |
|--|--|--------|-----------|------------------------------|
| Reset Configuration—Block Base Address 0x0_0900 | | | | |
| 0x0_0900 | Reset configuration word low register (RCWLR) | R | All zeros | 4.5.1.1/4-34 |
| 0x0_0904 | Reset configuration word high register (RCWHR) | R | All zeros | 4.5.1.2/4-34 |
| 0x0_0908 | Reserved, should be cleared | — | — | — |
| 0x0_090C | Reserved, should be cleared | — | — | — |
| 0x0_0910 | Reset status register (RSR) | R/W | All zeros | 4.5.1.3/4-34 |
| 0x0_0914 | Reset mode register (RMR) | R/W | All zeros | 4.5.1.4/4-36 |

Table 4-28. Reset Configuration and Status Registers Memory Map (continued)

| Address | Register | Access | Reset | Section/Page |
|-----------------------|--------------------------------------|--------|-----------|------------------------------|
| 0x0_0918 | Reset protection register (RPR) | R/W | All zeros | 4.5.1.5/4-36 |
| 0x0_091C | Reset control register (RCR) | R/W | All zeros | 4.5.1.6/4-37 |
| 0x0_0920 | Reset control enable register (RCER) | R/W | All zeros | 4.5.1.7/4-38 |
| 0x0_0924– 0x0_09FC | Reserved, should be cleared. | — | — | — |

4.5.1.1 Reset Configuration Word Low Register (RCWLR)

The reset configuration word low register (RCWLR) is shown in [Figure 4-3](#) and described in [Section 4.3.2.1, “Reset Configuration Word Low Register \(RCWLR\).”](#)

4.5.1.2 Reset Configuration Word High Register (RCWHR)

The reset configuration word high register (RCWHR) is shown in [Figure 4-4](#) and described in [Section 4.3.2.2, “Reset Configuration Word High Register \(RCWHR\).”](#)

4.5.1.3 Reset Status Register (RSR)

RSR, shown in [Figure 4-8](#), captures various reset events in the device. The RSR accumulates reset events. For example, because software watchdog expiration results in a hard reset, SWRS and HRS are all set after a software watchdog reset. This register returns to its reset value only when power-on reset occurs.

Address 0x0_0910

Access: User read/write

| | | | | | | | | | | | | | | |
|-------|-----------|------|----|----|------|----|----|------|------|------|----|-----|----|----|
| | 0 | 3 | 4 | | | | | | | 14 | 15 | | | |
| R | RSTSRC | | — | | | | | | BSF | | | | | |
| W | n1 | | — | | | | | | 0 | | | | | |
| Reset | n1 | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | | | |
| | 16 | 17 | 18 | 19 | 20 | 22 | 23 | 24 | 26 | 27 | 28 | 29 | 30 | 31 |
| R | — | SWHR | — | | JSRS | — | | CSHR | SWRS | BMRS | — | HRS | | |
| W | — | SWHR | — | | JSRS | — | | CSHR | SWRS | BMRS | — | HRS | | |
| Reset | All zeros | | | | | | | | | | | | | |

- ¹ The reset value of this field is determined according to the reset configuration input signals CFG_RESET_SOURCE[0:2] sampled during the reset flow.

Figure 4-8. Reset Status Register (RSR)

[Table 4-29](#) defines the reset status register bit fields.

Table 4-29. Reset Status Register Field Descriptions

| Bits | Name | Description |
|------|--------|--|
| 0–3 | RSTSRC | Reset configuration word source. Reflects the value of CFG_RESET_SOURCE input signal during the reset flow. See Section 4.3.1.1, “Reset Configuration Word Source.” Changing this field has no effect. |
| 4–14 | — | Reserved, should be cleared. |

Table 4-29. Reset Status Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 15 | BSF | Boot sequencer fail. If set, indicates that the I ² C boot sequencer has failed while loading the reset configuration words. Cleared by writing a 1 to it (writing zero has no effect). |
| 16–18 | — | Reserved, should be cleared. |
| 19 | SWHR | Software hard reset. If set, indicates a software hard reset. SWHR is cleared by writing a 1 to it (writing zero has no effect). |
| 20–22 | — | Reserved, should be cleared. |
| 23 | JSRS | JTAG soft reset status. Set when the JTAG reset request is set and remains set until software clears it. JSRS is cleared by writing a 1 to it (writing zero has no effect). 0 No JTAG reset event. 1 JTAG reset event. |
| 24–26 | — | Reserved, should be cleared. |
| 27 | CSHR | Check stop reset status. When the core enters a checkstop state and the checkstop reset is enabled by the RMR[CSRE], CSRS is set and it remains set until software clears it. CSRS is cleared by writing a 1 to it (writing zero has no effect). 0 No enabled check stop reset event. 1 Enabled check stop reset event. |
| 28 | SWRS | Software watchdog reset status. When a software watchdog expire event (which causes a reset) is detected, SWRS is set and remains that way until the software clears it. SWRS is cleared by writing a 1 to it (writing zero has no effect). 0 No software watchdog reset event. 1 Software watchdog reset event. |
| 29 | BMRS | Bus monitor reset status. When a bus monitor expire event (which causes a reset) is detected, BMRS is set and remains set until the software clears it. BMRS can be cleared by writing a 1 to it (writing zero has no effect). 0 No bus monitor reset event. 1 Bus monitor reset event. |
| 30 | — | Reserved |
| 31 | HRS | Hard reset status. When an external or internal hard reset event is detected, HRS is set and remains set until software clears it. HRS is cleared by writing a 1 (writing zero has no effect). 0 No hard reset event. 1 Hard reset event. |

4.5.1.4 Reset Mode Register (RMR)

RMR, shown in [Figure 4-9](#), enables a hard reset sequence on the device when the e300 core enters checkstop state.




Figure 4-9. Reset Mode Register (RMR)

[Table 4-30](#) describes the RMR fields.

Table 4-30. RMR Field Descriptions

| Bits | Name | Function |
|------|------|--|
| 0–30 | — | Reserved, should be cleared. |
| 31 | CSRE | Checkstop reset enable. The core can enter checkstop mode as the result of several exception conditions. Setting CSRE configures the device to perform a hard reset sequence when the core enters checkstop state. 0 Reset not generated when core enters checkstop state. 1 Reset generated when core enters checkstop state. |

4.5.1.5 Reset Protection Register (RPR)

RPR, shown in [Figure 4-10](#), prevents unintended software reset requests caused by writes to the reset control register (RCR). To disable a write to the reset control register (RCR), the user should write a 1 to RCER[CRE].




Figure 4-10. Reset Protection Register (RPR)

Table 4-31 defines the bit fields of RPR.

Table 4-31. RPR Bit Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–31 | RCPW | Reset control protection word. Prevents unintended software reset requests because of a write to the RCR. The user should write the value 0x5253_5445 (RSTE in ASCII) to enable. Enable indication appears in the reset control enable register (RCER[CRE]). Reading this register always returns all zeros. |

4.5.1.6 Reset Control Register (RCR)

RCR, shown in Figure 4-11, can be used by software to initiate a hard reset sequence. To allow writing to this register, the user must enable it by writing the value 0x5253_5445 to the RPR.



Figure 4-11. Reset Control Register (RCR)

Table 4-32 defines the bit fields of RCR.

Table 4-32. RCR Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–29 | — | Reserved, should be cleared. |
| 30 | SWHR | Software hard reset. Setting this bit causes the device to begin a hard reset flow. This bit returns to its reset state during the reset sequence, so reading it always returns all zeros. |
| 31 | — | Reserved. |

4.5.1.7 Reset Control Enable Register (RCER)

RCER, shown in Figure 4-12, indicates by the CRE field that the RPR is accessed with a value that enables RCR.



Figure 4-12. Reset Control Enable Register (RCER)

Table 4-33 defines the bit fields of RCER.

Table 4-33. RCER Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–30 | — | Reserved, should be cleared. |
| 31 | CRE | Control register enabled. When set, indicates that the RPR was accessed with a value that enables the RCR. Writing 1 to this bit disables the RCR and clears this bit. Writing zero has no effect. |

4.5.2 Clock Configuration Registers

The clock configuration and status registers are shown in Table 4-34.

Table 4-34. Clock Configuration Registers Memory Map

| Address | Register | Access | Reset | Section/Page |
|--|--------------------------------------|--------|--------------------------------------|--------------|
| Reset Configuration—Block Base Address 0x0_0900 | | | | |
| 0x0_0A00 | System PLL mode register (SPMR) | R | 0xn _{nnn} _n _{nnn} | 4.5.2.1/4-38 |
| 0x0_0A04 | Output clock control register (OCCR) | R/W | All zerosC0C0 | 4.5.2.2/4-40 |
| 0x0_0A08 | System clock control register (SCCR) | R/W | 0x5155_1410 | 4.5.2.3/4-41 |
| 0x0_0A0C– 0x0_0AFC | Reserved, should be cleared | — | — | — |

4.5.2.1 System PLL Mode Register (SPMR)

SPMR is shown in Figure 4-13, gets its values according to the $\overline{\text{CFG_CLKIN_DIV}}$ reset configuration input signal and the reset configuration word low loaded during the reset flow. Note that this register is

4.5.2.2 Output Clock Control Register (OCCR)

The OCCR, shown in [Figure 4-14](#), controls the device output clocks. It is possible to control some output clock modes by writing to this memory-mapped register as described below.

Addr: 0x0_0A04

Access: Read/Write




Figure 4-14. Output Clock Control Register (OCCR)

[Table 4-36](#) defines the bit fields of OCCR.

Table 4-36. OCCR Bit Settings

| Bits | Name | Description |
|-------|------------------|---|
| 0 | PCICOE0 | PCI_CLK_OUT0 enable. 0 PCI_CLK_OUT0 signal is disabled (drive constant zero). 1 PCI_CLK_OUT0 signal is enabled to toggle. |
| 1 | PCICOE1 | PCI_CLK_OUT1 enable. 0 PCI_CLK_OUT1 signal is disabled (drive constant zero). 1 PCI_CLK_OUT1 signal is enabled to toggle. |
| 2 | PCICOE2 | PCI_CLK_OUT2 enable. 0 PCI_CLK_OUT2 signal is disabled (drive constant zero). 1 PCI_CLK_OUT2 signal is enabled to toggle. |
| 3–15 | — | Reserved, should be cleared |
| 16 | MCK0OE, MCK_B0OE | Enable/Disable MCK[0] pins clock out 0 Disable MCK[0] and MCK[0] 1 Enable MCK[0] and MCK[0] |
| 17 | MCK1OE, MCK_B1OE | Enable/Disable MCK[1] pins clock out 0 Disable MCK[1] and MCK[1] 1 Enable MCK[1] and MCK[1] |
| 18–23 | — | Reserved, should be cleared |
| 24 | LCLK0E | Enable/Disable LCLK[0] pin clock out 0 Disable LCLK[0] 1 Enable LCLK[0] |

Table 4-36. OCCR Bit Settings (continued)

| Bits | Name | Description |
|-------|--------|---|
| 25 | LCLK1E | Enable/Disable LCLK[1] pin clock out 0 Disable LCLK[1] 1 Enable LCLK[1] |
| 28–31 | — | Reserved |

4.5.2.3 System Clock Control Register (SCCR)

SCCR, shown in [Figure 4-15](#), controls device units that have a configurable clock ratio.

Address 0x0_0A08

Access: Read/Write

| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|-------|---------|---------|---------|-------|---------|-----------|-----------|----|-------|----|----|----|----|----|----|----|
| R | TSEC1CM | TSEC2CM | — | ENCCM | USBDRCM | PCIEXP1CM | PCIEXP2CM | — | PCICM | | | | | | | |
| W | | | | | | | | | | | | | | | | |
| Reset | 0 | 1 | 0 | 1 | 0 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | 0 | 1 | 0 | 1 |
| | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 25 | 26 | 27 | 28 | 31 | | | | |
| R | — | SATA1CM | SATA2CM | — | TDMCM | | | | | | | | | | | |
| W | | | | | | | | | | | | | | | | |
| Reset | 0 | 0 | 0 | 1 | 0 | 1 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | |

Figure 4-15. System Clock Control Register (SCCR)

[Table 4-37](#) defines the bit fields of SCCR.

Table 4-37. SCCR Bit Descriptions

| Bits | Name | Description |
|------|---------|--|
| 0–1 | TSEC1CM | TSEC1 clock mode. 00 TSEC1 clock is disabled. 01 TSEC1 clock/ <i>csb_clk</i> ratio is 1:1. 10 TSEC1 clock/ <i>csb_clk</i> ratio is 1:2 (<i>csb_clk</i> has higher frequency than TSEC1). 11 TSEC1 clock/ <i>csb_clk</i> ratio is 1:3 (<i>csb_clk</i> has higher frequency than TSEC1). |
| 2–3 | TSEC2CM | TSEC2 clock mode. 00 TSEC2 clock is disabled. 01 TSEC2 clock/ <i>csb_clk</i> ratio is 1:1. 10 TSEC2 clock/ <i>csb_clk</i> ratio is 1:2 (<i>csb_clk</i> has higher frequency than TSEC2). 11 TSEC2 clock/ <i>csb_clk</i> ratio is 1:3 (<i>csb_clk</i> has higher frequency than TSEC2). |
| 4–5 | — | Reserved |
| 6–7 | ENCCM | Encryption core and I ² C clock mode. 00 Encryption core clock is disabled. 01 Encryption core clock/ <i>csb_clk</i> ratio is 1:1. 10 Encryption core clock/ <i>csb_clk</i> ratio is 1:2 (<i>csb_clk</i> has higher frequency than the encryption core). 11 Encryption core clock/ <i>csb_clk</i> ratio is 1:3 (<i>csb_clk</i> has higher frequency than the encryption core). Note: The encryption core must have the same clock ratio as the USB unit, unless one of them has its clock disabled. |

Table 4-37. SCCR Bit Descriptions (continued)

| Bits | Name | Description |
|-------|-----------|---|
| 8–9 | USB DRCM | USB DR clock mode. 00 USB DR clock is disabled. 01 USB DR clock/ <i>csb_clk</i> ratio is 1:1. 10 USB DR clock/ <i>csb_clk</i> ratio is 1:2 (<i>csb_clk</i> has higher frequency than the USB DR). 11 USB DR clock/ <i>csb_clk</i> ratio is 1:3 (<i>csb_clk</i> has higher frequency than the USB DR). Note: The USB DR unit must have the same clock ratio as the encryption core unit, unless one of them has its clock disabled. |
| 10–11 | PCIEXP1CM | PCIEXP1 clock mode. Define the clock mode for the PCI Express 1 controller. 00 PCIEXP1 clock is disabled. 01 PCIEXP1 clock is enabled. |
| 12–13 | PCIEXP2CM | PCIEXP2 clock mode. Define the clock mode for the PCI Express 2 controller. 00 PCIEXP2 clock is disabled. 01 PCIEXP2 clock is enabled. |
| 14 | — | Reserved |
| 15 | PCICM | PCI clock mode. Define the clock mode for all of the PCI complex - PCI and DMA. 0 PCI complex clocks are disabled. 1 PCI complex clocks are enabled. |
| 16–17 | — | Reserved |
| 18–19 | SATA1CM | SATA1 clock mode. Define the clock mode for the SATA1 controller. 00 SATA1 clock is disabled. 01 SATA1 clock/ <i>csb_clk</i> ratio is 1:1. 10 SATA1 clock/ <i>csb_clk</i> ratio is 1:2 (<i>csb_clk</i> has higher frequency than the SATA1 clock). 11 SATA1 clock/ <i>csb_clk</i> ratio is 1:3 (<i>csb_clk</i> has higher frequency than the SATA1 clock). |
| 20–21 | SATA2CM | SATA2 clock mode. Define the clock mode for the SATA2 controller. 00 SATA2 clock is disabled 01 SATA2 clock/ <i>csb_clk</i> ratio is 1:1. 10 SATA2 clock/ <i>csb_clk</i> ratio is 1:2 (<i>csb_clk</i> has higher frequency than the SATA2 clock). 11 SATA2 clock/ <i>csb_clk</i> ratio is 1:3 (<i>csb_clk</i> has higher frequency than the SATA2 clock). Note: All SATA controllers must have the same clock ratio, unless one of them has its clock disabled. |
| 22–25 | — | Reserved |
| 26–27 | TDMCM | TDM clock mode. Defines the clock mode for TDM controller. 00 TDM clock is disabled 01 TDM clock: <i>csb_clk</i> ratio is 1:1 10 TDM clock: <i>csb_clk</i> ratio is 1:2 (<i>csb_clk</i> has higher frequency than TDM clock) 11 TDM clock: <i>csb_clk</i> ratio is 1:3 (<i>csb_clk</i> has higher frequency than TDM clock) |
| 28–31 | — | Reserved |

Chapter 5

System Configuration

5.1 Introduction

This chapter describes several functions that control the local access windows, system configuration, protection, and general utilities. These functions are discussed in the following sections:

- [Section 5.2, “Local Memory Map Overview and Example”](#)
- [Section 5.3, “System Configuration”](#)
- [Section 5.4, “Software Watchdog Timer \(WDT\)”](#)
- [Section 5.5, “Real Time Clock Module \(RTC\)”](#)
- [Section 5.6, “Periodic Interval Timer \(PIT\)”](#)
- [Section 5.7, “General-Purpose Timers \(GTM\)”](#)
- [Section 5.8, “Power Management Control \(PMC\)”](#)

5.2 Local Memory Map Overview and Example

The device provides a flexible local memory map. The local memory map refers to the 32-bit address space seen by the processor as it accesses memory and I/O space. Internal DMA engines also see this same local memory map. All memory accessed by the DDR SDRAM and local bus memory controllers exists in this memory map, as do all memory-mapped configuration, control, and status registers.

The local memory map is defined by a set of eleven local access windows. Each of these windows maps a region of memory to a particular target interface, such as the DDR SDRAM controller or the PCI controller. The DSP subsystem is not operational in the MSC7104. Note that the local access windows do not perform any address translation. The size of each window can be configured from 4 Kbytes to 2 Gbytes. Each local access window is assigned to a specific target interface as specified in [Table 5-1](#).

Table 5-1. Local Access Windows Target Interface

| Window Number | Target Interface | Comments |
|---------------|--------------------------------|---------------------------|
| 0 | Configuration registers (IMMR) | Fixed 1-Mbyte window size |
| 1 | Local bus | — |
| 2 | Local bus | — |
| 3 | Local bus | — |
| 4 | Local bus | — |
| 5 | PCI | — |

Table 5-1. Local Access Windows Target Interface (continued)

| Window Number | Target Interface | Comments |
|---------------|------------------|----------|
| 6 | PCI | — |
| 7 | DDR SDRAM | — |
| 8 | DDR SDRAM | — |
| 9 | PCI Express 1 | — |
| 10 | PCI Express 2 | — |

Figure 5-1 shows an example memory map.




Figure 5-1. Local Memory Map Example

Table 5-2 shows one example of local access window settings.

Table 5-2. Local Access Windows Example

| Window | Base Address | Size | Target Interface |
|--------|--------------|------------|------------------|
| 7 | 0x0000_0000 | 2 Gbytes | DDR SDRAM |
| 2 | 0x8000_0000 | 1 Mbyte | Local bus |
| 5 | 0xA000_0000 | 512 Mbytes | PCI |
| 9 | 0x8100_0000 | 1 Mbyte | PCI Express 1 |
| 10 | 0x8200_0000 | 1 Mbyte | PCI Express 2 |

Table 5-2. Local Access Windows Example (continued)

| Window | Base Address | Size | Target Interface |
|---------|--------------|------------|--------------------------------|
| 3 | 0xC000_0000 | 256 Mbytes | Local bus |
| 0 | 0xFF40_0000 | 1 Mbyte | Configuration registers (IMMR) |
| 1 | 0xFF80_0000 | 8 Mbytes | Local bus boot ROM Flash |
| 4, 6, 8 | Unused | | |

In this example, the local access window of the boot ROM is defined as window number 1, on a local bus device, in the highest 8 Mbytes of memory as set by the reset configuration word high during the reset sequence (see [Section 4.3.2.2.4, “Boot ROM Location”](#)) and [Section 5.2.4.3.1, “LBLAWBAR0\[BASE_ADDR\] Reset Value.”](#) The local access window, which describes the range of memory used for memory-mapped registers (IMMR), is a fixed 1-Mbyte space pointed to by the IMMRBAR register, using its default value (0xFF40_0000). See [Section 5.2.4.1, “Internal Memory Map Registers Base Address Register \(IMMRBAR\).”](#)

5.2.1 Address Translation and Mapping

In addition to any address translation performed by the e300c3 core MMU, three distinct types of translation and mapping operations are performed on transactions at the integrated device level. These are as follows:

- Mapping a local address to a target interface
- Translating the local 32-bit address to an external address space
- Translating external addresses to the local 32-bit address space

The local access windows perform target mapping for transactions within the local address space. The local access windows do not perform any address translation.

Outbound windows perform the mapping from the local 32-bit address space to the address space of PCI, which may be much larger than the local space.

Inbound windows perform address translation from the external address spaces of PCI to the local address space.

The target mappings created by an inbound window must be consistent with those of the local access windows. That is, if an inbound window maps a transaction to a given local address, a valid local access window for that address must be set independently.

All of the configuration registers that define mapping of local access windows follow the same register format. [Table 5-3](#) summarizes the general format of these window definitions.

Table 5-3. Format of Window Definitions

| Register | Function |
|------------------------|--|
| Base address | High-order address bits defining location of the window in the initial address space |
| Window size/attributes | Window enable, window size ¹ |

¹ An exception is the IMMR window, which is always enabled and has a fixed 1-Mbyte size.

Windows must be a power-of-two size. To perform a mapping function, the address of the transaction is compared with the base address register of each window. The number of bits used in the comparison is dictated by each window's size attribute. When an address hits within a window, the transaction is directed to the appropriate target.

5.2.2 Window into Configuration Space

The internal memory map registers' base address register (IMMRBAR) defines a window that is used to access all memory-mapped configuration, control, and status registers, referred to as internal memory map registers or IMMR. This window is always enabled with a fixed size of 1 Mbyte, and no other attributes are attached so there is no associated size/attribute register. This window always takes precedence over all local access windows. The IMMRBAR always come out of reset with a default base address value of 0xFF40_0000, and this base address value can be modified by writing to this register. For more information, see [Section 5.2.4.1, "Internal Memory Map Registers Base Address Register \(IMMRBAR\)."](#)

NOTE

Although it is legal to use the 3-Mbyte space consecutive to the 1 Mbyte of the IMMR (for example, if IMMRBAR is 0xFF40_0000, the 3-Mbyte address space consecutive to it is 0xFF50_0000–0xFF7F_FFFF), it is not recommended. This space may be used in future derivatives of the device that require a larger internal memory space.

5.2.3 Local Access Windows

As demonstrated in the address map overview in [Section 5.2, "Local Memory Map Overview and Example,"](#) local access windows associate a range of the local 32-bit address space with a particular target interface. This allows the internal interconnections of the device to route a transaction from its source to the proper target. No address translation is performed. The base address defines the high order address bits that give the location of the window in the local address space. The window attributes enable the window and define its size, while the window number specifies the target interface.

With the exception of configuration space (mapped by IMMRBAR), all addresses used by the system must be mapped by a local access window. This includes addresses that are mapped by PCI inbound windows.

The local access window registers exist as part of the local access block in the system configuration registers. See [Section 5.3.2, "System Configuration Registers."](#) A detailed description of the local access window registers is given in the following sections. Note that the minimum size of a window is 4 Kbytes, so the low order 12 bits of the base address cannot be specified.

5.2.3.1 Local Access Register Memory Map

Table 5-4 shows the memory map for the local access registers.

Table 5-4. Local Access Register Memory Map

| Local Access—Block Base Address 0x0_0000 | | | | |
|--|---|--------|--------------------------|---------------|
| Local Memory Offset (Hex) | Register | Access | Reset | Section/Page |
| 0x0_0000 | Internal memory map base address register (IMMRBAR) | R/W | 0xFF40_0000 | 5.2.4.1/5-6 |
| 0x0_0004 | Reserved | — | — | — |
| 0x0_0008 | Alternate configuration base address register (ALTCBAR) | R/W | 0x0000_0000 | 5.2.4.2/5-7 |
| 0x0_000C– 0x0_001C | Reserved | — | — | — |
| 0x0_0020 | eLBC local access window 0 base address register (LBLAWBAR0) | R/W | 0x0000_0000 ¹ | 5.2.4.3/5-8 |
| 0x0_0024 | eLBC local access window 0 attribute register (LBLAWAR0) | R/W | 0x0000_0000 ² | 5.2.4.4/5-9 |
| 0x0_0028 | eLBC local access window 1 base address register (LBLAWBAR1) | R/W | 0x0000_0000 | 5.2.4.3/5-8 |
| 0x0_002C | eLBC local access window 1 attribute register (LBLAWAR1) | R/W | 0x0000_0000 | 5.2.4.4/5-9 |
| 0x0_0030 | eLBC local access window 2 base address register (LBLAWBAR2) | R/W | 0x0000_0000 | 5.2.4.3/5-8 |
| 0x0_0034 | eLBC local access window 2 attribute register (LBLAWAR2) | R/W | 0x0000_0000 | 5.2.4.4/5-9 |
| 0x0_0038 | eLBC local access window 3 base address register (LBLAWBAR3) | R/W | 0x0000_0000 | 5.2.4.3/5-8 |
| 0x0_003C | eLBC local access window 3 attribute register (LBLAWAR3) | R/W | 0x0000_0000 | 5.2.4.4/5-9 |
| 0x0_0040– 0x0_005C | Reserved | — | — | — |
| 0x0_0060 | PCI local access window 0 base address register (PCILAWBAR0) | R/W | 0x0000_0000 ³ | 5.2.4.5/5-10 |
| 0x0_0064 | PCI local access window 0 attribute register (PCILAWAR0) | R/W | 0x0000_0000 ⁴ | 5.2.4.6/5-11 |
| 0x0_0068 | PCI local access window 1 base address register (PCILAWBAR1) | R/W | 0x0000_0000 ⁵ | 5.2.4.5/5-10 |
| 0x0_006C | PCI local access window 1 attribute register (PCILAWAR1) | R/W | 0x0000_0000 | 5.2.4.6/5-11 |
| 0x0_0070– 0x0_007C | Reserved | — | — | — |
| 0x0_0080 | PCI Express 1 local access window base address register (PCIEXP1LAWBAR) | R/W | 0x0000_0000 | 5.2.4.7/5-12 |
| 0x0_0084 | PCI Express 1 local access window attribute register (PCIEXP1LAWAR) | R/W | 0x0000_0000 | 5.2.4.8/5-12 |
| 0x0_0088 | PCI Express 2 local access window base address register (PCIEXP2LAWBAR) | R/W | 0x0000_0000 | 5.2.4.9/5-13 |
| 0x0_008C | PCI Express 2 local access window attribute register (PCIEXP2LAWAR) | R/W | 0x0000_0000 | 5.2.4.10/5-14 |
| 0x0_00890– 0x0_009C | Reserved | — | — | — |
| 0x0_00A0 | DDR local access window 0 base address register (DDRLAWBAR0) | R/W | 0x0000_0000 ⁶ | 5.2.4.11/5-14 |

Table 5-4. Local Access Register Memory Map (continued)

| Local Access—Block Base Address 0x0_0000 | | | | |
|--|--|--------|--------------------------|---------------|
| Local Memory Offset (Hex) | Register | Access | Reset | Section/Page |
| 0x0_00A4 | DDR local access window 0 attribute register (DDRLAWAR0) | R/W | 0x0000_0000 ⁷ | 5.2.4.12/5-15 |
| 0x0_00A8 | DDR local access window 1 base address register (DDRLAWBAR1) | R/W | 0x0000_0000 | 5.2.4.11/5-14 |
| 0x0_00AC | DDR local access window 1 attribute register (DDRLAWAR1) | R/W | 0x0000_0000 | 5.2.4.12/5-15 |
| 0x0_00B0–0x0_00FC | Reserved | — | — | — |

- ¹ Depends on reset configuration word high values. See [Section 5.2.4.3.1, “LBLAWBAR0\[BASE_ADDR\] Reset Value,”](#) for details.
- ² Depends on reset configuration word high values. See [Section 5.2.4.4.1, “LBLAWAR0\[EN\] and LBLAWAR0\[SIZE\] Reset Value,”](#) for details.
- ³ Depends on reset configuration word high values. See [Section 5.2.4.5.1, “PCILAWBAR0\[BASE_ADDR\] Reset Value,”](#) for details.
- ⁴ Depends on reset configuration word high values. See [Section 5.2.4.11.1, “DDRLAWBAR0\[BASE_ADDR\] Reset Value,”](#) for details.
- ⁵ Depends on reset configuration word high values. See [Section 5.2.4.6.1, “PCILAWAR0\[EN\] and PCILAWAR0\[SIZE\] Reset Value,”](#) for details.
- ⁶ Depends on reset configuration word high values. See [Section 5.2.4.11.1, “DDRLAWBAR0\[BASE_ADDR\] Reset Value,”](#) for details.
- ⁷ Depends on reset configuration word high values. See [Section 5.2.4.12.1, “DDRLAWAR0\[EN\] and DDRLAWAR0\[SIZE\] Reset Value,”](#) for details.

5.2.4 Local Access Register Descriptions

5.2.4.1 Internal Memory Map Registers Base Address Register (IMMRBAR)

The IMMR window contains configuration, control, and status registers, as well as internal device memory arrays. The internal memory map occupies a 1-Mbyte region of memory space. Its location is programmable using the internal memory map register (IMMR). The default base address for the internal memory map register is 0xFF40_0000. Because IMMRBAR is at offset 0x0 from the beginning of the local access registers, IMMRBAR always points to itself.

5.2.4.1.1 Updating IMMRBAR

Updates to IMMRBAR that relocate the entire 1-Mbyte region of the internal memory block require special treatment. The effect of the update must be guaranteed to be visible by the mapping logic before an access to the new location is seen. To make sure this happens, the following guidelines should be followed:

- IMMRBAR should be updated during initial configuration of the device when only one host or controller has access to the device as follows:
 - If an external host on PCI is configuring the device, it should set IMMRBAR to the desired final location before the e300c3 core is released to boot.
 - If the core is initializing the device, it should set IMMRBAR to the desired final location before enabling other I/O devices to access the device.

- When the e300 core is writing to IMMRBAR, it should use the following sequence:
 - Read the current value of IMMRBAR using a load word instruction followed by an **isync**. This forces all accesses to configuration space to complete.
 - Write the new value to IMMRBAR.
 - Perform a load of an address that does not access configuration space or the on-chip SRAM, but has an address mapping already in effect (for example, boot ROM). Follow this load with an **isync**.
 - Read the contents of IMMRBAR from its new location, followed by another **isync**.

The IMMRBAR is shown in [Figure 5-2](#).




Figure 5-2. Internal Memory Map Registers' Base Address Register (IMMRBAR)

[Table 5-5](#) defines the bit fields of IMMRBAR.

Table 5-5. IMMRBAR Bit Settings

| Bits | Name | Description |
|-------|-----------|--|
| 0–11 | BASE_ADDR | Identifies the 12 most-significant address bits of the base of the 1-Mbyte internal memory window. |
| 12–31 | — | Reserved. Software must write all zeros. |

5.2.4.2 Alternate Configuration Base Address Register (ALTCBAR)

The alternate configuration base address register (ALTCBAR) is used to define the base address for an alternate 1-Mbyte region of configuration space to be used by the boot sequencer. By loading the proper boot sequencer command in the serial ROM, the base address in the ALTCBAR can be combined with the 20 bits of address offset supplied from the serial ROM to generate a 32-bit address. Thus, by configuring this register, the boot sequencer has access to the entire memory map, one 1-Mbyte block at a time. See [Section 20.4.5, “Boot Sequencer Mode,”](#) for more information.

NOTE

ALTCBAR is not considered a local access window on its own, so the boot sequencer must configure one of the other local access windows properly to reach the desired target peripherals.


Table 5-8 defines the reset value of LBLAWBAR0[BASE_ADDR].

Table 5-8. LBLAWBAR0[BASE_ADDR] Reset Value

| RCWHR[BMS] | BASE_ADDR Reset Value |
|------------|-----------------------|
| 0 | 0x00000 |
| 1 | 0xFF800 |

5.2.4.4 LBC Local Access Window *n* Attributes Registers (LBLAWAR0–LBLAWAR3)

The LBC local access window *n* attributes registers (LBLAWAR0–LBLAWAR3) are shown in Figure 5-5.



- 1 The LBLAWAR0[EN] reset value depends on the reset configuration word high values. See Section 5.2.4.4.1, “LBLAWAR0[EN] and LBLAWAR0[SIZE] Reset Value,” for a detailed description.
- 2 The LBLAWAR0[SIZE] reset value is always 0b010110, meaning an 8-Mbyte local access window. See Section 5.2.4.4.1, “LBLAWAR0[EN] and LBLAWAR0[SIZE] Reset Value,” for a detailed description.

Figure 5-5. LBC Local Access Window *n* Attributes Registers (LBLAWAR0–LBLAWAR3)

Table 5-9 defines the bit fields of LBLAWAR0–LBLAWAR3.

Table 5-9. LBLAWAR0–LBLAWAR3 Bit Settings

| Bits | Name | Description |
|-------|------|--|
| 0 | EN | 0 Local bus local access window <i>n</i> is disabled. 1 Local bus local access window <i>n</i> is enabled and other LBLAWAR0 and LBLAWBAR0 fields combine to identify an address range for this window. |
| 1–25 | — | Reserved. Write has no effect, read returns 0. |
| 26–31 | SIZE | Identifies the size of the window from the starting address. Window size is $2^{(SIZE+1)}$ bytes. 000000–001010 Reserved. Window is undefined. 001011 4 Kbytes 001100 8 Kbytes 001101 16 Kbytes $2^{(SIZE+1)}$ bytes 011110 2 Gbytes 011111–111111 Reserved. Window is undefined. |

5.2.4.4.1 LBLAWAR0[EN] and LBLAWAR0[SIZE] Reset Value

The core may use a local bus peripheral device to fetch its boot vector. For this purpose an 8-Mbyte ($2^{(22+1)}$) local access window is defined by the LBLAWBAR0[SIZE] reset value, and LBLAWAR0 is enabled according to the value set in the reset configuration word high ROMLOC field.


Table 5-10 defines the reset value for LBLAWAR0[EN].

Table 5-10. LBLAWAR0[EN] Reset Value

| RCWHR[RLEXT]/ RCWHR[ROMLOC] | LBLAWAR0[EN] Reset Value | Description |
|--------------------------------|-----------------------------|--|
| 00 / 000–100 | 0 | e300c3 core boot not performed from a local bus device. |
| 00 / 101–111 | 1 | e300c3 core boot performed from a local bus device. Local bus 8-Mbyte ($2^{(22+1)}$) local access window is enabled. |

5.2.4.5 PCI Local Access Window *n* Base Address Register (PCILAWBAR0–PCILAWBAR1)

The PCI local access window *n* base address registers (PCILAWBAR0–PCILAWBAR1) are shown in Figure 5-6.



¹ The reset value of PCILAWBAR0[BASE_ADDR] depends on the reset configuration word high values. See Section 5.2.4.5.1, “PCILAWBAR0[BASE_ADDR] Reset Value,” for a detailed description.

Figure 5-6. PCI Local Access Window *n* Base Address Registers (PCILAWBAR0–PCILAWBAR1)

Table 5-11 defines the bit fields of PCILAWBAR0–PCILAWBAR1.

Table 5-11. PCILAWBAR0–PCILAWBAR1 Bit Settings

| Bits | Name | Description |
|-------|-----------|--|
| 0–19 | BASE_ADDR | Identifies the 20 most-significant address bits of the base of local access window <i>n</i> . The specified base address should be aligned to the window size, as defined by PCILAWAR <i>n</i> [SIZE]. |
| 20–31 | — | Reserved. Write has no effect, read returns 0. |

5.2.4.5.1 PCILAWBAR0[BASE_ADDR] Reset Value

The core may use a PCI peripheral device to fetch its boot vector. For this purpose, the PCILAWBAR0[BASE_ADDR] reset value is set according to the value set in the reset configuration word high BMS field.

Table 5-12 defines the reset value of PCILAWBAR0[BASE_ADDR].

Table 5-12. PCILAWBAR0[BASE_ADDR] Reset Value

| RCWHR[BMS] | PCILAWBAR0[BASE_ADDR] Reset Value |
|------------|--------------------------------------|
| 0 | 0x00000 |
| 1 | 0xFF800 |

Table 5-14. PCILAWAR0[EN] Reset Value (continued)

| RCWHR[RLEXT]/RCWHR[ROMLOC] | PCILAWR0[EN] Reset Value | Description |
|----------------------------|--------------------------|--|
| 01–11 / 000–111 | 0 | e300c3 core boot not performed from a PCI device. |
| 00 / 001 | 1 | e300c3 core boot performed from a PCI device. PCI 8-Mbyte ($2^{(22+1)}$) local access window is enabled. |

5.2.4.7 PCI Express 1 Local Access Window Base Address Register (PCIEXP1LAWBAR)

The PCI Express 1 local access window base address register is shown in [Figure 5-8](#).




Figure 5-8. PCI Express 1 Local Access Window Base Address Register (PCIEXP1LAWBAR)

[Table 5-15](#) defines the bit fields of PCIEXP1LAWBAR.

Table 5-15. PCIEXP1LAWBAR Bit Settings

| Bits | Name | Description |
|-------|-----------|--|
| 0–19 | BASE_ADDR | Identifies the 20 most-significant address bits of the base of local access window. The specified base address should be aligned to the window size, as defined by PCIEXP1LAWAR[SIZE]. |
| 20–31 | — | Reserved. Write has no effect, read returns 0. |

5.2.4.8 PCI Express 1 Local Access Window Attributes Registers (PCIEXP1LAWAR)

The PCI Express 1 local access window attributes register is shown in [Figure 5-9](#).




Figure 5-9. PCI Express 1 Local Access Window Attributes Register (PCIEXP1LAWAR)

Table 5-21 defines the bit fields of DDRLAWAR0–DDRLAWAR1.

Table 5-21. DDRLAWAR0–DDRLAWAR1 Bit Settings

| Bits | Name | Description |
|-------|------|--|
| 0 | EN | 0 The DDR local access window <i>n</i> is disabled. 1 The DDR local access window <i>n</i> is enabled and other DDRLAWAR n and DDRLAWBAR n fields combine to identify an address range for this window. |
| 1–25 | — | Reserved. Write has no effect, read returns 0. |
| 26–31 | SIZE | Identifies the size of the window from the starting address. Window size is $2^{(SIZE+1)}$ bytes. 000000–001010 Reserved. Window is undefined. 001011 4 Kbytes 001100 8 Kbytes 001101 16 Kbytes $2^{(SIZE+1)}$ bytes 011110 2 Gbytes 011111–111111 Reserved. Window is undefined. |

5.2.4.12.1 DDRLAWAR0[EN] and DDRLAWAR0[SIZE] Reset Value

The core may use a DDR SDRAM device to fetch its boot vector. For this purpose an 8-Mbyte ($2^{(22+1)}$) local access window is defined by DDRLAWBAR0[SIZE] reset value, and DDRLAWAR0 is enabled according to the value set in the reset configuration word high ROMLOC field.

Table 5-22 defines the reset value DDRLAWAR0[EN] and DDRLAWAR0[SIZE].

Table 5-22. DDRLAWAR0[EN] Reset Value

| RCWHR[RLEXT]/RCWHR[ROMLOC] | DDRLAWAR0[EN] Reset Value | Description |
|----------------------------|---------------------------|--|
| 00 / 000 | 1 | e300c3 core boot performed from a DDR SDRAM device. DDR 8-Mbyte ($2^{(22+1)}$) local access window is enabled. |
| Else | 0 | e300c3 core boot not performed from a DDR SDRAM device. |

5.2.5 Precedence of Local Access Windows

If two local access windows overlap, the lower numbered window takes precedence (see Table 5-1 for window numbers). For instance, if two windows are set up as shown in Table 5-23, local access window 1 governs the mapping of the 1-Mbyte region from 0x7FF0_0000 to 0x7FFF_FFF, even though the window described in local access window 7 also encompasses that memory region.

Table 5-23. Overlapping Local Access Windows

| Window | Base Address | Size | Target Interface |
|--------|--------------|----------|------------------|
| 1 | 0x7FF0_0000 | 1 Mbyte | Local bus |
| 7 | 0x0000_0000 | 2 Gbytes | DDR SDRAM |

5.2.6 Configuring Local Access Windows

After a local access window is enabled, it should not be modified while any device in the system may be using the window. Accordingly, a new window should not be used until the effect of the write to the window is visible to all blocks that use the window. This can be guaranteed by completing a read of the last local access window configuration register before enabling any other devices to use the window. For instance, if local bus local access windows 1–3 are being configured in order during the initialization process, the last write (to LBLAWAR3) should be followed by a read of LBLAWAR3 before any devices try to use any of these windows. If the configuration is being done by the local e300c3 core, the read of LBLAWAR3 should be followed by an **isync** instruction.

5.2.7 Distinguishing Local Access Windows from Other Mapping Functions

It is important to distinguish between the mapping function performed by the local access windows and the additional mapping functions that happen at the target interface. The local access windows define how a transaction is routed through the device internal interconnects from the transaction's source to its target. Once the transaction has arrived at its target interface, that interface controller may perform additional mapping. For instance, the DDR SDRAM controller has chip select registers that map a memory request to a particular external device. The local bus controller has base registers that perform a similar function. The PCI interface has outbound address translation units that map the local address into an external address space.

These other mapping functions are configured by programming the configuration, control, and status registers of the individual interfaces. Note that there is no need to have a one-to-one correspondence between local access windows and chip select regions or outbound windows. A single local access window can be further decoded to any number of chip selects or to any number or outbound windows at the target interface.

5.2.8 Outbound Address Translation and Mapping Windows

Outbound address translation and mapping refers to the translation of addresses from the local 32-bit address space to the external address space and attributes of a particular I/O interface. On this device, the PCI block has an outbound address translation unit.

The PCI controller has six outbound windows plus a default window. See [Section 4.5, “Memory Map/Register Definitions,”](#) for a detailed description of the PCI outbound windows.

5.2.9 Inbound Address Translation and Mapping Windows

Inbound address translation and mapping refers to the translation of an address from the external address space of an I/O interface (such as PCI address space) to the local address space understood by the internal interfaces of this processor. It also refers to the mapping of transactions to a particular target interface and the assignment of transaction attributes. The PCI controller has inbound address translation unit.

5.2.9.1 PCI Inbound Windows

The PCI controller has three general inbound windows plus a dedicated window for memory-mapped configuration accesses (PIMMR). These windows have a one-to-one correspondence with the base address registers in the PCI programming model. Updating one automatically updates the other. There is no default inbound window; if a PCI address does not match one of the inbound windows, this processor does not respond with an assertion of `PCI_DEVSEL`. See [Section 13.4.6, “PCI Inbound Address Translation,”](#) for a detailed description of the PCI inbound windows.

5.2.10 Internal Memory Map

All of the memory-mapped configuration, control, and status registers in the device are contained within a 1-Mbyte address region, referred as the IMMR. To allow for flexibility, the internal memory map block can be relocated in the local address space. The local address map location of this register block is controlled by the internal memory map registers' base address register (IMMRBAR); see [Section 5.2.4.1, “Internal Memory Map Registers Base Address Register \(IMMRBAR\).”](#) The default value for the IMMRBAR is `0xFF40_0000`.

NOTE

The internal memory map window is always the highest priority local access window.

5.2.11 Accessing Internal Memory from External Masters

In addition to being accessible by the e300 processor, the IMMR memory window is accessible from external interfaces. This allows external masters on the I/O ports to configure the device.

External masters do not need to know the location of the IMMR memory in the local address map. Rather, they access this region of the local memory map through a window defined by a register in the interface's programming model that is accessible to the external master from its external memory map.

The PCI base address for accessing the local IMMR memory is selectable through the PCI internal memory map register (PIMMR), at offset `0x10`, described in [Section 13.3.2.12, “PCI Inbound Base Address Registers \(PIBARn\).”](#) When the device is a PCI agent, an external PCI master sets this register by running a PCI configuration cycle. Subsequent memory accesses by a PCI master to the PCI address range indicated by PIMMR are translated to the local address indicated by the current setting of IMMRBAR.

5.3 System Configuration

The following sections describe some general information and configuration options that affect system behavior and performance.

5.3.1 System Configuration Register Memory Map

Table 5-24 shows the memory map for the system configuration registers.

Table 5-24. System Configuration Register Memory Map

| Local Memory Offset (Hex) | Register | Access | Reset | Section/Page |
|---|--|--------|----------------------------|----------------|
| System Configuration—Block Base Address 0x0_0000 | | | | |
| 0x00100 | System general purpose register low (SGPRL) | R/W | 0x0000_0000 | 5.3.2.1/5-19 |
| 0x00104 | System general purpose register high (SGPRH) | R/W | 0x0000_0000 | 5.3.2.2/5-20 |
| 0x00108 | System part and revision ID register (SPRIDR) | R | 0x80B4_0010 | 5.3.2.3/5-20 |
| 0x0010C | Reserved | — | — | — |
| 0x00110 | System priority configuration register (SPCR) | R/W | 0x0000_0000 | 5.3.2.4/5-21 |
| 0x00114 | System I/O configuration register low (SICRL) | R/W | 0xF00F00_n0n0 ¹ | 5.3.2.5/5-23 |
| 0x00118 | System I/O configuration register high (SICRH) | R/W | 0xFFFF_C0n | 5.3.2.6/5-26 |
| 0x0011C–0x001FC | Reserved | — | — | — |
| 0x00128 | DDR control driver register (DDRCDR) | R/W | 0x7B84_0001 | 5.3.2.8/5-29 |
| 0x0012C | DDR debug status register (DDRDSR) | R | 0x3B80_0000 | 5.3.2.9/5-31 |
| 0x00140 | PCI Express control register 1 (PECR1) | R/W | 0x0000_0000 | 5.3.2.10/55-31 |
| 0x00144 | PCI Express control register 2 (PECR2) | R/W | 0x0000_0000 | 5.3.2.10/55-31 |
| 0x00148–0x001FC | Reserved | — | — | — |

¹ Depends on the reset configuration word high configuration values.

5.3.2 System Configuration Registers

5.3.2.1 System General Purpose Register Low (SGPRL)

The system general purpose register low (SGPRL), shown in Figure 5-14, can be used by software for any purpose. The values set in this register have no effect on the internal hardware.




Figure 5-14. System General Purpose Register Low (SGPRL)

Table 5-25 defines the bit fields of SGPRL.

Table 5-25. SGPRL Bit Settings

| Bits | Name | Description |
|------|------|-----------------|
| 0–31 | GP | General purpose |

5.3.2.2 System General Purpose Register High (SGPRH)

The system general purpose register high (SGPRH), shown in Figure 5-15, can be used by software for any purpose. The values set in this register have no effect on the internal hardware.




Figure 5-15. System General Purpose Register High (SGPRH)

Table 5-26 defines the bit fields of SGPRH.

Table 5-26. SGPRH Bit Settings

| Bits | Name | Description |
|------|------|-----------------|
| 0–31 | GP | General purpose |

5.3.2.3 System Part and Revision ID Register (SPRIDR)

SPRIDR, shown in Figure 5-16, provides information about the device and revision numbers.




Figure 5-16. System Part and Revision ID Register (SPRIDR)

Table 5-27 defines the bit fields of SPRIDR.

Table 5-27. SPRIDR Bit Settings

| Bits | Name | Description |
|-------|--------|---|
| 0–15 | PARTID | Part identification. This read-only field is mask-programmed with a code corresponding to the device number. It is intended to help factory test and user code that is sensitive to device changes. The device number changes according to manufacturing considerations. See Table 5-28 for values of this field. |
| 16–31 | REVID | Revision identification. This read-only field is mask-programmed with a code corresponding to the revision number of the part defined in PARTID field. It is intended to help factory test and user code that is sensitive to device changes. The mask number is programmed in a commonly changed layer, and changes with each mask set change. See Table 5-29 for values of this field. |

5.3.2.3.1 SPRIDR[PARTID] Coding

Table 5-28 defines the reset values of SPRIDR[PARTID].

Table 5-28. PARTID Coding

| PARTID | Device Name | Package Type |
|--------|-------------|--------------|
| 0x80B4 | MPC8315E | PBGA |
| 0x80B5 | MPC8315 | PBGA |
| 0x80B6 | MPC8314E | PBGA |
| 0x80B7 | MPC8314 | PBGA |

Table 5-29 defines the reset values of SPRIDR[REVID].

Table 5-29. REVID Coding

| REVID | Device Revision |
|--------|-----------------|
| 0x0010 | 1.0 |

5.3.2.4 System Priority and Configuration Register (SPCR)

The system priority and configuration register (SPCR), shown in Figure 5-17, controls the priority of requests for transactions on the internal system bus. This priority is considered by the system arbiter

System Configuration

whenever an internal unit requests mastership of the coherent system bus (CSB). The SPCR also includes some other control functions.




Figure 5-17. System Priority Configuration Register (SPCR)

Table 5-30 defines the bit fields of SPCR.

Table 5-30. SPCR Bit Settings

| Bits | Name | Description |
|-------|--------|---|
| 0–2 | — | Reserved. Should be cleared. |
| 3 | PCIHPE | PCI highest priority enable. If this bit is set, the PCI bridge is permitted to request the coherent system bus (CSB) with highest priority, regardless of SPCR[PCIPR] value, when it needs to complete a posted write transaction from an external PCI master. To follow PCI ordering rules specifications, the PCI bridge must flush any outstanding write transactions before it can start a new read transaction. Setting this bit allows faster flushing of the outstanding write transactions coming from the PCI bus onto the CSB and to the device targets, such as DDR SDRAM and local bus memories. |
| 4–5 | — | Reserved. Should be cleared. |
| 6–7 | PCIPR | PCI bridge CSB request priority. The level of priority can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) Note: DMA has the same priority as PCI. |
| 8 | OPT | Optimize. Setting this bit may enhance the performance of transactions issued to the internal coherent system bus (CSB) by the security engine (SEC) and the USB controller. Performance is enhanced by reading more bytes on the bus than actually needed by the master in the case that this is more efficient. The user may set this bit only if it is known that USB transactions sent to the internal CSB are not accessing devices in which speculative reads may change the state of the device (for example, FIFOs in which reading a byte may advance some internal counter). 0 No performance enhancement. 1 Performance enhancement by speculative reading is enabled. |
| 9 | TBEN | e300c3 core time base unit enable 0 Time base unit is disabled. 1 Time base unit is enabled. |
| 10–11 | COREPR | e300c3 core CSB request priority. The priority level for the core in accessing the CSB can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) |

Table 5-30. SPCR Bit Settings (continued)

| Bits | Name | Description |
|-------|---------|---|
| 12–17 | — | Reserved. Should be cleared. |
| 12–21 | — | Reserved. Should be cleared. |
| 18–19 | TSECDP | eTSEC data priority. Selects the CSB request priority driven by eTSEC1 and eTSEC2 when they require to transfer data on this bus. The level of priority can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) |
| 20–21 | TSECBDP | eTSEC buffer descriptor priority. Selects the CSB request priority driven by eTSEC1 and eTSEC2 when they require to transfer a buffer descriptor (BD) on this bus. The level of priority can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) |
| 22–23 | TSECEP | TSEC emergency priority. Selects the CSB request priority driven by eTSEC1 and eTSEC2 when an emergency condition occurs. The level of priority can be chosen from 4 possible levels. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) |
| 24–31 | — | Reserved, should be cleared. |

5.3.2.5 System I/O Configuration Register Low (SICRL)

The system I/O configuration register low (SICRL) controls the multiplexing of some of the device I/O pins. Each bit or set of bits in the SICRL selects which function is used by a certain group of the device pins.

The reset value of this register depends on TSEC1M fields setting in the reset configuration word high. This is used to avoid contention in systems not using TBI or RTBI modes. In these systems, the device pins are not driven during and after reset. The function of these pins can be changed by writing to this register during system initialization. TSOBI1 reset value also depends on the TSEC1M field settings in the reset configuration word high in order to select the correct output buffer impedance for full or reduced TSEC pin mode.

Figure 5-18 shows SICRL.

Offset 0x00114

Access: Read/Write

| | | | | | | | | |
|-------|--------------------------|------------------------------|----------|----|------------------------|----|--------------------------|--------------------------|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| R | DMA_CH0 | | SPI | | UART | | $\overline{\text{IRQ4}}$ | $\overline{\text{IRQ5}}$ |
| W | DMA_CH0 | | SPI | | UART | | $\overline{\text{IRQ4}}$ | $\overline{\text{IRQ5}}$ |
| Reset | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 |
| | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
| R | $\overline{\text{IRQ5}}$ | $\overline{\text{IRQ[6-7]}}$ | IIC1 | | TDM | | TDM_SHARED | |
| W | $\overline{\text{IRQ5}}$ | $\overline{\text{IRQ[6-7]}}$ | IIC1 | | TDM | | TDM_SHARED | |
| Reset | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 |
| | 16 | 17 | 18 | 19 | 20 | 23 | | |
| R | PCI_A | | ELBC_A | | — | | | |
| W | PCI_A | | ELBC_A | | — | | | |
| Reset | All zeros | | | | | | | |
| | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 |
| R | ETSEC1_A ¹ | | ETSEC1_B | | ETSEC1_C ¹ | | — | TSEXPOBI |
| W | ETSEC1_A ¹ | | ETSEC1_B | | ETSEC1_C ¹ | | — | TSEXPOBI |
| Reset | depends on RCWH[16-18] | | 0 | 0 | depends on RCWH[19-21] | | 0 | 0 |

¹ The SICRL[ETSEC1_A] depends on the value of RCWH[16-18] and SICRL[ETSEC1_C] reset value depends on the value of RCWH[19-21] which is loaded into the reset configuration word.

Figure 5-18. System I/O Configuration Register Low (SICRL)

Table 5-31 defines the bit fields of SICRL. Each Pin Function column lists the name of the multi-function pin used in this option. Some groups have only two options (shown as Pin Function 0 and Pin Function 1) and therefore, only one control bit. In this case they can only have a value of 0b0 or 0b1. Other groups may have four options (shown as Pin Function 0, Pin Function 1, Pin Function 2, and Pin Function 3) and therefore, two control bits. In this case they can have a value of 0b00, 0b01, 0b10, or 0b11. Use the notations ‘0bN’ or ‘0bNN’ according to whether a group has one or two control bits, respectively.

Table 5-31. SICRL Bit Settings

| SICRL[Bits] Value | | 0b00 | 0b01 | 0b10 | 0b11 | Reset |
|-------------------|---------|--------------------------------|----------------|----------------|----------------|-------|
| Bits | Group | Pin Function 0 | Pin Function 1 | Pin Function 2 | Pin Function 3 | |
| 0-1 | DMA_CH0 | $\overline{\text{DMA_DREQ0}}$ | — | — | GPIO_12 | 11 |
| | | $\overline{\text{DMA_DACK0}}$ | — | — | GPIO_13 | |
| | | DMA_DONE0 | — | — | GPIO_14 | |
| 2-3 | SPI | SPI_MOSI | — | — | GPIO_15 | 11 |
| | | SPI_MISO | — | — | GPIO_16 | |
| | | SPISEL | — | — | GPIO_17 | |

Table 5-31. SICRL Bit Settings (continued)

| SICRL[Bits] Value | | 0b00 | 0b01 | 0b10 | 0b11 | Reset |
|-------------------|------------------------------|----------------|-------------------------------------|----------------|----------------|-------|
| Bits | Group | Pin Function 0 | Pin Function 1 | Pin Function 2 | Pin Function 3 | |
| 4-5 | UART | UART_SOUT1 | MSRCID0 (DDRID) | LSRCID0 | — | 00 |
| | | UART_SIN1 | MSRCID1 (DDRID) | LSRCID1 | — | |
| | | UART_CTS1 | MSRCID2 (DDRID) | LSRCID2 | — | |
| | | UART_RTS1 | MSRCID3 (DDRID) | LSRCID3 | — | |
| | | UART_SOUT2 | MSRCID4 (DDRID) | LSRCID4 | — | |
| | | UART_SIN2 | MDVAL | LDVAL | — | |
| | | UART_CTS2 | — | — | — | |
| | | UART_RTS2 | — | — | — | |
| 6 | $\overline{\text{IRQ4}}$ | IRQ4 | — | — | — | 0 |
| 7-8 | $\overline{\text{IRQ5}}$ | IRQ5 | $\overline{\text{CORE_SRESET_IN}}$ | — | — | 00 |
| 9 | $\overline{\text{IRQ}}[6:7]$ | IRQ6 | CKSTOP_OUT | — | — | 0 |
| | | IRQ7 | CKSTOP_IN | — | — | |
| 10-11 | IIC1 | IIC1_SDA | CKSTOP_OUT | — | — | 00 |
| | | IIC1_SCL | CKSTOP_IN | — | — | |
| 12-13 | TDM | TDM_RD | — | — | GPIO_20 | 11 |
| | | TDM_TCK | — | — | GPIO_21 | |
| | | TDM_TFS | — | — | GPIO_22 | |
| | | TDM_TD | — | — | GPIO_23 | |
| 14-15 | TDM Shared | TDM_RCK | — | — | GPIO_18 | 11 |
| | | TDM_RFS | — | — | GPIO_19 | |
| 16-17 | PCI_A | PCI_INTA | — | — | — | 00 |
| | | PCI_RESET_OUT | — | — | — | |
| | | PCI_AD[0:31] | — | — | — | |
| | | PCI_C_BE[0:3] | — | — | — | |
| | | PCI_PAR | — | — | — | |
| | | PCI_FRAME | — | — | — | |
| | | PCI_TRDY | — | — | — | |
| | | PCI_IRDY | — | — | — | |
| | | PCI_STOP | — | — | — | |
| | | PCI_DEVSEL | — | — | — | |
| | | PCI_IDSEL | — | — | — | |

Table 5-31. SICRL Bit Settings (continued)

| SICRL[Bits] Value | | 0b00 | 0b01 | 0b10 | 0b11 | Reset |
|-------------------|-----------------------|---|----------------|-----------------|----------------|---|
| Bits | Group | Pin Function 0 | Pin Function 1 | Pin Function 2 | Pin Function 3 | |
| 18–19 | ELBC_A | LAD[0:15] | — | — | — | 00 |
| | | LA[16:25] | — | — | — | |
| 20–23 | — | — | — | — | — | 0000 |
| 24–25 | ETSEC1_A | TSEC1_COL | USBDR_TXDRXD0 | | GPIO_24 | if ETSEC1 is in RTBI mode, then chosen by RCWH [TSEC1M] |
| | | TSEC1_CRS | USBDR_TXDRXD1 | | GPIO_25 | |
| | | TSEC1_GTX_CLK | USBDR_TXDRXD2 | | | |
| | | TSEC1_RX_CLK | USBDR_TXDRXD3 | | | |
| | | TSEC1_RX_DV | USBDR_TXDRXD4 | | | |
| 26 | ETSEC1_B | TSEC1_TX_CLK | USBDR_CLK | | | |
| 27 | ETSEC1_B ¹ | TSEC1_TX_CLK | USBDR_CLK | | | 00 |
| | | TSEC1_RXD3 | USBDR_TXDRXD5 | TSEC_TMR_CLK | | |
| | | TSEC1_RXD2 | USBDR_TXDRXD6 | TSEC_TMR_TRIG1 | | |
| | | TSEC1_RXD1 | USBDR_TXDRXD7 | TSEC_TMR_TRIG2 | | |
| | | TSEC1_RXD0 | USBDR_NXT | | | |
| | | TSEC1_RX_ER | USBDR_DIR | | | |
| 28–29 | ETSEC1_C | TSEC1_TXD3 | USBDR_STP | TSEC_TMR_GCLK | GPIO_28 | 00 if RTBI, else 11 |
| | | TSEC1_TXD2 | | TSEC_TMR_PP1 | GPIO_29 | |
| | | TSEC1_TXD1 | | TSEC_TMR_PP2 | GPIO_30 | |
| | | TSEC1_TXD0 | | TSEC_TMR_PP3 | | |
| | | TSEC1_TX_EN | | TSEC_TMR_ALARM1 | GPIO_31 | |
| | | TSEC1_TX_ER | | TSEC_TMR_ALARM2 | | |
| 30 | — | | | | | 0 |
| 31 | TSEXPOBI | TSEC output buffer impedance for col, rx_er, tx_clk & tx_er pads. This bit controls the output buffer impedance of the above mentioned pads of TSEC1 & TSEC2 output pads. 0 Output buffer is set for 40 Ω, 3.3 V. 1 Output buffer is set for 40 Ω, 2.5 V. | | | | 0 |

¹ TSEC1_TX_CLK is selected when SICRL[27] is logic 0 irrespective of SICRL[26].

5.3.2.6 System I/O Configuration Register High (SICRH)

The system I/O configuration register high, shown in [Figure 5-19](#), controls the multiplexing of the rest of the device I/O pins. Each bit or set of bits in this register select which function is used by a certain group of the device pins.

Offset 0x00118

Access: Read/Write




Figure 5-19. System I/O Configuration Register High (SICRH)

Table 5-32 defines the bit fields of SICRH. Each Pin Function column lists the name of the multi-function pin used in this option. Some groups have only two options (shown as Pin Function 0 and Pin Function 1) and therefore, only one control bit. In this case they can have the value of 0b0 or 0b1 only. Other groups may have three options (shown as Pin Function 0, Pin Function 1, and Pin Function 2) and therefore, two control bits. In this case they can have the value of 0b00, 0b01, or 0b10. A value of 0b11 selects GPIO mode of the appropriate pin. Use the 0bN or 0bNN notations according to a group having one or two control bits, respectively.

Note that bits 30 and 31 (TSOBI1 and TSOBI2), which control TSEC output buffer impedance, are described in Table 5-33.

Table 5-32. SICRH Bit Settings

| SICRH[Bits] Value | | 0b00 | 0b01 | 0b10 | 0b11 | Reset Value |
|-------------------|--------|--------------------------------|---|----------------|----------------|-------------|
| Bits | Group | Pin Function 0 | Pin Function 1 | Pin Function 2 | Pin Function 3 | |
| 0–1 | GPIO_0 | $\overline{\text{DMA_DREQ1}}$ | $\overline{\text{GTM1_TOUT1}}$ | — | GPIO_0 | 11 |
| 2–3 | GPIO_1 | $\overline{\text{DMA_DACK1}}$ | GTM1_TIN2/GTM2_TIN1 | — | GPIO_1 | 11 |
| 4–5 | GPIO_2 | DMA_DONE1 | $\overline{\text{GTM1_TGATE2/}}$ $\overline{\text{GTM2_TGATE1}}$ | — | GPIO_2 | 11 |
| 6–7 | GPIO_3 | — | GTM1_TIN3/GTM2_TIN4 | — | GPIO_3 | 11 |
| 8–9 | GPIO_4 | — | $\overline{\text{GTM1_TGATE3/}}$ $\overline{\text{GTM2_TGATE4}}$ | — | GPIO_4 | 11 |

Table 5-32. SICRH Bit Settings (continued)

| SICRH[Bits] Value | | 0b00 | 0b01 | 0b10 | 0b11 | Reset Value |
|-------------------|-----------------------------------|---|--|---------------------------------|----------------|-------------|
| Bits | Group | Pin Function 0 | Pin Function 1 | Pin Function 2 | Pin Function 3 | |
| 10–11 | GPIO_5 | — | $\overline{\text{GTM1_TOUT3}}$ | $\overline{\text{GTM2_TOUT1}}$ | GPIO_5 | 11 |
| 12–13 | GPIO_6 | — | GTM1_TIN4/GTM2_TIN3 | — | GPIO_6 | 11 |
| 14–15 | GPIO_7 | — | $\overline{\text{GTM1_TGATE4/}}/\overline{\text{GTM2_TGATE3}}$ | — | GPIO_7 | 11 |
| 16–17 | GPIO_8 | USBDR_DRIVE_VBUS | GTM1_TIN1/GTM2_TIN2 | — | GPIO_8 | 11 |
| 18–19 | GPIO_9 | USBDR_PWRFAULT | $\overline{\text{GTM1_TGATE1/}}/\overline{\text{GTM2_TGATE2}}$ | — | GPIO_9 | 11 |
| 20–21 | GPIO_10 | USBDR_PCTL0 | $\overline{\text{GTM1_TOUT2}}$ | $\overline{\text{GTM2_TOUT1}}$ | GPIO_10 | 11 |
| 22–23 | GPIO_11 | USBDR_PCTL1 | $\overline{\text{GTM1_TOUT4}}$ | $\overline{\text{GTM2_TOUT3}}$ | GPIO_11 | 11 |
| 24–25 | ETSEC2_A | TSEC2_COL | — | — | GPIO_26 | 11 |
| | | TSEC2_CRS | — | — | GPIO_27 | |
| 26 | — | | | | | 0 |
| 27 | TSCOMOB1 | See Table 5-33 for description and reset value. | | | | |
| 28 | GTX_CLK1_DELAY/ USBDR_TX_DRXD2 | | | | | |
| 29 | GTX_CLK2_DELAY | | | | | |
| 30 | TSOBI1 | | | | | |
| 31 | TSOBI2 | | | | | |

Table 5-33. SICRH[27–31] Bit Settings

| Bits | Name | Description | Reset Value |
|------|-----------------------------------|---|-------------|
| 27 | TSCOMOB1 | Input buffer configuration for common I/O's between TSEC1 and TSEC2: GTX_CLK125, MDC, MDIO pads. 0 3.3 V input signal 1 2.5 V input signal | 0 |
| 28 | GTX_CLK1_DELAY/ USBDR_TX_DRXD2 | TSEC1_GTX_CLK I/O delay configuration. This bit controls the output buffer hold time of the TSEC1_GTX_CLK signal. The reset value of this bit depends on the reset configuration word high TSEC1M field (~RCWH[16]). 0 Output buffer delay is 00, no delay unit is added. 1 Output buffer delay is 10, one delay unit is added. | 0 |
| 29 | GTX_CLK2_DELAY | TSEC2_GTX_CLK I/O delay configuration. This bit controls the output buffer hold time of the TSEC2_GTX_CLK signal. The reset value of this bit depends on the reset configuration word high TSEC2M field (~RCWH[18]). 0 Output buffer delay is 00, no delay unit is added. 1 Output buffer delay is 10, one delay unit is added. | 0 |

Table 5-33. SICRH[27–31] Bit Settings (continued)

| Bits | Name | Description | Reset Value |
|------|--------|---|--------------------------------------|
| 30 | TSOBI1 | TSEC1 output buffer impedance. This bit controls the output buffer impedance of the TSEC1 output signals, used for reduced pin mode interfaces (RTBI/RGMII). The output buffer impedance should be correlated to the voltage supplied to the TSEC1 I/O pins (LVDD1). For non-eTSEC mode of operation, this bit must be cleared. 0 Output buffer is set for 40 Ω, 3.3 V. 1 Output buffer is set for 40 Ω, 2.5 V. | 0 Else 1 RGMII, RTBI ¹ |
| 31 | TSOBI2 | TSEC2 output buffer impedance. This bit controls the output buffer impedance of TSEC2 output signals, used for reduced pin mode interfaces (RTBI/RGMII). The output buffer impedance should be correlated to the voltage supplied to the TSEC2 I/O pins (LVDD2). 0 Output buffer is set for 40 Ω, 3.3 V. 1 Output buffer is set for 40 Ω, 2.5 V. | 0 Else 1 RGMII, RTBI |

¹ If RCWH[ETSEC1M] is RGMII/RTBI then the reset value is 1, otherwise it is 0.

5.3.2.7 Debug Configuration

Debug information may be driven on the device pins. This information can identify the internal source of a transaction that reached the DDR SDRAM or local bus interfaces. The device can be configured to drive the MSRCID[0:4] and MDVAL, or LSRCID[0:4] and LDVAL signals, respectively on other device pins. The coding of the source ID debug information is the same as the coding of the MSTR_ID field in the AEATR register of the arbiter (See [Section 6.2.7, “Arbiter Event Attributes Register \(AEATR\)”](#)).

5.3.2.7.1 DDR Debug Configuration

The DDR debug configuration enables a DDR memory controller to enter debug mode in which the DDR SDRAM source ID field and data valid strobe are driven onto one of two optional sets of pins:

- UART pins. UART operation is disabled, and any signals driven by UART devices must be electrically disconnected from the UART I/O pins. Set SICRL[4–5] to 0b01 to select this mode.

5.3.2.7.2 Local Bus Debug Configuration

The local bus debug configuration enables a LBC debug mode in which the SDRAM source ID field and data valid strobe for LBC memory accesses are driven onto UART pins. UART operation must be disabled, and any signals driven by UART devices must be electrically disconnected from the UART I/O pins in this case. Set SICRL[4–5] to 0b10 to select this mode.

5.3.2.8 DDR Control Driver Register (DDRCDR)

The DDR control driver register (DDRCDR) contains bits that allow control over the driver of the DDR SDRAM controller.

5.3.2.9 DDR Debug Status Register (DDRDSR)

Figure 5-21 contains the debug status bits from the DDR SDRAM controller.




Figure 5-21. DDR Debug Status Register (DDRDSR)

Table 5-35 shows the bit settings of the DDRDSR.

Table 5-35. DDRDSR Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–1 | — | Reserved |
| 2–5 | PZ | Current setting of PFET driver impedance 0000 Half strength—highest Z 1000 Higher Z than nominal 1100 Nominal impedance setting 1110 Lower Z than nominal 1111 Much lower Z than nominal |
| 6–9 | NZ | Current setting of NFET driver impedance 0000 Half strength—highest Z 1000 Higher Z than nominal 1100 Nominal impedance setting 1110 Lower Z than nominal 1111 Much lower Z than nominal |
| 10–31 | — | Reserved |

5.3.2.10 PCI Express Control Registers (PECR1 and PECR2)

The PCI Express control registers can be used to control various settings that affect the system response to PCI Express controller DMA operations or PIO inbound operation. Those registers also control soft reset assertion and negation to various logic parts of the PCI Express controllers. Note that PECR1 programming controls the behavior of PCI Express controller for port 1, and PECR2 controls the behavior of PCI Express controller for port 2. PECR1 is located at offset 0x140 and PECR2 is located at offset 0x144.

Table 5-36. PECR Field Description (continued)

| Bits | Name | Description |
|-------|---------|--|
| 28–29 | PRI_DES | DMA descriptor priority. This field will be used to present priority level for CSB arbitration for PCI Express controller's DMA requests. Bits 28–29 will be used when the request belongs to descriptor fetch or update. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) |
| 30–31 | PRI_PIO | PIO priority. This field will be used to present priority level for CSB arbitration for PCI Express controller's PIO inbound requests. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) |

5.4 Software Watchdog Timer (WDT)

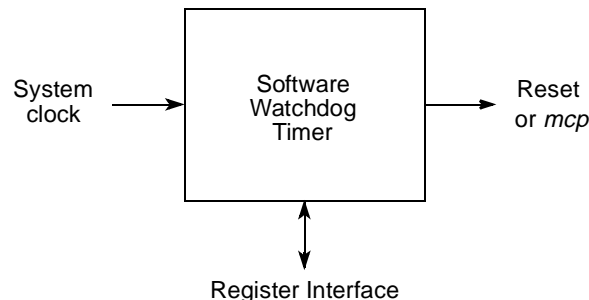
The following sections describe the theory of operation of the software watchdog timer (WDT) in the device, including a definition of the external signals and the functions they serve. Additionally, the configuration, control, and status registers are also described. Note that individual chapters in this book describe specific initialization aspects for each individual block.

5.4.1 WDT Overview

The device provides a software watchdog timer (WDT) feature to prevent system lock in case the software becomes trapped in loops with no controlled exit. Watchdog timer operations are configured in the system watchdog control register (SWCRR).

The watchdog counter is a free-running down-counter that generates a reset or a non-maskable interrupt on underflow. To prevent a reset, software must periodically restart the countdown. The WDT is responsible for asserting a hardware reset or machine-check interrupt (*mcp*) if the software fails to service the software watchdog timer for a certain period of time (for example, because software is lost or trapped in a loop with no controlled exit).

Figure 5-23 shows a high-level block diagram of the WDT.


Figure 5-23. Software Watchdog Timer High-Level Block Diagram

The software watchdog timer is enabled after reset to cause a hardware reset if it times out. The user has the option of disabling the software watchdog if it is not needed. If used, the software watchdog timer requires a special service sequence to be executed periodically. Without this periodic servicing, the software watchdog timer times out and issues a reset or a non-maskable interrupt.

5.4.2 WDT Features

The WDT includes the following key features:

- Based on 16-bit prescaler and 16-bit down-counter
- Provides a selectable range for the time-out period
- Provides ~32.2-sec maximum software time-out delay for 133-MHz input clock
- Functional and programming compatibility with MPC8260 watchdog timer

5.4.3 WDT Modes of Operation

The WDT unit can operate in the following modes:

- WDT enable/disable mode:
If the software watchdog timer is not needed, the user can disable it with software after a system reset. When the watchdog timer is disabled, the watchdog counter and prescaler counter are held in a stopped state.
- WDT output reset/interrupt mode:
Without software periodic servicing, the software watchdog timer times out and issues a reset or a nonmaskable interrupt (*mcp*)
- WDT prescaled/non-prescaled clock mode:
The WDT counter clock can be prescaled by programming the SWCRR[SWPR] bit, which controls the divide-by-65,536 of the WDT counter.

5.4.4 WDT Memory Map/Register Definition

The WDT programmable register map occupies 16 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All WDT registers are 16- or 32-bit wide, located on 16-bit address boundaries, and should be accessed as 16- or 32-bit quantities. All addresses used in this chapter are offsets from the WDT base, as defined in Chapter 3, “Memory Map.”

Table 5-37 shows the WDT memory map.

Table 5-37. WDT Register Address Map

| Offset | Register | Access | Reset Value | Section/ Page |
|---|----------|--------|-------------|---------------|
| Watchdog Timer (WDT)—Block Base Address 0x0_0200 | | | | |
| 0x00–0x03 | Reserved | — | — | — |


Table 5-37. WDT Register Address Map (continued)

| Offset | Register | Access | Reset Value | Section/ Page |
|-----------|--|--------|--|---------------|
| 0x04 | System watchdog control register (SWCRR) | R/W | 0xFFFF_0003 or 0xFFFF_0007 ¹ | 5.4.4.1/5-35 |
| 0x08 | System watchdog count register (SWCNR) | R | 0x0000_FFFF | 5.4.4.2/5-36 |
| 0x0C–0x0D | Reserved | — | — | — |
| 0x0E | System watchdog service register (SWSRR) | R/W | 0x0000 | 5.4.4.3/5-36 |

¹ SWCRR[SWEN] reset value directly depends on RCWHR[SWEN] (reset configuration word high).

5.4.4.1 System Watchdog Control Register (SWCRR)

The system watchdog control register (SWCRR), shown in [Figure 5-24](#), controls the software watchdog period and configures watchdog timer operation. SWCRR can be read at any time but can be written only once after system reset.



¹ SWCRR[SWEN] reset value directly depends on RCWHR[SWEN] (reset configuration word high).

Figure 5-24. System Watchdog Control Register (SWCRR)

[Table 5-38](#) defines the bit fields of SWCRR.

Table 5-38. SWCRR Bit Settings

| Bits | Name | Description |
|-------|------|---|
| 0–15 | SWTC | Software watchdog time count The SWTC field contains the modulus that is reloaded into the watchdog counter by a service sequence. When a new value is loaded into SWCRR[SWTC], the software watchdog timer is not updated until the servicing sequence is written to the SWSRR. If SWCRR[SWEN] is loaded with 0, the modulus counter does not count. The new value is also used at the next and all subsequent reloads. Reading the SWCRR register returns the value in the system watchdog control register. Reset initializes the SWCRR[SWTC] field to 0xFFFF. Note: The prescaler counter is reset any time a new value is loaded into the watchdog counter and also during reset. |
| 16–28 | — | Write reserved, read = 0 |
| 29 | SWEN | Watchdog enable bit Enables the watchdog timer. The reset value directly depends on the value of the RCWHR[SWEN] bit. It should be cleared by software after a system reset to disable the software watchdog timer. When the watchdog timer is disabled, the watchdog counter and prescaler counter are held in a stopped state. 0 Watchdog timer disabled 1 Watchdog timer enabled Note: After software writes the SWRI bit, the state of SWEN cannot be changed. |

Table 5-38. SWCRR Bit Settings (continued)

| Bits | Name | Description |
|------|------|--|
| 30 | SWRI | Software watchdog reset/interrupt select bit A WDT time out causes either a hard reset or machine check interrupt to the core. 0 Software watchdog timer causes a machine check interrupt to the core 1 Software watchdog timer causes a hard reset |
| 31 | SWPR | Software watchdog counter prescale bit Controls the divide-by-65,536 WDT counter prescaler 0 The WDT counter is not prescaled. 1 The WDT counter clock is prescaled. |

5.4.4.2 System Watchdog Count Register (SWCNR)

The system watchdog count register (SWCNR), shown in [Figure 5-25](#), provides visibility to the watchdog counter value. SWCNR is a read-only register. Writes to SWCNR have no effect and terminate without transfer error exception.




Figure 5-25. System Watchdog Count Register (SWCNR)

[Table 5-39](#) defines the bit fields of SWCNR.

Table 5-39. SWCNR Bit Settings

| Bits | Name | Description |
|-------|-------|--|
| 0–15 | — | Write reserved, read = 0 |
| 16–31 | SWCNR | Software watchdog count field. The read-only SWCNR[SWCNR] field reflects the current value in the watchdog counter. Writing to the SWCNR register has no effect, and write cycles are terminated normally. Reset initializes the SWCNR[SWCNR] field to 0xFFFF. Note: Reading the 16 least-significant bits of 32-bit SWCNR register with two 8-bit reads is not guaranteed to return a coherent value. |

5.4.4.3 System Watchdog Service Register (SWSRR)

The system watchdog service register (SWSRR) is shown in [Figure 5-26](#). When the watchdog timer is enabled, a write of 0x556C followed by a write 0xAA39 to the SWSRR register before the watchdog counter times out prevents a device reset. If the SWSRR register is not serviced before the timeout, a signal from the watchdog timer to the reset or interrupt controller module asserts a system reset or interrupt (depending on the setting of SWCRR[SWRI]).

Both writes must occur before the timeout in the order listed, but any number of instructions can be executed between the two writes. However, writing any value other than 0x556C or 0xAA39 to the SWSRR register resets the servicing sequence, requiring both values to be written to keep the watchdog

timer from causing a reset. Reset initializes the SWSRR[WS] field to 0x0000. SWSRR can be written at any time, but returns all zeros when read.




Figure 5-26. System Watchdog Service Register (SWSRR)

Table 5-40 defines the bit fields of SWCNR.

Table 5-40. SWSRR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–15 | WS | Software watchdog service field. The user should periodically write 0x556C followed by 0xAA39 to this register to prevent a software watchdog timer timeout. SWSRR[WS] can be written at any time, but returns all zeros when read. |

5.4.5 Functional Description

5.4.5.1 Software Watchdog Timer Unit

The device provides a software watchdog timer (WDT) feature to prevent system lock in case the software becomes trapped in loops with no controlled exit. Watchdog timer operations are configured in the system watchdog control register (SWCRR).

The software watchdog timer is enabled after reset to cause a soft reset or non-maskable interrupt (MCP) if it times out. If the software watchdog timer is not needed, the user must clear SWCRR[SWEN] to disable it. If used, the software watchdog timer requires a special service sequence to be executed periodically. Without this periodic servicing, the software watchdog timer times out and issues a reset or a nonmaskable interrupt, as programmed in SWCRR[SWRI]. Once software writes SWRI, the state of SWEN cannot be changed.

The software watchdog timer service sequence consists of the following two steps:

- Write 0x556C to the system watchdog service register (SWSRR)
- Write 0xAA39 to SWSRR

The service sequence reloads the watchdog timer and the timing process begins again. If a value other than 0x556C or 0xAA39 is written to the SWSRR, the entire sequence must start over. Although the writes must occur in the correct order before a time-out, any number of instructions can be executed between the

writes. This allows interrupts and exceptions to occur between the two writes when necessary. [Figure 5-27](#) shows a state diagram for the watchdog timer.




Figure 5-27. Software Watchdog Timer Service State Diagram

Although most software disciplines permit or even encourage the watchdog concept, some systems require a selection of time-out periods. For this reason, the software watchdog timer must provide a selectable range for the time-out period. [Figure 5-28](#) shows how to handle this need.




Figure 5-28. Software Watchdog Timer Functional Block Diagram

[Figure 5-28](#) shows that the range is determined by SWCRR[SWTC]. The value in SWTC is then loaded into a 16-bit decremter clocked by the system clock. An additional divide-by-65,536 prescaler value is used when needed.

The decremter begins counting when loaded with a value from SWTC. After the timer reaches 0x0, a software watchdog expiration request is issued to the reset or *mcp* (machine check) control logic. Upon reset, SWTC is set to the maximum value and is again loaded into the system watchdog service register (SWSRR), starting the process over. When a new value is loaded into SWTC, the software watchdog timer is not updated until the servicing sequence is written to the SWSRR. If SWCRR[SWEN] is loaded with 0, the modulus counter does not count.

5.4.5.2 Modes of Operation

The WDT unit can operate in the following modes:

- WDT enable/disable mode:

If the software watchdog timer is not needed, the user can disable it. The SWCRR[SWEN] bit enables the watchdog timer. It should be cleared by software after a system reset to disable the software watchdog timer. When the watchdog timer is disabled, the watchdog counter and prescaler counter are held in a stopped state.

 - WDT enable mode (SWCRR[SWEN] = 1)
 - WDT disable mode (SWCRR[SWEN] = 0)
- WDT reset/interrupt output mode

Without software periodic servicing, the software watchdog timer times out and issues a reset or a nonmaskable interrupt (*mcp*), programmed in SWCRR[SWRI].

According to the value of SWCRR[SWRI], the WDT timer causes a hard reset or machine check interrupt to the core.

 - Reset mode (SWCRR[SWRI] = 1).

Software watchdog timer causes a hard reset (this is the default value after hard reset).
 - Interrupt mode (SWCRR[SWRI] = 0).

Software watchdog timer causes a machine check interrupt to the core.
- WDT prescaled/non-prescaled clock mode

The WDT counter clock can be prescaled by programming the SWCRR[SWPR] bit that controls the divide-by-65,536 of the WDT counter.

 - Prescale mode (SWCRR[SWPR] = 1)

The WDT clock is prescaled.
 - Non-prescale mode (SWCRR[SWPR] = 0)

The WDT clock is not prescaled.

5.4.6 Initialization/Application Information

5.4.6.1 WDT Programming Guidelines

The software watchdog timer is enabled (by the default value of SWCRR[SWEN]) after reset. The following initialization sequence of WDT is required:

- WDT disabling

If the software watchdog timer is not needed, the user must clear SWCRR[SWEN] bit to disable the WDT not later than its timer times out (~32.2 sec. for a 133-MHz system clock).
- WDT initial servicing

If the software watchdog timer is to be used, the special service sequence, described in [Section 5.4.5.1, “Software Watchdog Timer Unit,”](#) must be executed after system reset and not later than the first WDT time-out (~32.2 sec. for a 133-MHz system clock).

Subsequently, periodical WDT servicing should be performed according to the programming guidelines given in [Section 5.4.5.1, “Software Watchdog Timer Unit.”](#)

5.5 Real Time Clock Module (RTC)

The following sections describe the theory of operation of the real time clock module (RTC) including a definition of the external signals and the functions it serves. Additionally, the configuration, control, and status registers are described. Note that individual chapters in this reference manual describe additional specific initialization aspects for each individual block.

5.5.1 RTC Overview

The device platform provides a real time clock (RTC) timer suitable for timestamping or time and calendar generation. It can maintain a one-second count which is unique over a period of approximately 136 years.

The RTC can be initialized by software with an initial count value using the real time counter load register (RTLDR). It can also be programmed to generate an interrupt every second. The real time counter control register (RTCTR) is used to enable or disable the various timer functions. The real time counter event register (RTEVR) is used to report the interrupt source. The RTC counter is initialized by software and can be disabled if needed.

[Figure 5-29](#) shows the high level RTC block diagram.




Figure 5-29. Real Time Clock Module High Level Block Diagram

5.5.2 RTC Features

The key features of the RTC include the following:

- Maintains a one-second count, unique over a period of thousand of years
- 32-bit RTC counter can be initialized by software to specific initial count value
- Provides an alarm function with programmable and maskable alarm interrupt
- Provides programmable and maskable every second interrupt
- Uses two possible clock sources: the CSB bus clock or an external RTC clock
- RTC function can be disabled if needed

5.5.3 RTC Modes of Operation

The RTC unit can operate in the following modes:

- RTC enable/disable mode
- RTC every-second interrupt enable/disable mode
- RTC alarm interrupt enable/disable mode
- RTC internal/external input clock mode

5.5.4 RTC External Signal Description

This section provides an overview and detailed descriptions of the RTC signals.

There is one distinct external RTC clock input signal, defined in [Table 5-41](#).

Table 5-41. RTC Signal Properties

| Name | Port | Function | I/O | Reset | Pull Up |
|---------|---------|------------------------|-----|-------|---------|
| RTC_CLK | RTC_CLK | Real time clock input. | I | N/A | — |

[Table 5-42](#) provides a detailed description of the external RTC signal.

Table 5-42. RTC External Signal—Detailed Signal Description

| Signal | I/O | Description |
|---------|-----|---|
| RTC_CLK | I | This signal is used as the timebase for the real time clock module. |
| | | State Meaning — |
| | | Timing — |

5.5.5 RTC Memory Map/Register Definition

The RTC programmable register map occupies 32 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All RTC registers are 32-bit wide that are located on 32-bit address boundaries and should only be accessed as a 32-bit quantities.

All addresses used in this section are offsets from the RTC base, as defined in [Chapter 3, “Memory Map.”](#)

[Table 5-37](#) shows the memory map of the RTC.

Table 5-43. RTC Register Address Map

| Offset | Register | Access | Reset Value | Section/ Page |
|--|--|--------|-------------|------------------------------|
| Real Time Clock (RTC)—Block Base Address 0x0_0300 | | | | |
| 0x00 | Real time counter control register (RTCNR) | R/W | 0x0000_0000 | 5.5.5.1/5-42 |
| 0x04 | Real time counter load register (RTLDR) | R/W | 0x0000_0000 | 5.5.5.2/5-43 |

5.5.5.2 Real Time Counter Load Register (RTLDR)

The real time counter load register (RTLDR), shown in [Figure 5-31](#), contains the 32-bit value to be loaded in the 32-bit RTC counter.




Figure 5-31. Real Time Counter Load Register (RTLDR)

[Table 5-45](#) defines the bit fields of RTLDR.

Table 5-45. RTLDR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–31 | CLDV | Contains the 32-bit value to be loaded in the 32-bit RTC counter. |

5.5.5.3 Real Time Counter Prescale Register (RTPSR)

The real time counter prescale register (RTPSR), shown in [Figure 5-32](#), is a read/write register used to configure the RTC prescaler’s value.




Figure 5-32. Real Time Counter Prescale Register (RTPSR)

[Table 5-46](#) defines the bit fields of RTPSR.

Table 5-46. RTPSR Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–31 | PRSC | RTC prescaler bits. Select the input clock divider for the RTC counter clock. The prescaler is programmed to divide the RTC clock input by values from 1 to 4,294,967,296. The value 0x0000 divides the clock by 1 and 0xFFFF_FFFF divides the clock by 4,294,967,296. To accurately predict the timing of the next count, change the RTPSR[PRSC] field only when the enable bit RTCNR[CLE] is clear. Changing the RTPSR[PRSC] bits resets the prescaler counter. System reset and the loading of a new value into the counter both reset the prescaler counter. Clearing RTCNR[CLE] stops the prescaler counter. |

5.5.5.4 Real Time Counter Register (RTCTR)

The real time counter register (RTCTR), shown in [Figure 5-33](#), is a read-only register that shows the current value in the RTC counter.

The CNTV value is not affected by reads or writes to RTCTR.

System Configuration




Figure 5-33. Real Time Counter Register (RTCTR)

Table 5-47 defines the bit fields of RTCTR.

Table 5-47. RTCTR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–31 | CNTV | RTC counter value field. RTCTR[CNTV] contains the current value of the time counter. This is a read-only field. Writes have no effect on RTCTR[CNTV]. |

5.5.5.5 Real Time Counter Event Register (RTEVR)

The real time counter event register (RTEVR), shown in Figure 5-34, is used to report the source of the interrupts. The register can be read at any time.




Figure 5-34. Real Time Counter Event Register (RTEVR)

RTEVR bits are cleared by writing ones. Writing zeros does not affect the value of the status bits.

Table 5-48 defines the bit fields of RTEVR.

Table 5-48. RTEVR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–29 | — | Write reserved, read = 0 |
| 30 | AIF | Alarm interrupt flag bit. Used to indicate the alarm interrupt. The bit is set if the RTC issues an interrupt when the RTC counter value equals RTALR[ALRM]. |
| 31 | SIF | Second interrupt flag bit. Used to indicate the every-second interrupt. This status bit is set each time that the prescaler count reaches zero and should be cleared by software. |

5.5.5.6 Real Time Counter Alarm Register (RTALR)

The real time counter alarm register (RTALR), shown in [Figure 5-35](#), contains the 32-bit alarm (ALRM) value. When the value of the RTC counter equals the RTALR[ALRM] value, a maskable interrupt is generated.




Figure 5-35. Real Time Counter Alarm Register (RTALR)

[Table 5-49](#) defines the bit fields of RTALR.

Table 5-49. RTALR Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–31 | ALRM | RTC alarm value. The alarm interrupt is generated when the value of the RTC counter equals RTALR[ALRM]. |

5.5.6 Functional Description

5.5.6.1 Real Time Counter Unit

The real time clock (RTC) timer is suitable for time stamping or time and calendar generation. It can maintain a one-second count which is unique over a period of approximately 136 years. Software can convert this count into time-of-day or calendar information if required. An alarm function is also provided. The RTC can be clocked by the internal system bus clock or by an external clock source. The RTC consists of 32-bit up-counter which is incremented by an one-second count clock derived from the RTC input clock. The RTC can be programmed to generate a maskable interrupt when the time value matches the value in its associated alarm register.

The RTC can be initialized by software with an initial count value in the real time counter load register (RTLDR). It can also be programmed to generate an interrupt every second. The real time counter control register (RTCTR) is used to enable or disable the various timer functions. The real time counter event register (RTEVR) is used to report the interrupt source. The RTC counter is reset to zero on hard reset but is not affected by soft reset. It is initialized by the software. The RTC function can be disabled.

Figure 5-36 shows the functional RTC block diagram.




Figure 5-36. Real Time Clock Module Functional Block Diagram

5.5.6.2 RTC Operational Modes

The RTC unit can operate in the following modes:

- RTC enable/disable mode:

RTCNR[CLEN] enables the RTC timer. It should be set by software after a system reset to enable the RTC timer.

 - RTC disable mode (RTCNR[CLEN] = 0)

When the RTC's clock is disabled, counter maintains its old value (default).
 - RTC enable mode (RTCNR[CLEN] = 1)

When the counter's clock is enabled, it continues counting using the previous value.
- RTC every-second interrupt enable/disable mode:
 - RTC every-second interrupt enable mode (RTCNR[SIM] = 1)

In this mode the RTC set the RTEVR[SIF] flag and generate an interrupt after the RTC's 32-bit counter reaches zero.
 - RTC every-second interrupt disable mode (RTCNR[SIM] = 0)

In this mode the RTC sets the RTEVR[SIF] flag but does not generate an interrupt after the RTC's 32-bit counter reaches zero.
- RTC alarm interrupt enable/disable mode:
 - RTC alarm interrupt enable mode (RTCNR[AIM] = 1)

In this mode, the RTC sets the RTEVR[AIF] flag and generates an interrupt each time when the RTC's 32-bit counter reaches the RTALR[ALRM] value.
 - RTC alarm interrupt disable mode (RTCNR[AIM] = 0)

In this mode the RTC sets the RTEVR[AIF] flag but does not generate an interrupt when the RTC's 32-bit counter reaches the RTALR[ALRM] value.

- RTC internal/external input clock mode:
The input clock to the RTC may be the CSB clock or an external 32.768-kHz crystal.
 - RTC uses the internal input clock mode (RTCNR[CLIN] = 0)
 - RTC uses the external 32.768-kHz crystal clock (RTCNR[CLIN] = 1)

5.5.7 RTC Programming Guidelines

The following initialization sequence for the RTC is recommended:

1. Write to RTPSR to set the RTC prescaler to the desired value
2. Write to RTLDR to initialize the RTC initial value
3. Write to RTALR to program the RTC alarm value, if needed
4. Write to RTCNR to configure and start the RTC operation: RTC input clock source, second/alarm interrupt mask, RTC clock enable.

5.6 Periodic Interval Timer (PIT)

The following sections describe theory of operation of the periodic interval timer (PIT) including a definition of the external signals and the functions it serves. Additionally, the configuration, control, and status registers are described. Note that individual chapters in this reference manual describe additional specific initialization aspects for each individual block.

5.6.1 PIT Overview

The periodic interval timer (PIT) that generates periodic interrupts for a real-time operating system or an application software.

The PIT consists of a 32-bit down-counter which is decremented by a clock derived from a CSB clock or from an external 32.768-kHz crystal. The 32-bit counter decrements to zero when loaded with a initial value from the periodic interval timer load register (PTLDR). The periodic interval timer control register (PTCTR) is used to enable or disable the various timer functions. The periodic interval timer event register (PTEVR) is used to report the interrupt source. The PIT function can be disabled if needed.

Figure 5-37 shows the functional PIT block diagram.




Figure 5-37. Periodic Interval Timer High Level Block Diagram

5.6.2 PIT Features

The key features of the PIT include the following:

- Maintains a 32-bit down-counter, clocked by a 32-bit prescaled input clock
- 32-bit PIT counter can be initialized by software to specific initial count value
- Provides programmable and maskable periodic interrupt
- Uses two possible clock sources: the CSB clock or an external PIT clock
- PIT function can be disabled

5.6.3 PIT Modes of Operation

The PIT unit can operate in the following modes:

- PIT enable/disable mode
- PIT periodic interrupt enable/disable mode
- PIT internal/external input clock mode

5.6.4 PIT External Signal Description

This section provides an overview and detailed descriptions of the PIT signals.

There is one distinct external input signal (PIT clock), defined in [Table 5-50](#).

Table 5-50. PIT Signal Properties

| Name | Port | Function | I/O | Reset | Pull Up |
|---------|---------|--------------------------|-----|-------|---------|
| PIT_CLK | PIT_CLK | Periodic interval timer. | I | N/A | — |

[Table 5-51](#) describes of the external PIT signal.

Table 5-51. PIT External Signal—Detailed Signal Descriptions

| Signal | I/O | Description | |
|---------|-----|---|---|
| PIT_CLK | I | This signal is used as the timebase for the periodic interval timer module. | |
| | | State Meaning | — |
| | | Timing | — |

5.6.5 PIT Memory Map/Register Definition

The PIT programmable register map occupies 32 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All PIT registers are 32-bit wide and reside on 32-bit address boundaries and should only be accessed as 32-bit quantities.

All addresses used in this chapter are offsets from PIT base, as defined in [Chapter 3, “Memory Map.”](#)

5.6.5.2 Periodic Interval Timer Load Register (PTLDR)

The periodic interval timer load register (PTLDR), shown in [Figure 5-39](#), contains the 32-bit value to be loaded in a 32-bit PIT counter.




Figure 5-39. Periodic Interval Timer Load Register (PTLDR)

[Table 5-54](#) defines the bit fields of PTLDR.

Table 5-54. PTLDR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–31 | CLDV | Contains the 32-bit value to be loaded in a 32-bit PIT counter. |

5.6.5.3 Periodic Interval Timer Prescale Register (PTPSR)

The periodic interval timer prescale register (PTPSR), shown in [Figure 5-40](#), is a read/write register that used to configure the PIT prescaler’s value.




Figure 5-40. Periodic Interval Timer Prescale Register (PTPSR)

[Table 5-55](#) defines the bit fields of PTPSR.

Table 5-55. PTPSR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–31 | PRSC | PIT prescaler bits. Selects the input clock divider to generate the PIT counter clock. The prescaler is programmed to divide the PIT clock input by values from 1 to 4,294,967,296. The value 0x0000 divides the clock by 1 and 0xFFFF_FFFF divides the clock by 4,294,967,296. To accurately predict the timing of the next count, change the PRSC bit only when the enable bit PTCNR[CLE] is clear. Changing PRSC resets the prescaler counter. System reset and the loading of a new value into the counter also reset the prescaler counter. Clearing the PTCNR[CLE] bit stops the prescaler counter. |

5.6.5.4 Periodic Interval Timer Counter Register (PTCTR)

The periodic interval timer counter register (PTCTR), shown in [Figure 5-41](#), is a read-only register that shows the current value in the PIT counter. The PTCTR counter is not affected by reads or writes.




Figure 5-41. Periodic Interval Timer Counter Register (PTCTR)

[Table 5-56](#) defines the bit fields of PTCTR.

Table 5-56. PTCTR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–31 | CNTV | PIT counter value field. Contains the current value of the time counter. This is a read-only field. Writes have no effect on PTCTR[CNTV]. |

5.6.5.5 Periodic Interval Timer Event Register (PTEVR)

The periodic interval timer event register (PTEVR), shown in [Figure 5-42](#), is used to report the source of the interrupts. The register can be read at any time.

PTEVR bits are cleared by writing ones. Writing zeros does not affect the value of the status bits.




Figure 5-42. Periodic Interval Timer Event Register (PTEVR)

[Table 5-57](#) defines the bit fields of PTEVR.

Table 5-57. PTEVR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–30 | — | Write reserved, read = 0 |
| 31 | PIF | Periodic interrupt flag bit. It is asserted after the SPMPIT counter counts to zero. This status bit should be cleared by software. |

5.6.6 Functional Description

5.6.6.1 Periodic Interval Timer Unit

The PIT generates periodic interrupts for use with a real-time operating system or the application software. It consists of a 32-bit down-counter which is decremented by a clock derived from the CSB clock or from

the PIT clock. The 32-bit counter decrements to zero when loaded with a initial value from the periodic interval timer load register (PTLDR). After the timer reaches zero, PTEVR[PIF] is set and an interrupt is generated if PTCNR[PIM] = 1. At the next count cycle, the value in the PTLDR[CLDV] is loaded into the counter and the process repeats. When a new value is loaded into the PTLDR[CLDV], the PIT is updated, the prescaler counter is reset, and the counter begins counting. Setting of PTEVR[PIF] generates an interrupt, that remains pending until PTEVR[PIF] is cleared. If PTEVR[PIF] is set again before being cleared, the interrupt remains pending until PTEVR[PIF] is cleared. Any write to the PTLDR[CLDV] stops the current countdown and the count resumes with the new value in PTLDR[CLDV]. If PTCNR[CLEN] = 0, the PIT cannot count and retains the old count value. PTCTR contain the PIT current value. The PIT function can be disabled if needed.

Figure 5-43 shows the functional PIT block diagram.




Figure 5-43. Periodic Interval Timer Functional Block Diagram

5.6.6.2 PIT Operational Modes

The PIT unit can operate in the following modes:

- PIT enable/disable mode:
 - The PTCNR[CLEN] bit enables the PIT timer. It should be set by software after a system reset to enable the PIT timer.
 - PIT disable mode (PTCNR[CLEN] = 0). When the PIT’s clock is disabled, counter maintains its old value.
 - PIT enable mode (PTCNR[CLEN] = 1). When the counter’s clock is enabled, it continues counting using the previous value.
- PIT periodic interrupt enable/disable mode:
 - PIT periodic interrupt enable mode (PTCNR[PIM] = 1). After the PIT’s 32-bit counter reaches zero, the PIT sets the PTEVR[PIF] flag and generates an interrupt.
 - PIT periodic interrupt disable mode (PTCNR[PIM] = 0). After the PIT’s 32-bit counter reaches zero, the PIT sets the PTEVR[PIF] flag but does not generate an interrupt.
- PIT internal/external input clock mode:
 - The input clock to the PIT may be an internal system clock or the PIT clock.
 - PIT use the internal input clock mode (PTCNR[CLIN] = 0)

— PIT use the PIT clock (PTCNR[CLIN] = 1)

5.6.7 PIT Programming Guidelines

The following initialization sequence of PIT is recommended:

1. Write to PTPSR to set the PIT prescaler to the desired value
2. Write to PTLDR to initialize the PIT initial value
3. Write to PTCNR to configure and start the PIT operation: PIT input clock source, periodic interrupt mask, PIT clock enable.

For real-time clock programming guidelines, see [Section 5.5.7, “RTC Programming Guidelines.”](#)

5.7 General-Purpose Timers (GTM)

The following sections describe theory of operation of the general purpose (global) timer module, including a definition of the external signals and the functionality. Additionally, the configuration, control, and status registers are described. Note that individual chapters in this book describe additional specific initialization aspects for each individual block.

5.7.1 GTM Overview

Each global timer module (GTM) includes four identical 16-bit general-purpose timers, two 32-bit timers or one 64-bit timer. Each GTM timer consists of a timer prescale register (GTPSR), a timer mode register (GTMDR), a timer capture register (GTCPR), a timer counter register (GTCNR), a timer reference register (GTRFR), a timer event register (GTEVR), and a timer global configuration register (GTCFR). The GTPSRs and the GTMDRs contain the primary and secondary prescalers, programmed by the user.

Figure 5-44 shows the functional GTM block diagram.




Figure 5-44. Global Timers Block Diagram

5.7.2 GTM Features

The key features of the timer include the following:

- The maximum input clock is the system bus clock
- Four 16-bit programmable timers
- Two timers cascaded internally or externally to form a 32-bit timer
- One timer cascaded internally or externally to form a 64-bit timer
- Maximum period of ~515 seconds (at 133-MHz bus clock in slow go mode, primary and secondary prescaler = 256) for 16-bit timer
- Maximum period of ~8246 seconds (at 133-MHz bus clock and prescaler = 256) for 32-bit timer
- Maximum period of thousands of years (at 133-MHz bus clock and prescaler = 256) for 64-bit timer
- 7.5-nanosecond timer resolution (at 133-MHz bus clock and no prescaler)
- Resolution and maximum period can be traded off by selecting prescaler divisor
- Three programmable input clock sources for the timer prescalers
- Input capture capability

- Output compare with programmable mode for the output pin
- Free run and restart modes
- Functional and programming compatibility with MPC8260 timers

5.7.3 GTM Modes of Operation

The GTM unit can operate in the following modes:

- Cascaded modes
- Clock source modes
- Reference modes
- Capture modes

5.7.3.1 Cascaded Modes

$GTCFRn[PCAS]$ and $GTCFR2[SCAS]$ are used to put the timers into different cascaded modes:

- Non-cascaded mode: Each timer (timer 1, timer 2, timer 3, and timer 4), function as a independent 16-bit timer with a 16-bit GTRFR, GTCPR, GTMDR, and GTCNR. In this mode, the non-cascaded GTRFR, GTCPR, and GTCNR should be referenced with corresponding 16-bit bus cycles.
- Pair-cascaded mode: In this mode, two 16-bit timers can be internally cascaded to form a 32-bit counter: timer 1 can be internally cascaded to timer 2 and timer 3 may be internally cascaded to timer 4. Because the decision to cascade timers is made independently, the user has the option of selecting two 16-bit timers and one 32-bit timer, or two 32-bit timers. When working in the pair-cascaded mode, the cascaded GTRFR, GTCPR, and GTCNR should be referenced with 32-bit bus cycles.
- Super-cascaded mode: In this mode, all four 16-bit timers can be internally cascaded to form a 64-bit counter. When working in the super-cascaded mode, the cascaded GTRFR, GTCPR, and GTCNR should be referenced with two 32-bit bus cycles.

5.7.3.2 Clock Source Modes

The clock input to the timer's prescaler can be selected from three sources:

- The system clock
- The system slow go clock (system bus clock internally divided by 16)
- The corresponding TINx pin

5.7.3.3 Reference Modes

Each timer can be configured to count until a reference is reached and then either begin a new time count immediately or continue to run. The FRR bit of the corresponding GTMRR selects each mode.

- Free run reference mode. The corresponding timer count continues to increment after the reference value is reached.
- Reset reference mode. The corresponding timer count is reset immediately after the reference value is reached.

5.7.3.4 Capture Modes

Each timer has a 16-bit field in GTCPR, used to latch the value of the counter when a defined transition of TINx is sensed by the corresponding input capture edge detector.

- Normal gate mode enables the count on a falling edge of the $\overline{\text{TGATE}}$ pin and disables the count on the rising edge of $\overline{\text{TGATE}}$. This mode allows the timer to count conditionally, based on the state of $\overline{\text{TGATE}}$.
- The restart gate mode performs the same function as normal mode, except it also resets the counter on the falling edge of the $\overline{\text{TGATE}}$ pin. This mode has applications in pulse interval measurement and bus monitoring.

5.7.4 GTM External Signal Description

This section provides an overview and detailed descriptions of the GTM signals.

There are four distinct external input timer capture signals (TIN1, TIN2, TIN3, and TIN4), four distinct external input timer gate signals ($\overline{\text{TGATE1}}$, $\overline{\text{TGATE2}}$, $\overline{\text{TGATE3}}$, and $\overline{\text{TGATE4}}$), and four distinct external timer output signals ($\overline{\text{TOUT1}}$, $\overline{\text{TOUT2}}$, $\overline{\text{TOUT3}}$, and $\overline{\text{TOUT4}}$). The GTM interface signals are defined in [Table 5-58](#).

Table 5-58. GTM Signal Properties

| Name | Port | Function | I/O | Reset | Require Pull Up |
|----------------------------|----------------------------|--|-----|-------|-----------------|
| TIN1 | TIN1 | Global timer 1 capture control signal | I | 0 | No |
| TIN2 | TIN2 | Global timer 2 capture control signal | I | 0 | No |
| TIN3 | TIN3 | Global timer 3 capture control signal | I | 0 | No |
| TIN4 | TIN4 | Global timer 4 capture control signal | I | 0 | No |
| $\overline{\text{TGATE1}}$ | $\overline{\text{TGATE1}}$ | Global timer 1 counter gate control signal | I | 0 | No |
| $\overline{\text{TGATE2}}$ | $\overline{\text{TGATE2}}$ | Global timer 2 counter gate control signal | I | 0 | No |
| $\overline{\text{TGATE3}}$ | $\overline{\text{TGATE3}}$ | Global timer 3 counter gate control signal | I | 0 | No |
| $\overline{\text{TGATE4}}$ | $\overline{\text{TGATE4}}$ | Global timer 4 counter gate control signal | I | 0 | No |
| $\overline{\text{TOUT1}}$ | $\overline{\text{TOUT1}}$ | Global timer 1 counter output signal | O | 1 | No |
| $\overline{\text{TOUT2}}$ | $\overline{\text{TOUT2}}$ | Global timer 2 counter output signal | O | 1 | No |
| $\overline{\text{TOUT3}}$ | $\overline{\text{TOUT3}}$ | Global timer 3 counter output signal | O | 1 | No |
| $\overline{\text{TOUT4}}$ | $\overline{\text{TOUT4}}$ | Global timer 4 counter output signal | O | 1 | No |

Table 5-59 provides detailed descriptions of the external GTM signals.

Table 5-59. GTM External Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|----------------------|-----|--|
| TIN_n | I | Global timer capture control signal. Used to latch the value of the counter when a defined transition of TIN_n is sensed by the corresponding input capture edge detector. |
| | | State Meaning Asserted/Negated—According to the programmed polarity by the corresponding $GTMDR_n[CE]$. Each timer has a 16-bit GTCPR used to latch the value of the counter when a defined transition of TIN_n is sensed by the corresponding input capture edge detector. Upon a capture or reference event, the corresponding GTEVR bit is set and a maskable interrupt request is issued to the interrupt controller. |
| | | Timing Assertion/Negation—Asynchronous to internal bus clock. TIN_n is internally synchronized to the system bus clock. If TIN_n meets the asynchronous input setup time, the value of counter is captured after one system bus clock when working with the internal clock. |
| \overline{TGATE}_n | I | Global timer counter gate control signal. Used to gate/restart the counter when a defined transition of \overline{TGATE}_n is sensed by the corresponding input capture edge detector. |
| | | State Meaning Asserted/Negated—According to the programmed polarity by the corresponding $GTCFR[GMx]$ bits. In a restart gate mode ($GTCFR[GMn] = 0$), the \overline{TGATE}_n pin is used to enable/disable count. A falling \overline{TGATE}_n pin enables and restarts the count and a rising edge of \overline{TGATE}_n disables the count. In a normal gate mode ($GTCFR[GMn] = 1$), the \overline{TGATE}_n have similar functionality, except the falling edge of \overline{TGATE}_n does not restart the appropriate count value in $GTCNR_n[CNVn]$. |
| | | Timing Assertion/Negation—Asynchronous to internal bus clock. \overline{TGATE}_n is internally synchronized to the system bus clock. If \overline{TGATE}_n meets the asynchronous input setup time, the counter begins counting or stops counting (depending on the signal state and the configured mode) after one system bus clock when working with the internal clock. |
| \overline{TOUT}_n | O | Global timer counter output signal. The GTM output a signal on the timer output pin \overline{TOUT}_n when the reference value is reached. |
| | | State Meaning Asserted/Negated—According to the programmed polarity by the corresponding $GTMDR_n[OMn]$. <ol style="list-style-type: none"> Active-low pulse on \overline{TOUT}_n for one timer input clock cycle as defined by the $GTMDR_n[ICLK_n]$ bits ($GTMDR_n[OMn] = 1$). Thus, \overline{TOUT}_n may be low for one general system clock period, one general system slow go clock period, or one TIN_n pin clock cycle period. Toggle the \overline{TOUT}_n pin ($GTMDR_n[OMn] = 0$). \overline{TOUT}_n begins or stops counting, depending on the signal state and the configured mode. |
| | | Timing Assertion/Negation— \overline{TOUT}_n changes occur on the rising edge of the timer input clock. |

5.7.5 GTM Memory Map/Register Definition

The GTM programmable register map occupies 64 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All GTM registers are 8- or 16-bit wide, located on 8-bit or 16-bit address boundaries, and should only be accessed as 8-bit or 16-bit quantities. All addresses used in this chapter are offsets from GPT Base, as defined in Chapter 3, “Memory Map.”

Table 5-60 shows the memory map of the GTM.

Table 5-60. GTM Register Address Map

| Offset | Register | Access | Reset Value | Section/ Page |
|---|---|--------|-------------|------------------------------|
| General Purpose (Global) Timer Module 1—Block Base Address 0x0_0500 | | | | |
| 0x00 | Timer 1 and 2 global timers configuration register (GTCFR1) | R/W | 0x00 | 5.7.5.1/5-59 |
| 0x01–0x03 | Reserved | — | — | — |
| 0x04 | Timer 3 and 4 global timers configuration register (GTCFR2) | R/W | 0x00 | 5.7.5.1/5-59 |
| 0x05–0x0F | Reserved | — | — | — |
| 0x10 | Timer 1 global timers mode register (GTMDR1) | R/W | 0x0000 | 5.7.5.2/5-62 |
| 0x12 | Timer 2 global timers mode register (GTMDR2) | | | |
| 0x14 | Timer 1 global timers reference register (GTRFR1) | R/W | 0xFFFF | 5.7.5.3/5-63 |
| 0x16 | Timer 2 global timers reference register (GTRFR2) | | | |
| 0x18 | Timer 1 global timers capture register (GTCPR1) | R/W | 0x0000 | 5.7.5.4/5-63 |
| 0x1A | Timer 2 global timers capture register (GTCPR2) | | | |
| 0x1C | Timer 1 global timers counter register (GTCNR1) | R/W | 0x0000 | 5.7.5.5/5-64 |
| 0x1E | Timer 2 global timers counter register (GTCNR2) | | | |
| 0x20 | Timer 3 global timers mode register (GTMDR3) | R/W | 0x0000 | 5.7.5.2/5-62 |
| 0x22 | Timer 4 global timers mode register (GTMDR4) | | | |
| 0x24 | Timer 3 global timers reference register (GTRFR3) | R/W | 0xFFFF | 5.7.5.3/5-63 |
| 0x26 | Timer 4 global timers reference register (GTRFR4) | | | |
| 0x28 | Timer 3 global timers capture register (GTCPR3) | R | 0x0000 | 5.7.5.4/5-63 |
| 0x2A | Timer 4 global timers capture register (GTCPR4) | | | |
| 0x2C | Timer 3 global timers counter register (GTCNR3) | R/W | 0x0000 | 5.7.5.5/5-64 |
| 0x2E | Timer 4 global timers counter register (GTCNR4) | | | |
| 0x30 | Timer 1 global timers event register (GTEVR1) | w1c | 0x0000 | 5.7.5.6/5-64 |
| 0x32 | Timer 2 global timers event register (GTEVR2) | | | |
| 0x34 | Timer 3 global timers event register (GTEVR3) | | | |
| 0x36 | Timer 4 global timers event register (GTEVR4) | | | |
| 0x38 | Timer 1 global timers prescale register (GTPSR1) | R/W | 0x0003 | 5.7.5.7/5-65 |
| 0x3A | Timer 2 global timers prescale register (GTPSR2) | | | |
| 0x3C | Timer 3 global timers prescale register (GTPSR3) | | | |
| 0x3E | Timer 4 global timers prescale register (GTPSR4) | | | |
| General Purpose (Global) Timer Module 2: All registers defined for GTM1 are also defined for GTM2; the base address of GTM2 registers is 0x0_06nn. | | | | |

5.7.5.1 Global Timers Configuration Registers (GTCFR n)

The global timers configuration registers (GTCFR1 and GTCFR2), shown in [Figure 5-45](#) and [Figure 5-46](#), contain configuration parameters used by the timers. These registers allow simultaneous starting, stopping and resetting of a pair of timers (1 and 2 or 3 and 4) or of a groups of timers (1, 2, 3, and 4) if one bus cycle is used. GTCFR is cleared by reset.

NOTE

For proper operation of the timers, do not change the modes of operation and enable the timer in the same register write operation. The modes can be changed when GTCFR n [RST n] is cleared. However, when GTCFR n [RST n] are set, they are the only bits that can be changed.




Figure 5-45. Global Timers Configuration Register 1 (GTCFR1)

[Table 5-61](#) defines the bit fields of GTCFR1.

Table 5-61. GTCFR1 Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0 | PCAS | Pair-cascade mode 0 Normal operation 1 Timers 1 and 2 cascade to form a 32-bit timer. Note: This bit is ignored in super-cascade mode (GTCFR2[SCAS] = 1). Note: It is allowed to change the value of this bit only when the corresponding timers are in reset mode. Thus, the user should first clear the RST1 and RST2 bits (without changing PCAS) and then, in a separate write to the register, change the value of PCAS. |
| 1 | BCM | Backward compatible mode 0 Provide backward compatibility to PowerQUICC II family timers. In this mode GTCFR1[GM2] bit will control the gate mode for timers 1 and 2 and GTCFR2[GM4] bit will control the gate mode for timers 3 and 4. GTCFR1[GM1] and GTCFR2[GM3] bits are ignored. 1 Normal operational mode |
| 2 | STP2 | Stop timer 2 0 Normal operation 1 Reduce power consumption of the corresponding timer. This bit stops all clocks to the timer 2, except the Register Interface clock, which allows to read and write timer registers. The clocks to the timer remain stopped until the user clears this bit or a hardware reset occurs. |
| 3 | RST2 | Reset timer 2 0 Reset the timer 2, including GTMDR2, GTRFR2, GTCNR2, GTCPR2, and GTEVR2 (a software reset is identical to an external reset). 1 Enable the corresponding timer if the STP2 bit is cleared. |

Table 5-61. GTCFR1 Bit Settings (continued)

| Bits | Name | Description |
|------|------|--|
| 4 | GM2 | <p>Gate mode for $\overline{\text{TGATE2}}$</p> <p>0 Restart gate mode. The $\overline{\text{TGATE2}}$ pin is used to enable/disable count. A low level of $\overline{\text{TGATE2}}$ enables and a falling edge of $\overline{\text{TGATE2}}$ restarts the count (reset the dynamic counter's count value to 0) and a high level of $\overline{\text{TGATE2}}$ disables the count.</p> <p>1 Normal gate mode. This mode is the same as 0, except the falling edge of $\overline{\text{TGATE2}}$ does not restart the appropriate count value in GTCNR2[CNV2].</p> |
| 5 | GM1 | <p>Gate mode for $\overline{\text{TGATE1}}$</p> <p>0 Restart gate mode. The $\overline{\text{TGATE1}}$ is used to enable/disable count. A low level of $\overline{\text{TGATE1}}$ enables and a falling edge of $\overline{\text{TGATE1}}$ restarts the count (reset the dynamic counter's count value to 0) and a high level of $\overline{\text{TGATE1}}$ disables the count.</p> <p>1 Normal gate mode. This mode is the same as 0, except the falling edge of $\overline{\text{TGATE1}}$ does not restart the appropriate count value in GTCNR1[CNV1].</p> <p>Note: In backward compatible mode (GTCFR1[BCM] = 0) this bit is ignored. GTCFR1[GM2] bit will control the gate mode for timers 1 and 2.</p> |
| 6 | STP1 | <p>Stop timer 1</p> <p>0 Normal operation</p> <p>1 Reduce power consumption of the corresponding timer. This bit stops all clocks to the timer 1, except the register interface clock, which allows to read and write timer registers. The clocks to the timer remain stopped until the user clears this bit or a hardware reset occurs.</p> |
| 7 | RST1 | <p>Reset timer 1</p> <p>0 Reset the timer 1, including GTMDR1, GTRFR1, GTCNR1, GTCPR1, and GTEVR1 (a software reset is identical to an external reset).</p> <p>1 Enable the corresponding timer if the STP1 bit is cleared.</p> |

The GTCFR2 register is shown in [Figure 5-46](#).




Figure 5-46. Global Timers Configuration Register 2 (GTCFR2)

Table 5-62 defines the bit fields of GTCFR2.

Table 5-62. GTCFR2 Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0 | PCAS | Pair-cascade mode 0 Normal operation. 1 Timers 3 and 4 cascade to form a 32-bit timer. Note: This bit is ignored in super-cascade mode (GTCFR2[SCAS] = 1). Note: It is allowed to change the value of this bit only when the corresponding timers are in reset mode. Thus, the user should first clear the RST3 and RST4 bits (without changing PCAS) and then, in a separate write to the register, change the value of PCAS. |
| 1 | SCAS | Super cascade mode 0 Normal operation 1 Timers 1, 2, 3 and 4 cascade to form a 64-bit timer. Note: In super-cascade mode (GTCFR2[SCAS] = 1) the pair-cascade mode bits are ignored, (GTCFR1/2[PCAS] = Don't Care). Note: It is allowed to change the value of this bit only when the corresponding timers are in reset mode. Thus, the user should first clear the RST1, RST2, RST3, and RST4 bits (without changing SCAS) and then, in a separate write to the register, change the value of SCAS. |
| 2 | STP4 | Stop timer 4 0 Normal operation 1 Reduce power consumption of the corresponding timer. This bit stops all clocks to the timer 4, except the register interface clock, which allows to read and write timer registers. The clocks to the timer remain stopped until the user clears this bit or a hardware reset occurs. |
| 3 | RST4 | Reset timer 4 0 Reset the timer 4, including GTMDR4, GTRFR4, GTCNR4, GTCPR4, and GTEVR4 (a software reset is identical to an external reset). 1 Enable the corresponding timer if the STP4 bit is cleared. |
| 4 | GM4 | Gate mode for $\overline{\text{TGATE4}}$ 0 Restart gate mode. The $\overline{\text{TGATE4}}$ is used to enable/disable count. A low level of $\overline{\text{TGATE4}}$ enables and a falling edge of $\overline{\text{TGATE4}}$ restarts the count (reset the dynamic counter's count value to 0) and a high level of $\overline{\text{TGATE4}}$ disables the count. 1 Normal gate mode. This mode is the same as 0, except the falling edge of $\overline{\text{TGATE4}}$ does not restart the appropriate count value in GTCNR4[CNV4]. |
| 5 | GM3 | Gate mode for $\overline{\text{TGATE3}}$ 0 Restart gate mode. The $\overline{\text{TGATE3}}$ is used to enable/disable count. A low level of $\overline{\text{TGATE3}}$ enables and a falling edge of $\overline{\text{TGATE3}}$ restarts the count (reset the dynamic counter's count value to 0) and a high level of $\overline{\text{TGATE3}}$ disables the count. 1 Normal gate mode. This mode is the same as 0, except the falling edge of $\overline{\text{TGATE3}}$ does not restart the appropriate count value in GTCNR3[CNV3]. Note: In backward compatible mode (GTCFR1[BCM] = 0) this bit is ignored. The GTCFR2[GM4] bit controls the gate mode for timers 3 and 4. |
| 6 | STP3 | Stop timer 3 0 Normal operation 1 Reduce power consumption of the corresponding timer. This bit stops all clocks to the timer 3, except the register interface clock, which allows to read and write timer registers. The clocks to the timer remain stopped until the user clears this bit or a hardware reset occurs. |
| 7 | RST3 | Reset timer 3 0 Reset the timer 3, including GTMDR3, GTRFR3, GTCNR3, GTCPR3, and GTEVR3 (a software reset is identical to an external reset). 1 Enable the corresponding timer if the STP3 bit is cleared. |

5.7.5.2 Global Timers Mode Registers (GTMDR1–GTMDR4)

The global timers mode registers (GTMDR1, GTMDR2, GTMDR3, and GTMDR4) are shown in Figure 5-47.

Erratic behavior may occur if GTCFR1 and GTCFR2 are not initialized before the GTMDR n . Only GTCFR n [RST n] and GTCFR n [STP n] can be modified at any time.




Figure 5-47. Global Timers Mode Registers (GTMDR1–GTMDR4)

Table 5-63 defines the bit fields of GTMDR.

Table 5-63. GTMDR Bit Settings


| Bits | Name | Description |
|------|------|--|
| 0–7 | SPS | Secondary prescaler value The secondary prescaler is programmed to divide the clock input to corresponding timer by values from 1 to 256. The value 0x00 divides the clock by 1 and 0xFF divides the clock by 256. |
| 8–9 | CE | Capture edge and enable interrupt 00 Disable interrupt on capture event; capture function is disabled 01 Capture on rising TIN n edge only and enable interrupt on capture event. 10 Capture on falling TIN n edge only and enable interrupt on capture event. 11 Capture on any TIN n edge and enable interrupt on capture event. Note: The frequency of TIN n should be slower than system clock (TIN n is sampled internally by system clock to detect TIN n 's rising/falling edge before updating the counter) |
| 10 | OM | Output mode 0 Toggle $\overline{\text{TOUT}}_n$ every time when the corresponding timer matches its reference value. 1 Active-low pulse on $\overline{\text{TOUT}}_n$ for one timer input clock cycle (4 input clock cycles for the system clock) as defined by the ICLK n bits. Thus, $\overline{\text{TOUT}}_n$ may be low for four general system clocks, one general system slow go clock period, or one TIN n pin clock cycle period. Note: $\overline{\text{TOUT}}_n$ changes are internally synchronized to the rising edge of the system clock |
| 11 | ORI | Output reference interrupt enable 0 Disable interrupt for reference reached (does not affect interrupt on capture function). 1 Enable interrupt on reaching the reference value. |
| 12 | FRR | Free run/restart mode 0 Free run. The timer count continues to increment after the reference value is reached. 1 Restart. The timer count is reset immediately after the reference value is reached. |

Table 5-63. GTMDR Bit Settings (continued)

| Bits | Name | Description |
|-------|------|---|
| 13–14 | ICLK | Input clock source for the timer. 00 Internally cascaded input. This selection means: For ICLK1, the timer 1 input is the output of timer 2. For ICLK2, the timer 1 input is the output of timer 2, the timer 2 input is the output of timer 3, the timer 3 input is the output of timer 4. For ICLK3, the timer 3 input is the output of timer 4. For ICLK4 this selection means no input clock is provided to the timer. 01 Internal general system bus clock. 10 Internal slow go clock (divided by 16 system bus clock). 11 TIN n : corresponding TIN1, TIN2, TIN3, or TIN4 pin (falling edge). |
| 15 | GE | Gate enable 0 The $\overline{\text{TGATE}}_n$ signal is ignored. 1 The $\overline{\text{TGATE}}_n$ signal is used to control the timer. |

5.7.5.3 Global Timers Reference Registers (GTRFR1–GTRFR4)

Global timers reference registers, shown in [Figure 5-48](#), are 16-bit memory-mapped, read/write registers containing the 16-bit reference values for each timer’s timeout. The reference value is not reached until $\text{GTCNR}_n[\text{CNV}]$ increments to the value in $\text{GTRFR}_n[\text{TRV}]$.


Figure 5-48. Global Timers Reference Registers (GTRFR1–GTRFR4)

[Table 5-64](#) defines the bit fields of GTRFR.

Table 5-64. GTRFR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–15 | TRV | Timeout reference value. 16-bit timeout reference value for the corresponding timer. Set to all ones by reset. |

5.7.5.4 Global Timers Capture Registers (GTCPR1–GTCPR4)

Global timers capture registers (GTCPR_1 , GTCPR_2 , GTCPR_3 , and GTCPR_4), shown in [Figure 5-49](#), are used to latch the value of the counters according to $\text{GTMDR}_n[\text{CE}]$.



Figure 5-49. Global Timers Capture Registers (GTCPR1–GTCPR4)

Table 5-67 defines the bit fields of $GTEVR_n$.

Table 5-67. $GTEVR_n$ Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–13 | — | Reserved, should be cleared. |
| 14 | REF | Output reference event 0 No event 1 The counter reached the $GTRFR_n[TRV]$ value. $GTMDR_n[ORI]$ is used to enable the interrupt request caused by this event. |
| 15 | CAP | Counter capture event Corresponding timer's 16-bit read/write up-counter value. 0 No event 1 The counter value has been latched into the $GTCPR_n[LCV]$. $GTMDR_n[CE]$ is used to enable generation of this event. |

5.7.5.7 Global Timers Prescale Registers (GTPSR1–GTPSR4)

The global timers prescale registers (GTPSR1, GTPSR2, GTPSR3, and GTPSR4) are shown in Figure 5-52.

Erratic behavior may occur if $GTPSR_n$ is not initialized before the corresponding $GTMDR_n$.




Figure 5-52. Global Timers Prescale Registers (GTPSR1–GTPSR4)

Table 5-68 defines the bit fields of $GTPSR_n$.

Table 5-68. $GTPSR_n$ Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–7 | — | Reserved, should be cleared. |
| 8–15 | PPS | Primary prescaler bits The primary prescaler is programmed to divide the clock input to corresponding timer by values from 1 to 256. The value 0x00 divides the clock by 1 and 0xFF divides the clock by 256. |

NOTE

The total timer prescale value is calculated as follows:

$$GTM_n_{prescaler} = (GTPSR_n[PPS] + 1) \cdot (GTMDR_n[SPS] + 1)$$

This gives a total prescale range from 1 ($GTPSR_n[PPS] = 0x00$, $GTMDR_n[SPS] = 0x00$) to 65,536 ($GTPSR_n[PPS] = 0xFF$, $GTMDR_n[SPS] = 0xFF$).

5.7.6 Functional Description

5.7.6.1 General-Purpose Timer Units

The clock input to the timer's prescaler can be selected from the following sources:

- The system clock
- The system slow go clock (internally divided by 16)

The general system clock is generated in the clock synthesizer and defaults to the system frequency. However, the general system clock has the option to be divided before it leaves the clock synthesizer. This mode, called slow go, is used to save power. Whatever the resulting frequency of the general system clock, the user can either choose that frequency or the frequency divided by 16 as the input to the prescaler of each timer. Alternatively, the user may prefer TIN_n to be the clock source. TIN_n is internally synchronized to the internal clock. If the user has chosen to internally cascade two 16-bit timers to a 32-bit timer, then a timer can use the clock generated by the output of another timer.

The clock input source is selected by the corresponding $GTMDR_n[ICLK]$ bits. The prescalers ($GTMDR_n[SPS]$ and $GTPSR_n[PPS]$) can be programmed to divide the clock input by values from 1 to 65,536 and the output of the prescaler is used as an input to the 16-bit counters. The best resolution of the timer is one clock cycle (7.5 ns at 133 MHz system clock, for example). The maximum period (when the reference value is all ones and the prescaler divides by 256) for one 16-bit timer is ~515 s at 133 MHz.

5.7.6.2 Reference Modes

Each timer can be configured to count until a reference is reached and then either begin a new time count immediately or continue to run. The FRR bit of the corresponding $GTMRR$ selects each mode.

- Free run reference mode ($GTMDR_n[FRR] = 0$)
The corresponding timer count continues to increment after the reference value is reached.
- Reset reference mode ($GTMDR_n[FRR] = 1$)
The corresponding timer count is reset immediately after the reference value is reached.

Upon reaching the reference value, the corresponding $GTEVR_n[REF]$ bit is set and an interrupt is issued if $GTMDR_n[ORI] = 1$. The timers can output a signal on the timer output pin \overline{TOUT}_n if the reference value is reached (selected by the corresponding $GTMDR_n[OM]$). This signal can be an active-low pulse or a toggle of the current output. The output can also be connected internally to the input of another timer, resulting in a 32- or 64-bit timer.

5.7.6.3 Capture Modes

In addition, each timer has a 16-bit field in $GTCPR$, used to latch the value of the counter when a defined transition of TIN_n is sensed by the corresponding input capture edge detector. The timers may be gated/restarted by an external gate signals (\overline{TGATE}_n) that controls the timers. The type of transition triggering the capture is selected by the corresponding $GTMDR_n[CE]$ bits. Upon a capture or reference

event, corresponding $GTEVR_n[REF]$ or $GTEVR_n[CAP]$ is set and a maskable interrupt request is issued to the interrupt controller.

- Normal gate mode enables the count on a falling edge of \overline{TGATE} and disables the count on the rising edge of \overline{TGATE} . This mode allows the timer to count conditionally, based on the state of \overline{TGATE} .
- The restart gate mode performs the same function as normal mode, except it also resets the counter on the falling edge of \overline{TGATE} .

This mode has applications in pulse interval measurement and bus monitoring as follows:

- Pulse measurement—The restart gate mode can measure a low pulse on \overline{TGATE} . The rising edge of \overline{TGATE} completes the measurement and if \overline{TGATE}_n is connected externally to TIN_n , it causes the timer to capture the count value and generate a rising-edge interrupt.
- Bus monitoring—The restart gate mode can detect a signal that is stuck abnormally low. The bus signal should be connected to \overline{TGATE} . The timer count is reset on the falling edge of the bus signal and if the bus signal does not go high again within the number of user-defined clocks, an interrupt can be generated.

The gate function is enabled in the $GTMDR$; the gate operating mode is selected in the $GTCFR_n$.

NOTE

\overline{TGATE} is internally synchronized to the system clock. If \overline{TGATE} meets the asynchronous input setup time, the counter begins or stops counting after one system clock when working with the internal clock.

5.7.6.4 Cascaded Modes

$GTCFR_n[PCAS]$ and $GTCFR2[SCAS]$ are used to put the timers into different cascaded modes:

- Non-cascaded mode ($GTCFR_n[PCAS] = 0$ and $GTCFR2[SCAS] = 0$)
If $GTCFR_n[PCAS] = 0$ and $GTCFR2[SCAS] = 0$, the each timer (timer 1, timer 2, timer 3, and timer 4), function as a independent 16-bit timer with a 16-bit $GTRFR$, $GTCPR$, $GTMDR$, and $GTCNR$ for each one (Figure 5-53). When working in the none-cascaded mode, the non-cascaded $GTRFR$, $GTCPR$, and $GTCNR$ should be referenced with appropriate 16-bit bus cycles.




Figure 5-53. Timers Non-Cascaded Mode Block Diagram

- Pair-cascaded mode ($GTCFR1[PCAS] = 1$ and/or $GTCFR2[PCAS] = 1$, $GTCFR2[SCAS] = 0$)
 In this mode, two 16-bit timers can be internally cascaded to form a 32-bit counter: timer 1 may be internally cascaded to timer 2 and timer 3 may be internally cascaded to timer 4, as shown in [Figure 5-54](#). Since the decision to cascade timers is made independently, the user has the option of selecting two 16-bit timers and one 32-bit timer ($GTCFR1[PCAS] = 1$, $GTCFR2[PCAS] = 0$ or $GTCFR1[PCAS] = 0$, $GTCFR2[PCAS] = 1$), or two 32-bit timers ($GTCFR1[PCAS] = 1$ and $GTCFR2[PCAS] = 1$).

If $GTCFR1[PCAS] = 1$ and/or $GTCFR2[SCAS] = 1$, the two 16-bit timers (timer 1 and timer 2 or timer 3 and timer 4) function as a 32-bit timer with a 32-bit GTRFR, GTCPR, and GTCNR. In this case, GTMDR1/GTMDR3 is ignored, and the modes and functions are defined using GTMDR2/GTMDR4, and GTCFR1/GTCFR2. The capture are controlled from TIN2, and the interrupts are generated from GTEVR2. When working in the pair-cascaded mode, the cascaded GTRFR, GTCPR, and GTCNR should be referenced with 32-bit bus cycles.




Figure 5-54. Timer Pair-Cascaded Mode Block Diagram

- Super-cascaded mode ($GTCFR2[SCAS] = 1$)
 In this mode, all four 16-bit timers can be internally cascaded to form a 64-bit counter, as shown in [Figure 5-55](#).

If $GTCFR2[SCAS] = 1$, the all four 16-bit timers function as a 64-bit timer with a cascaded 32-bit GTRFR, GTCPR, and GTCNR. In this case, registers GTMDR1, GTMDR2, GTMDR3, and GTCFR1 are ignored, and the modes and functions are defined using GTMDR4 and GTCFR2 only. The capture are controlled from TIN4, and the interrupts are generated from GTEVR4. When working in the super-cascaded mode, the cascaded GTRFR, GTCPR, and GTCNR should be referenced with two 32-bit bus cycles.




Figure 5-55. Timers Super-Cascaded Mode Block Diagram

5.7.7 Initialization/Application Information

5.7.7.1 Programming Guidelines

5.7.7.1.1 GTM Registers

The following initialization sequence of GTM is recommended:

- Write to $GTCFR_n$ in order to reset, to stop or to configure the appropriate timer's operation: cascaded timers configuration, gate mode configuration.
- Write to $GTPSR_n[PPS]$ fields in order to program the appropriate timer's clock primary prescaler.
- Write to $GTMDR_n$ in order to choose an input clock, to program the secondary prescaler and to set a desirable appropriate timer's operational mode.

NOTE

Erratic behavior may occur if $GTCFR_n$ and $GTPSR$ are not initialized before the $GTMDR$. Only $GTCFR_n[RST_n]$ can be modified at any time

- Clear $GTEVR_n[REF]$ and $GTEVR_n[CAP]$ by writing 1s in order to clear the previous events.
- Write to $GTRFR$ and to $GTCNR_n$ according to appropriate timer's $GTMDR_n$ programming.

NOTE

A write cycle to a $GTCNR_n[CNV]$ fields sets the register to the written value, causing its corresponding primary and secondary prescalers, ($GTPSR_n[PPS]$ and $GTMDR_n[SPS]$), to be reset.

- Write to $GTCFR_n[STP_n]$ and to $GTCFR_n[RST_n]$ in order to initialize the appropriate timer's operation.

5.8 Power Management Control (PMC)

The device provides a power management control (PMC) unit, which enables the device to smoothly enter and exit low-power modes. Low-power modes may be used when internal units in the device temporarily or permanently do not perform any action.

The device uses one or more of the following methods for power saving:

- Dynamic power management
- Shutting down unused blocks
- Software-controlled power-down states
- Support for the PCI Power Management Interface Specification in both host and agent modes. When the device is in either mode, the PMC is capable of placing the device into one of the supported low-power states and supporting the power management event (PME) signaling protocol.
- A low-power mode (D3Warm) can be achieved using a split supply, which will make the blocks $VDDC$, $NVDD1_ON$, $NVDD2_ON$, and $LVDD2_ON$ power-on selectively while the rest of the blocks (VDD , $GVDD$, $NVDD3_OFF$, $NVDD2_OFF$, $NVDD4_OFF$, $NVDD1_OFF$, $VDD1IO$,

NVDD3_OFF, SATA_VDD, VDD33_ANA, VDD33_PLL, LVDD1_OFF, XPADVDD, XCOREVDD, USB_VDDA, USB_PLL_PWR1, USB_PLL_PWR3) are powered-off. Refer to the *MPC8315E Hardware Specification* for power supply details. Support from wake-up exist on the Ethernet (ETSEC_2), GPIO (GPIO_0, GPIO_1), interrupt (IRQ_1, IRQ_2), and general purpose timers (GTMs) in D3Warm state.

The MPC8315E will support the power states D0, D1, D2, D3Hot, D3Warm, D3Cold. D3Warm is not supported in PCI agent mode. See [Section 5.8.3, “Functional Description,”](#) for details about each state.

Table 5-69. MPC8315E Power States Supported

| | D0 (Full-power) | D1 (Doze) | D2 (Nap) | D3Hot (Sleep) | D3Warm (Standby) | D3Cold (Off) |
|-------|--------------------|--------------|-------------|------------------|---------------------|-----------------|
| Agent | Yes | Yes | Yes | Yes | No | Yes |
| Host | Yes | Yes | Yes | Yes | Yes | Yes |

5.8.1 External Signal Description

[Table 5-70](#) describes the power management signals.

Table 5-70. System Control Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|-----------------------------|-----|--|
| $\overline{\text{QUIESCE}}$ | O | Quiesce state. Indicates that the processor system and PowerPC core are in low power state. |
| | | State Meaning Asserted—The system and PowerPC core are in low power state. Negated—The system and PowerPC core are not in low power state. |
| | | Timing The timing between a quiesce request from the PowerPC core and the assertion of the external indication or between negation of the core's quiesce request and negation of the external indication depends on the current state of the internal system units and may vary accordingly. |
| EXT_PWR_CTRL | O | External power control. Enables the external switch to provide power to the device. |
| | | State Meaning Asserted—Power is supplied to the switchable power supply. Negated—Power is removed from the switchable power supply. |
| | | Timing — |
| PMC_PWR_OK | I | Stable power. Indicates whether the power supply on the board is stable. |
| | | State Meaning Asserted—The external power supply is stable to specifications. Negated—The external power supply is off or not stable to specifications. |
| | | Timing — |

Table 5-73. PMCER Bit Settings (continued)

| Bits | Name | Description |
|------|--------|--|
| 25 | USB | <p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on USB. This wake-up event was caused by a detection of a non idle state on USB interface. See Chapter 17, “Universal Serial Bus Interface,” for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p>Note: This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared. Also, this bit is not used for wake up from D3 warm.</p> |
| 26 | eTSEC1 | <p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on eTSEC1. This wake-up event was caused by a detection of a Magic Packet on the receive path of eTSEC1. See Chapter 19, “Enhanced Three-Speed Ethernet Controllers,” for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p>Note: This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared. Also, this bit is not used for wake up from D3 warm.</p> |
| 27 | eTSEC2 | <p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on eTSEC2. This wake-up event was caused by a detection of a Magic Packet on the receive path of eTSEC2. See Chapter 19, “Enhanced Three-Speed Ethernet Controllers,” for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p>Note: This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.</p> |
| 28 | TIMER | <p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on General Purpose Timer 1. This wake up event was caused by a detection of match between timer current value and timer reference value of the fourth 16 bit unit of GTM1. See Section 5.7, “General-Purpose Timers (GTMs),” for more details. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p>Note: This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.</p> |
| 29 | INT1 | <p>Wake-up event detected.</p> <p>0 A wake-up event did not occur from this wake-up source.</p> <p>1 A wake-up event occurred on external interrupt request 1. This wake-up event was caused by a detection of an active state of the IRQ1 external pin. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect).</p> <p>Note: This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared.</p> |

Table 5-73. PMCER Bit Settings (continued)

| Bits | Name | Description |
|------|------|--|
| 30 | INT2 | Wake-up event detected. 0 A wake-up event did not occur from this wake-up source. 1 A wake-up event occurred on external interrupt request 2. This wake-up event was caused by a detection of an active state of the $\overline{IRQ2}$ external pin. If the corresponding PMCMR bit is set, the PMC will assert interrupt request to the PowerPC core or external PME to the remote host, depending on the state of PMCCR1[PME_EN]. This bit can be cleared by writing a 1 to the bit location (writing zero has no effect). Note: This bit will not be affected by the wake-up event if the corresponding mask bit in PMCMR is cleared. |
| 31 | PMCI | Power management controller interrupt. When set, indicates that one of the following events has occurred: <ul style="list-style-type: none"> • One of the unmasked wake-up events (bits 23–30) occurred and PMCCR1[PME_vPEN] is cleared, or • PM current state (as indicated in PMCCR1[Curr_State]) is different than PM next state (as written to PCIPMR1[Power_State] and indicated in PMCCR1[Next_State]) and PMCCR1[Use_State] is set, or • CSB platform is in low-power mode and a new CSB bus request is detected If PMCMR[PMCI] is set, the PMC interrupt request to the PowerPC core is driven, causing the PowerPC core to exit its low power state. PMCI can be cleared by writing a 1 to it (writing zero has no effect). |

5.8.2.3 Power Management Controller Mask Register (PMCMR)

The power management controller mask register (PMCMR), shown in [Figure 5-58](#), controls through the PMCI bit whether the PMC interrupt request to the PowerPC core is enabled. The PMC interrupt request causes the PowerPC core to exit its low power state before any transaction on the system bus occurs. Bits 23–30 are mask bits for the defined low power wake-up events.




Figure 5-58. Power Management Controller Mask Register

¹ Not used for wake-up from D3Warm.

Table 5-74 defines the bit fields of PMCMR.

Table 5-74. PMCMR Bit Settings

| Bits | Name | Description |
|-------|--|--|
| 0–22 | — | Reserved. Write has no effect, read returns 0. |
| 23–30 | GPIO, PCI(PME), USB, eTSEC1, eTSEC2, Timer, Int1, Int2 | Wake-up event masking. 0 Mask wake-up events from Int2, Int1, Timer, eTSEC2, eTSEC1, USB, PCI, or GPIO, respectively. 1 Do not mask wake-up events |
| 31 | PMCI | Power management controller interrupt enable. 0 PMC interrupt request (PMCI) is disabled. 1 PMC interrupt request (PMCI) is enabled. |

NOTE

The user is also required to enable the PMC interrupt in the programmable interrupt controller by setting SIMSR_L[PMCI].

5.8.2.4 Power Management Controller Configuration Register 1 (PMCCR1)

The power management controller configuration register 1 (PMCCR1), shown in Figure 5-59, controls the sequencing of the device into its low power state including PME (power management event) signaling, toggling of the external power switch, and indication of current and desired power states. The power transition to D3Warm power state and the D3Warm state itself is recognized by the PMCCR[SLPEN] = 1 and PMCCR1[POWER_OFF] = 1 condition.




Figure 5-59. Power Management Controller Configuration Register 1

Table 5-75 defines the bit fields of PMCCR1.

Table 5-75. PMCCR1 Bit Settings

| Bits | Name | Description |
|------|------------|--|
| 0 | USE_STATE | Controls whether the next and current state values should be used. This typically depends on the device operating mode (PCI host or agent). For D3Warm state, this bit should be set to 0. 0 Ignore the next and current state values (host mode). 1 Use the next and current state values (agent mode). When next state does not equal current state, an interrupt will be sent to the e300 CPU. |
| 1–22 | — | Reserved, must be cleared. |
| 23 | LLPEN | SerDes low power enable bit. This bit dictates whether the SerDes is intended to be switched off or not when the chip goes into low-power mode. 0 SerDes will be switched off when the chip goes into low-power mode. 1 SerDes will not be switched off when the chip goes into low-power mode. (This would be needed when wake-up through SGMII is intended.) |
| 24 | PME_EN | PME (power management event) signaling enable. Clearing this bit typically indicates the device is acting as PCI host. Setting this bit typically indicates that the device is a PCI agent and will assert $\overline{\text{PCI_PME}}$ on wake-up. 0 Mask PME signaling when wake-up events occur. 1 Allow PME signaling when wake-up events occur. |
| 25 | ASSERT_PME | Normally $\overline{\text{PCI_PME}}$ output is asserted automatically by PMC when a defined wake-up event occurs assuming $\text{PMCCR1}[\text{PME_EN}] = 1$ and $\text{PCIPMCR1}[\text{PME_EN}] = 1$. A defined wake-up event refers to those events that are registered in bits 23–30 of PMCER. ASSERT_PME allows $\overline{\text{PCI_PME}}$ to be asserted manually, by the e300. This would be done to inform the host of a power state change when PMC does not generate $\overline{\text{PCI_PME}}$ automatically, for example waking from D1 due to CSB bus activity. ASSERT_PME is not qualified by $\text{PMCCR1}[\text{PME_EN}]$. ASSERT_PME is cleared when the $\overline{\text{PCI_PME}}$ signal is asserted. 0 $\overline{\text{PCI_PME}}$ will only be asserted automatically when a defined wake-up event occurs. 1 Assert the $\overline{\text{PCI_PME}}$ signal under software control (manually). |
| 26 | POWER_OFF | Distinguishes between D3Hot and D3Warm. If this bit is set, EXT_PWR_CTRL will be toggled in D3 causing VDD to be switched off (if implemented externally). 0 Do not use D3warm state. Always assert the EXT_PWR_CTRL output. 1 On transitions to D3 state, negate the EXT_PWR_CTRL signal to switch external power low ($\text{VDD} = 0$). On wake-up assert EXT_PWR_CTRL to switch VDD power on. |
| 27 | — | Reserved |

Table 5-75. PMCCR1 Bit Settings (continued)

| Bits | Name | Description |
|-------|------------|--|
| 28–29 | NEXT_STATE | Indicate the power state as programmed by the PCI host in PCIPMR1[Power_State]. The host will write the Power_State bits to request that the device enter a certain low power state. The host may also write the Power_State to request that the device return to D0 from some low power state. When the NEXT_STATE field is different than the CURR_STATE field, an interrupt is asserted to the e300 processor through the IPIC. This field is read-only. 00 Host's desired power state is D0 01 Host's desired power state is D1 10 Host's desired power state is D2 11 Host's desired power state is D3 only D3Hot |
| 30–31 | CURR_STATE | Indicate the current power state of the device. These bits are written by the e300 just before entering a requested low power state, or when the device has returned to the full on state (D0). Writing these bits causes the PCIPMR1[Power_State] field to be updated informing the host that the device has entered the requested power state. 00 Current power state is D0 01 Current power state is D1 10 Current power state is D2 11 Current power state is D3Hot |

5.8.2.5 Power Management Controller Configuration Register 2 (PMCCR2)

The power management controller configuration register 2 (PMCCR2), shown in [Figure 5-60](#), contains count values used for power-up and power-down timers.



Figure 5-60. Power Management Controller Configuration Register 2

Table 5-76 defines the bit fields of PMCCR2.

Table 5-76. PMCCR2 Bit Settings

| Bits | Name | Description |
|-------|-------|---|
| 0–3 | — | Reserved, must be cleared. |
| 4–15 | RCNT | <p>Reset count value. When waking up from D3Warm, power (VDD) is reapplied to a portion of the die. This value determines the duration of the reset signal applied to this logic when power is re-applied. If PMCCR1[POWER_OFF] = 0, this field has no effect.</p> <p>Reset is applied to the powered-off region upon entering D3Warm. The RCNT value is copied into a decremter that counts down once every 32,000 CSB clock cycles. When a wake-up event occurs and the PMCCR1[NEXT_STATE] is set to 00 (D0), PMC will wait for the PMC_PWR_OK signal to be asserted then begin decrementing the reset counter. When the counter reaches 0 reset is removed. In some systems the PMC_PWR_OK signal will not be provided externally and will be tied active. In this case the counter needs to include time for the VDD power supply to become stable.</p> <p>Software needs to set the RCNT value based on the PMC clock frequency, the amount of time required for the for the VDD power supply to ramp (if PMC_PWR_OK is not used), and the amount of time required for the e300 PLL to lock.</p> <p>WARNING: If the value placed in this register is too small, the reset may not assert long enough to allow the chip to function properly. The default value is larger than the time it takes for the e300 PLLs to re-lock.</p> |
| 16–19 | — | Reserved, must be cleared. |
| 20–31 | PDCNT | <p>Power-down count value. This counter establishes a minimum time for which power can be removed to the VDD supply during D3Warm. When the device enters D3Warm, the EXT_PWR_CTRL signal is negated. At this point this counter is loaded with the PDCNT value and begins to decrement, once every 32,000 CSB clock cycles. PMC will not respond to a wake-up request and assert EXT_PWR_CTRL until this counter has expired. The count value is reloaded each time the VDD power is removed. If PMCCR1[POWER_OFF] = 0 this field has no effect.</p> <p>Software needs to set this register based on the PMC clock frequency and the requirements of the power supply.</p> <p>WARNING: If the value placed in this register is too small, the power supply may cycle too quickly and the chip may not function properly.</p> |

5.8.3 Functional Description

The device has features to minimize power consumption at several levels. Dynamic power management locally minimizes power consumption when a block is idle. Software can also shut down clocks to individual blocks when they are not needed through a memory-mapped register in the clock unit (SCCR). Additionally, software running on the PowerPC core can access the core’s SPRs to put the device into doze, nap, or sleep power down states. Finally, software can access the PMCCR register to enable the device to go to low power state whenever the PowerPC core enters nap or sleep states. The PMC supports features that work in concert with the PCI power management (PM) block (PME context) to provide support for PCI power management capabilities such as asserting or responding to power management events (PMEs). These power management features are described in further detail in this section.

There are six power states in the MPC8315E:

- D0: Full-power state
- D1: Core in doze mode, core PLL running, PME signaling supported
- D2: Core in nap mode, e300 PLL running, PME signaling supported

- D3Hot: Core in sleep mode, e300 PLL not running, power applied to entire chip via VDDC and VDD power plane
 - Wake-on-LAN, GPIO, external interrupt, timer, USB, and PCI supported
 - No PME from PCI Express
- D3Warm: Core in standby mode; Ethernet, timers powered, power applied only to VDDC
 - Wake-on-LAN, GPIO (limited number), external interrupt, and timer only is supported
- D3Cold: Fully powered off

These features are not supported:

- Wake-on-PCI Express event in any of the low-power state
- Transition to and from D3Warm power state in agent mode
- Inbound PME events on PCI Express

There is flexibility on support for both wake-on-GPIO and external interrupt in D3Warm.

The MPC8315E PMC will provide for support of D3Warm state but with aggressive requirement of switching off the unused I/Os and PCI. The wake-up of the device from this state is allowed on:

- Reception of Magic Packet on eTSEC2
- IRQ1, IRQ2
- Expiry of timer (GTM)
- Activation via GPIO

The low-power state transition to and from D3Warm is controlled internally by the PMC and e300 core.

To achieve low power dissipation, a two power rail structure is defined. All the I/Os not needed for wake-up procedure from D3Warm will be switched-off. This low-power mode requires that most clocks in the MPC8315E are shut off. The exception is the logic required for wake-up. This approach includes:

- Clock gating of non-switching devices on the SOC
- Additional power plane
 - The power plane provides power to logic needed for wake-up.
 - All other logic and not needed interface including PCI switched off.
- Using Nap/Doze/Sleep modes on the e300 processor
- Clock gating at the functional block level to shut off peripheral functions that are not in use in the given system configuration, like switching off the clocks of eTSEC, GTM when only Wake-on-GPIO is required.
- Shut-off the I/Os (for example, DDR, PCI, eLBC, eTSEC1) which are not required by the wake-up logic.

There are two power planes <VDDs> and a common ground named VSS. The two power planes are defined in [Table 5-77](#).

Table 5-77. MPC8315E Power Rail Nomenclature

| Power Plane | Description |
|-------------|---|
| VDDC | This power plane supplies current in all the power states. The power plane is used to supply the logic needed for wake-up from the new low power down mode, which include ETSEC2, GPIO, PMC and other related logical blocks. |
| VDD | This supply is switched off in D3Warm state. |

Table 5-78. Software-Controller Power-Down States—Basic Description

| System Mode | Core Mode | Suggested PCI D-State | Description | Core Responds To | | DDR SDRAM str1 State | Quiesce Signal State |
|---|---|---|--|------------------|-----------|----------------------------|---|
| | | | | Snoop | Interrupt | | |
| Full On (PMCCR[SLPEN]=0) | Full On | D0 PCIPMR1[PowerState]=00 | All units operating normally | Yes | Yes | Active | Negated |
| | Doze | D1 PCIPMR1[PowerState]=01 | Core stops dispatching new instructions (core is halted), and most of the core functional units are disabled. System operates normally. | Yes | Yes | Active | Negated |
| | Nap | D2 | Core is stopped with its clocks off except to time base. System operates normally. | No | Yes | Active | Negated |
| | Sleep | PCIPMR1[PowerState]=10 | Core is stopped with its clocks off. Core clocks to all blocks (including core time base) are stopped except to the interrupt unit. System operates normally. | | | | |
| Low Power (PMCCR[SLPEN]=1) | Nap | D3 PCIPMR1[PowerState]=11 | Core operation as described above. System is in idle state, DDR SDRAM memory operates in self-refresh mode if enabled. | No | Yes | According to PMCCR [DLPEN] | Asserted |
| | Sleep | | | | | | |
| Lowest Power (PMCCR[SLEPEN]=1 && PMCCR1[POWER_OFF]=1) | Power Off (VDD switched down). Deep sleep | D3Warm (No communication via PCI. PCI switched off) | This state is an extension of the above low-power modes where power can be removed to a large portion of the device die using an external power switch on VDD. Software enters core sleep mode with PMCCR1[POWER_OFF] = 1. Wake-up is in response to defined wake-up events. | No | No | PMCCR[DLPEN] = 1 | Asserted. EXT_PWR_CTRL will switch off the power supply |

Table 5-79 describes the device power down states as mapped to OS state.

Table 5-79. Power State Mapped to OS State

| | Power State | OS State | Note |
|---|-------------|----------|---|
| 1 | D0 | NORMAL | This is active state of OS with everything powered on. |
| 2 | D1, D2 | IDLE | — |
| 3 | D3Hot | STANDBY | With the core is in sleep and the rest of the module is in low power dissipation mode with VDD on, this is "standby" mode of OS. Hence, before going to standby, OS should ensure that all the corresponding device drivers and applications are taken care of. OS should also save the configuration of various devices before going to sleep. The core, when it wakes up, would find itself exactly where it was before going to "standby". OS should use saved data to restore the configuration of all modules. |
| 4 | D3Warm | SUSPEND | This power state can be mapped to suspend state of OS. Here, all the modules including e300 are going to be powered off. DDR will be put in self refresh mode. OS should save LUT (look up tables) of MMU (memory management unit), its program counter and stack pointer to Flash. Also there should be some data stored in Flash to state that OS is going on "suspend" mode. e300 wakes up by a normal reset. OS should check in POWER_OFF whether it is supposed to initialize everything (like it does while switching on) or it should restore its state. To restore, it will have to use LUT, program counter and stack pointer from Flash to reach at the same point in DDR. From there on, OS will have to restore configuration of all the modules. |

5.8.3.1 Dynamic Power Management

Many blocks in the device can dynamically turn off clocks within the block when sections of the block are idle. This feature is always enabled and occurs automatically.

5.8.3.2 Shutting Down Unused Blocks

As described in [Section 4.5.2.3, “System Clock Control Register \(SCCR\),”](#) SCCR provides a way to shut down certain functional blocks within the device when they are not needed in a particular system. SCCR can be written by the PowerPC core or by an external master. Powering down a block in this way turns off all clocks to that block. It does not remove power. It is required that the SCCR is written to shut down a certain functional block only when that block is idle.

NOTE

Functional blocks disabled using SCCR cannot respond to configuration accesses. Any access to configuration, control, and status registers of a disabled block is a programming error.

5.8.3.3 Software-Controlled Power-Down States

PowerPC software can place the core in doze, nap, or sleep power-down states by writing to HID0 in the core, as described in detail in the section “Hardware Implementation Register 0 (HID0),” of the *e300 PowerPC Core Reference Manual*. In addition, if PMCCR[SLPEN] is set when the PowerPC core request to enter nap or sleep modes, it will also cause the system internal logic units to enter low-power mode. The basic relationships between core and system power management states is shown in [Table 5-82](#).

5.8.3.4 Software-Controlled Power Supply Switching

The device has additional low power features that allow power to be removed to a portion of the die, allowing significant additional power savings. This mode is referred to as D3Warm (described below). [Figure 5-61](#) illustrates the power segmentation provided on the device. Sequencing in and out of D3Warm will be described in the following sections.

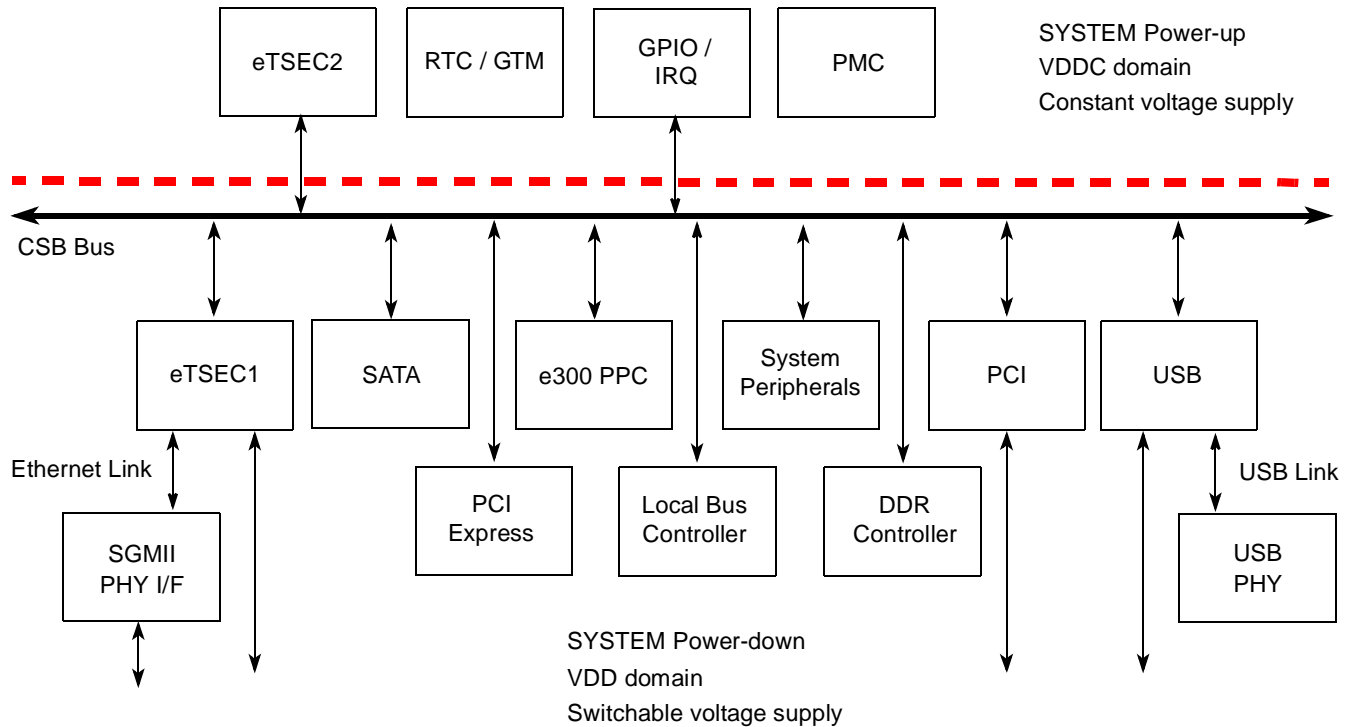


Figure 5-61. Power Segmentation in Deep Sleep Mode

5.8.3.5 Support of PCI Power Management Interface Specification

The device will support the PCI power states D0, D1, D2, D3Hot, and D3Cold as defined in the Rev. 1.2 of the PCI Power Management Interface Specification. [Table 5-80](#) defines these D_x states. The PCI power management specification defines a power management event, or PME, as the process by which a PCI agent signals a request to the host for a change in its power consumption state. When in agent mode, the device has the capability to generate PME signaling through the assertion of an external $\overline{PCI_PME}$ signal. As host, the device is able to respond to PME signaling as a wake-up event.

It is assumed that in D3Cold all power will be removed from the device, so $\overline{PCI_PME}$ signaling is not supported from D3Cold. Instead, an MPC8315E-specific D3Warm state is defined. The difference between D3Hot and D3Warm is that in D3Hot the device’s entire core region is supplied with the nominal 1-V VDD supply. In D3Warm, a portion of the core region can be powered off. This partial power-down mode allows the device to achieve significant reduction in power dissipation while still maintaining the capability to respond to wake-up events.

Table 5-80 defines the PCI power states.

Table 5-80. PCI Defined Power Management State Support

| PCI Dx Power State | Supported | How Supported |
|--------------------|-----------|--|
| D0 | Yes | Default state; full power; PME signaling supported through software control |
| D1 | Yes | e300 in Doze mode; e300 PLL running; PME signaling supported |
| D2 | Yes | e300 in Nap state; e300 PLL running; PME signaling supported |
| D3Warm | Yes | e300 in Standby mode; PME signaling is not supported |
| D3Hot | Yes | e300 in Sleep mode. In D3Hot VDD is still applied to the entire device. In this mode you have an option to power-off a portion of the device VDD. PME signaling is supported. |
| D3Cold | Yes* | *PME signaling is not supported in D3Cold. For PME signaling, use the D3Warm state, which is similar to PCI D3Hot state but with a portion of the chip powered off. According to PCI PM Specification 1.2, in D3Cold the PCI clock is removed and all devices will be reset when power is restored. The device can be placed into D3Cold state, but wake-up events cannot be recognized or generated and when power is restored, the device must go through a normal power-on reset boot sequence as it needs to be re-initialized. |

For completeness, Table 5-81 shows the PCI bus power management states (B0–B3) that the device supports. If the device is used as a PCI agent, the supported bus states are B0 and B1, since the PCI clock is running in these two states. Otherwise, as a host, the device could be configured to support all the bus states.

Table 5-81. PCI Bus Power Management State Support

| PCI Bx Bus State | PCIBus VDD | PCI Bus Clock | PCI Bus Activity | Support in PCI Agent Mode | Support in PCI Host Mode |
|------------------|------------|---------------|---|---------------------------------------|--------------------------|
| B0 (full on) | On | Yes | Any PCI transaction, function interrupt, or PME event | Yes | Yes |
| B1 | On | Yes | PME event; bus is idle | Yes | Yes ¹ |
| B2 | On | No | PME event | No ² (no bus clock) | Yes ³ |
| B3 | Off | No | PME event | No ⁴ (no bus clock or VDD) | Yes ⁵ |

¹ Requires the device to hold the PCI bus in idle state

² The device cannot process wake-up events with the PCI bus in this state and the bus must be returned to B0 state through a $\overline{\text{PORESET}}$.

³ Requires the device to turn off PCI bus clock

⁴ The device cannot process wake-up events with the PCI bus in this state and the bus must be returned to B0 state through a $\overline{\text{PORESET}}$.

⁵ Requires the device to turn off the PCI bus clock through the output clock control register (OCCR), and to turn off the PCI bus VDD through some customer-defined method (perhaps GPIO).

Below is a summary of the new *PCI Power Management Interface Specification* capabilities supported on the device:

- Generation of $\overline{\text{PCI_PME}}$ signal to an external PCI host when the device is operating in agent mode.
- Responding to $\overline{\text{PCI_PME}}$ as an input wake-up event when the device is operating as PCI host.

- Responding to other wake-up events from various sources: Ethernet Magic Packet, USB, GPIO, internal timer, external interrupt, PCI PME# (PCI_PME).
- Properly sequencing the device into and out its lowest power mode where VDD is removed to a portion of the die.
- Generating or responding to appropriate control signals relative to low-power modes: VDD switching control (EXT_PWR_CTRL) output signal and input signal indicating external power is stable (PMC_PWR_OK)

The relationship between the e300 core and the defined PMC power states is illustrated in [Table 5-82](#).

NOTE

The relationships shown below between e300 core doze, nap, and sleep modes and PCI D-states are only suggested. There is no forced hardware relationship between these states. For example, if a PCI host sets PCIPMR1[Power_State] = 01, it is still up to the software running on the e300 to put itself into doze mode.

Table 5-82. Software-Controller Power-Down States—Basic Description

| System Mode | Core Mode | Suggested PCI D- state | Description | Core Responds to | | DDR SDRAM strl state | Quiesce Signal State |
|--------------------------|-----------|------------------------------|--|------------------|-----------|----------------------|----------------------|
| | | | | Snoop | Interrupt | | |
| Full On (PMCCR[SLPEN]=0) | Full On | D0 PCIPMR1[PowerState]=00 | All units operating normally | Yes | Yes | Active | Negated |
| | Doze | D1 PCIPMR1[PowerState]=01 | Core stops dispatching new instructions (core is halted), and most of the core functional units are disabled. System operates normally. | Yes | Yes | Active | Negated |
| | Nap | D2 PCIPMR1[PowerState]=10 | Core is stopped with its clocks off except to time base. System operates normally. | No | Yes | Active | Negated |
| | Sleep | D2 PCIPMR1[PowerState]=10 | Core is stopped with its clocks off. Core clocks powered down to all blocks (including core time base) except to the interrupt unit. System operates normally. | No | Yes | Active | Negated |

Table 5-82. Software-Controller Power-Down States—Basic Description (continued)

| System Mode | Core Mode | Suggested PCI D- state | Description | Core Responds to | | DDR SDRAM strl state | Quiesce Signal State |
|--|-----------------------|--------------------------------|---|------------------|-----------|---------------------------|----------------------|
| | | | | Snoop | Interrupt | | |
| Low Power (PMCCR[SLPEN]=1) | Nap | D3 PCIPMR1[PowerState]= 11 | Core operation as described above. System is in idle state, DDR SDRAM memory operates in self-refresh mode if enabled. | No | Yes | According to PMCCR[DLPEN] | Asserted |
| | Sleep | D3 PCIPMR1[PowerState]= 11 | Core operation as described above. System is in idle state, DDR SDRAM memory operates in self-refresh mode if enabled. | No | Yes | According to PMCCR[DLPEN] | Asserted |
| Lowest Power (PMCCR[SLPEN]=1), PMCCR1[POWER_OFF] = 1 | Power Off, Deep Sleep | D3Warm PCIPMR1[PowerState]= 11 | This state is an extension of the above low-power modes where power can be removed to a portion of the device die using an external power switch. Software enters core “Sleep” mode with PMCCR1[POWER_OFF] = 1. Wake-up is in response to defined Wake-up events. | No | No | PMCCR[DLPEN] = 1 | Asserted |

5.8.3.5.1 Entering Low Power States—Core-Only Mode

Entering Doze mode is controlled only by the e300 PowerPC core itself, and does not involve the power management controller or other blocks. For a more detailed description, see [Table 7-2](#).

Entering Nap or Sleep modes occurs by writing to HID0 in the core, causing the core to make a quiesce request to the power management controller while PMCCR[SLPEN] is cleared. The core is immediately enabled to enter low power state, regardless of the system status. Note that since the core does not snoop the bus in this mode, it is the user’s responsibility to keep the cache coherent. Other device peripheral and internal units continue to operate in full-on mode while the core is in low power state in this mode.

5.8.3.5.2 Entering Low Power States—Core and System Mode

Core and system mode is achieved when the core makes a quiesce request to the power management controller after PMCCR[SLPEN] is set. To preserve cache coherency and otherwise avoid loss of system state, the core’s transition to low-power modes is coordinated with other functional blocks. The power management controller allows the core to enter power down mode only when the rest of the system is idle.

When the power management controller detects that the internal system bus is idle, and there are no outstanding transactions, it signals the internal logic units to enter low power state.

If PMCCR[DLPEN] is set, the DDR SDRAM is first set to self-refresh mode (if enabled by DDR_SDRAM_CFG[SREN] memory controller register) before the memory controller stops driving refresh commands. Self-refresh mode guarantees that the memory content will remain valid while the memory controller and its clocks are off. It stops driving clocks on MCK_n pins. Finally the DDR SDRAM memory controller enters low power state and acknowledges the power management controller.

The power management controller then signals the core and acknowledges it’s request to enter power down mode. Finally the $\overline{\text{QUIESCE}}$ output signal is asserted.

5.8.3.6 Exiting Core and System Low Power States

The device can exit low power state and return to full-on mode for one of the following reasons:

- The core internal time base unit invokes a request to exit low power state.
- The core has received an interrupt request.
- The power management controller has detected that the system is not idle and there are outstanding requests for transactions on the internal bus.

The actions taken to exit low power state depend on the mode and whether the system or only the core are in this state. The following sections describe the various scenarios.

5.8.3.6.1 Exiting Low Power States—Core-Only Mode

Exit from Doze mode is controlled only by the core itself and does not involve the power management controller or other blocks in the device. For a detailed description, see [Table 7-2](#).

Nap or Sleep modes are exited when the core has received an interrupt request, or according to the internal time base unit of the core (Nap mode only). The source of the interrupt can be an internal block or external signal. When the core returns to full-on state, it signals to the power management controller that it is ready and is immediately acknowledged to access the rest of the system.

5.8.3.6.2 Exiting Low Power States—Core and System Mode

The power management controller decides to exit low power state when it detects that the system is not idle anymore. The device may exit idle state when one of the peripheral interfaces makes a request to access the internal bus or when the core returns to full-on state, as described above, and makes a request to access the internal bus. For example, the TSEC receives an Ethernet frame, and requires to store it on the DDR SDRAM memory.

If the DDR SDRAM memory controller is in low power state (PMCCR[DxLPEN] was set when entering low power state), the power management controller initially enables the DDR SDRAM memory controller. DDR SDRAM clocks (MCK_n) are enabled and the memory controller exit self-refresh and returns to auto-refresh mode.

The power management controller then enables other internal units and interrupts the PowerPC core. When all internal units, including the core, are ready, the power management controller enables the device to return to full-on state, negate the $\overline{QUIESCE}$ output, and clear PMCCR[SLPEN]. Outstanding requests for transactions are now granted to execute on the internal bus.

NOTE

Software is required to enable PMCI interrupt by setting PMCMR[PMCI], otherwise exiting from low power state is not possible.

NOTE

It is the software's responsibility to clear PMCER[PMCI].

5.8.3.7 MPC8315E-Specific PMC Low Power States

This section will describe the MPC8315E D0–D3 low power states, including D3Warm, from a PCI power management perspective. The sequence of entering and exiting D0–D3 states will also be described.

The power management implementation assumes the following:

- The core region of the device will have a segmented power supply: a constant supply segment VDDC, and switchable supply segment VDD.
- When the device is used as a PCI agent, the PCI clock (PCI_CLK / PCI_SYNC_IN) must remain available, as this is the source of the device’s system clock. The PCI interface will reflect this requirement in the PCI’s PMC[PME_Clock] register bit.
- When the device is used as a PCI host, the PCI_CLK must also remain available.
- An external host is not required to reset the device through the PCI interface when transitioning from low-power states D1–D3Warm to the fully powered state. Reset will need to be applied if the device is put into D3Cold state (all power removed to the chip).

In the D3 Warm state, registers of the reset (Section 4.5.1, “Reset Configuration Register Descriptions”), clock (Section 4.5.2, “Clock Configuration Registers”), local memory map (Section 5.2.3.1, “Local Access Register Memory Map”), system configuration (Section 5.3.1, “System Configuration Register Memory Map”), GTM (Section 5.7.5, “GTM Memory Map/Register Definition”), PMC (Section 5.8.2, “PMC Memory Map/Register Definition”), eTSEC1, eTSEC2 (Section 19.5.2, “Detailed Memory Map”), and GPIO (Section 24.3, “Memory Map/Register Definition”) are powered ON. All the registers of these blocks remain intact during D3 Warm state.

5.8.3.7.1 Power State Transitions from an ACPI Perspective

The ACPI 3.0 specification defines the power management interface between host and agent from a board-level perspective. The device is intended to be usable as an agent in an ACPI-compliant system through the PCI interface. For the agent mode, the ACPI specification defines the power-down/wake-up sequence to be, generally:

1. The host determines the capabilities of the agent. The agent may or may not have the capability to inform the host of a wake-up event from a low-power state. This capability is determined through operating system software that identifies all PCI function power capabilities by traversing the capabilities structure in the PCI configuration space and reading the power management capabilities register.
2. The host enables the agent to generate wake-up signaling to the host.
3. The host signals the agent to transition into a low power state that it deems appropriate.
4. When the agent receives a wake-up event, it signals the host by asserting $\overline{\text{PCI_PME}}$. The agent then waits for the host to request that it transition into the active power state.
5. The host requests the agent to transition to full power state by writing the new power state into the PCI configuration space.

Table 5-83 defines the support for wake-up when the device is used as an agent. A “defined wake-up event” refers to wake-up events that are recognized by the PMC controller as agent: USB, internal timer, external interrupt, GPIO, or eTSEC. The events must be unmasked at the PMC to cause a wake-up.

Note that the device will also begin the wake-up process if the host sets PCIPMCR1[Power_State] = 00. In this case, the PMC will assert an interrupt to the e300 core, but this action will not set any of the wake-up events in the PMC event register (PMCER).

Table 5-83. MPC8315E Agent Mode Wake-Up Support

| Case | Power State | Remote Wake-up (PME) Signaling Enabled? | Wake-up Event Source | Action |
|------|-------------|---|---|--|
| 1 | D0 | Not available | Not applicable | Active state, normal system activity |
| 2 | D1 | Yes | Defined wake-up event: USB, internal timer, external interrupt, TSEC, GPIO | <p>PME signaling to the host from the D1 state is not likely to be used in an application. When a defined wake-up event occurs PMC will assert an interrupt to the e300 (if not masked) which will cause it to wake-up directly.</p> <p>To support PME signaling the following sequence should be followed: When a defined wake-up event occurs in D1 and PME signaling is enabled (PMCCR1[PME_EN] = 1), $\overline{PCI_PME}$ will be asserted to the host. Since the e300 has not been powered off, it will return to D0 (full-on) mode directly. To support PME signaling the e300 should not continue normal processing until the host has instructed the device to return to D0. The host would do this by writing PCIPMCR1[Power_State] = 00b in the PCI config space. This will cause the PMCCR1[NEXT_STATE] field to become 00b (D0) causing an interrupt to the e300. The e300 can then respond to the host that it has returned to D0 by writing PMCCR1[CURRENT_STATE] = 00, which will then be reflected in PCIPMCR1[POWER_STATE], indicating to the host that it has returned to D0. Normal processing can then continue.</p> |
| 3 | D2 | Yes | USB, internal timer, external interrupt, TSEC, GPIO | Same as D1 state (Case 2), except e300 transitions to Full On mode from Nap state. |
| 4 | D3Hot | Yes | USB, internal timer, external interrupt, TSEC, GPIO | Same as D1 state (Case 2), except e300 transitions to Full On mode from Sleep mode. |
| 5 | D3Warm | Yes | USB, internal timer, external interrupt, TSEC, GPIO | PME signaling not supported. |
| 6 | D1 | Yes | Other (for example, e300 decremter timer, e300 snoop hit, I ² C) | <p>PME signaling to the host from the D1 state is not likely to be used in an application.</p> <p>e300 is in Doze mode. If an e300 interrupt occurs, the e300 will transition to Full On mode directly. PME signalling to the host can be supported if desired (see Case 2).</p> |
| 7 | D2 | Yes | Other (for example, e300 decremter timer, I ² C) | Same as D1 state (Case 6), except e300 transitions to Full On mode from Nap state. |
| 8 | D3Hot | Yes | Other (for example, I ² C) | <p>Same as D1 (Case 6), with the following differences:</p> <ul style="list-style-type: none"> - e300 is in Sleep mode; - the e300 may not be awakened by a decremter timer interrupt. |
| 9 | D3Warm | Yes | Other | PME signaling not supported. |

Table 5-83. MPC8315E Agent Mode Wake-Up Support (continued)

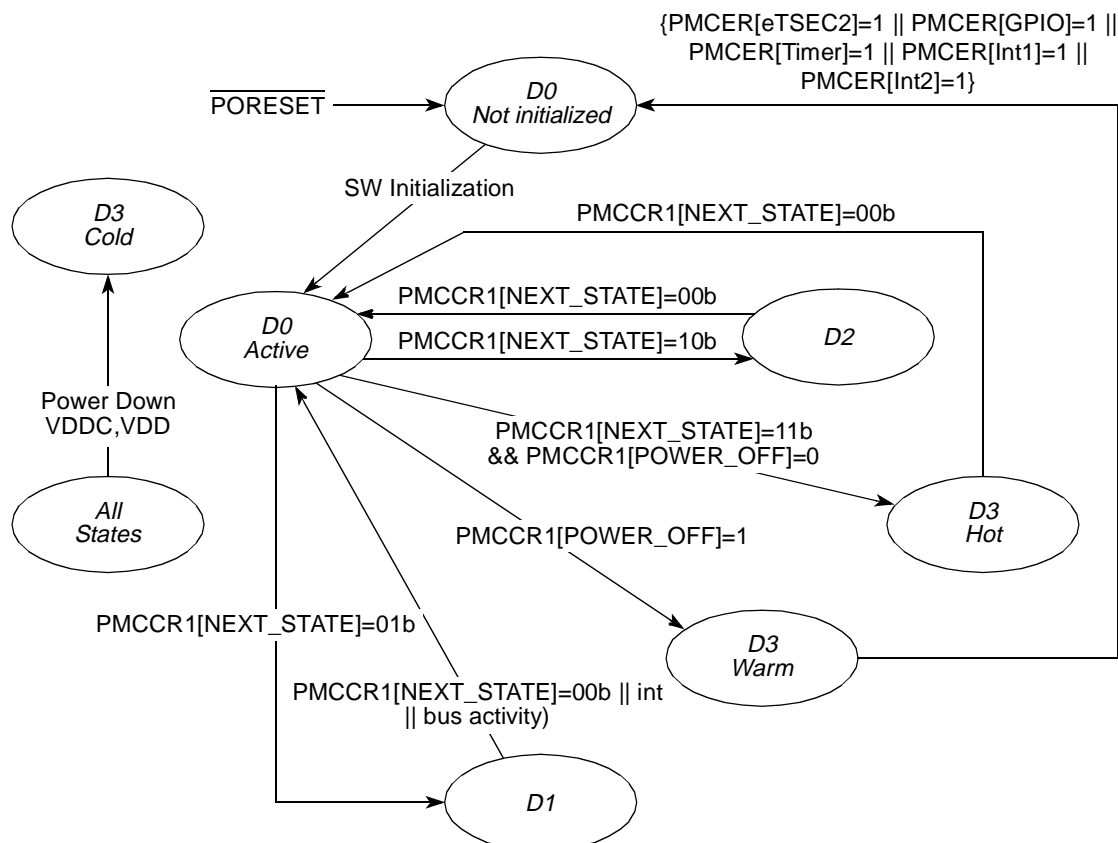
| Case | Power State | Remote Wake-up (PME) Signaling Enabled? | Wake-up Event Source | Action |
|------|-------------|---|---|---|
| 10 | D1 | No | USB, internal timer, external interrupt, TSEC, GPIO | When a defined wake-up event occurs PMC will assert an interrupt to the e300 (if not masked) which will cause it to wake-up directly (return to D0). Since PME signaling is disabled (PMCCR1[PME_EN] = 0) PCI_PME will not be asserted to the host. In this case the host will not be aware of the device power state change. To support PME signaling, see case 2. |
| 11 | D2 | No | USB, internal timer, external interrupt, TSEC, GPIO | Same as Case 10, except e300 transitions to D0 from nap mode. The state change is not reported to the host. |
| 12 | D3Hot | No | USB, internal timer, external interrupt, TSEC, GPIO | Same as Case 10, except e300 transitions to D0 from sleep mode. The state change is not reported to the host. |
| 13 | D3Warm | No | USB, internal timer, external interrupt, TSEC, GPIO | PME signaling not supported. |
| 14 | D1 | No | Other (for example, e300 decremter timer, e300 snoop hit, I ² C) | Same as Case 10. When PME signaling is disabled wake-up from Other sources is the same as wake-up from defined wake-up events. |
| 15 | D2 | No | Other (for example, e300 decremter timer, I ² C) | Same as Case 10. When PME signaling is disabled wake-up from Other sources is the same as wake-up from defined wake-up events. |
| 16 | D3Hot | No | Other (for example, I ² C) | Same as Case 10. When PME signaling is disabled wake-up from Other sources is the same as wake-up from defined wake-up events. |
| 17 | D3Warm | No | Other | Not applicable, because the e300, IPIC, and other non-essential blocks are powered off. |
| 18 | D3Cold | Not Available | Not applicable | It is assumed that in D3Cold VDDC and VDD power supplies are powered off by the host and PORESET is applied. Wake-up events or PME signaling is not supported in this case. Returning to D0 implies a POR reset and a complete initialization of the device. |

Table 5-84 defines the support for wake-up events when the device is operating as a host. A “defined wake-up event” refers to an interrupt that is recognized by the PMC in host mode: USB, GPIO, eTSEC, internal timer, external interrupt or PCI (PME input). Since the device is the host, PME signaling refers to external agents asserting the $\overline{\text{PCI_PME}}$ input to the device, which is one of the defined wake-up events.

Table 5-84. MPC8315E Host Mode Wake-Up Support

| Case | Power State | Wake-up Event Source | Action |
|------|-------------|---|--|
| 1 | D0 | Not applicable | Active state, normal system activity |
| 2 | D1 | Defined wake-up event: USB, internal timer, external interrupt, TSEC, GPIO, PCI (PCI_PME) | The wake-up event causes PMC to interrupt the e300 through the IPIC. The e300 transitions from Doze mode to Full On mode. If the wake-up source was the PCI_PME input the e300 OS may direct a power state change to the remote agent. |
| 3 | D2 | USB, internal timer, external interrupt, TSEC, GPIO, PCI (PCI_PME) | Same as Case 2, except e300 transitions to Full On mode from Nap mode. |
| 4 | D3Hot | USB, internal timer, external interrupt, TSEC, GPIO, PCI (PCI_PME) | Same as Case 2, except e300 transitions to Full On mode from Sleep mode |
| 5 | D3Warm | internal timer, external interrupt, TSEC, GPIO | Same as Case 2, except the e300 is powered off. External power will be re-applied and the e300 will be required to initialize a portion of the device (DDR, IPIC and so on) before the transition to D0 is complete. |
| 6 | D1 | Other (for example, e300 decremter timer, e300 snoop hit, I ² C) | e300 is in Doze mode; when an e300 interrupt occurs, the e300 transitions to Full On mode. |
| 7 | D2 | Other (for example, e300 decremter timer, I ² C) | e300 is in Nap mode; when an e300 interrupt occurs, the e300 transitions to Full On mode. |
| 8 | D3Hot | Other (for example, I ² C) | e300 is in Sleep mode; when an e300 interrupt occurs, the e300 transitions to Full On mode. |
| 9 | D3Warm | Other | Not applicable; e300, IPIC, and other non-essential blocks are powered off. |
| 10 | D3Cold | Not available | It is assumed that in D3Cold VDDC and VDD power supplies are powered off. Wake-up events are not recognized by the device in this case. Returning to D0 implies a POR reset and a complete initialization of the device. |

The Dx power state transitions are supported in the PMC as indicated by the following state machine, which includes both hardware and software implications. Note that when the device is agent and in one of its low-power modes, the host may decide to wake it up independent of a wake-up event or interrupt by writing the D0 state into the PCI configuration registers (PCIPMR1[Power_State] = 00. This action will cause an interrupt to the e300 (D1–D3Hot) or cause the PMC to begin the wake-up process (D3Warm).


NOTES:

- 1) Wake-up case #1: Wake-up will occur as a result of one of the following: e300 interrupt or bus activity, PMCCER[X] & PMCCMR[X] = 1, where X is one of {USB, TSEC, GPIO, PCI, interrupt, timer}, PMCCR1[NEXT_STATE] = 00. The device will wake-up directly as host or agent. PCI_PME can be generated manually to host if desired.
- 2) Wake-up case #2: As host (PMCCER[PME_EN] = 0), Wake-up will occur as a result of one of the following: PMCCER[X] & PMCCMR[X] = 1, where X is one of {USB, TSEC, GPIO, PCI, interrupt, timer}. As agent (PMCCER[PME_EN]=1) wake-up event will cause PCI_PME to be asserted; the device will wake up when the host sets PCIPMR1[Power_State] = 00, which causes PMCCR1[NEXT_STATE] = 00.
- 3) Before the transition to D3Warm, VDD power is switched off. VDDC remains constant.
- 4) Before the transition from D3Warm, VDD power is switched on.
- 5) D1, D2 & D3Hot modes are supported through use of the e300 Doze, Nap, and Sleep modes respectively. Transitions from D1–D3Hot are triggered as follows: the e300 internal time base unit invokes a request to exit low power state or e300 receives an interrupt request, including interrupt from a defined wake-up event.
- 6) Transitions to D3Cold implies power-down of VDDC and VDD supplies.
- 7) Transitions from D3Warm go to the D0-non-initialized state, meaning a portion of the chip (e300, DDR and so on) will need to be initialized. Portions of the chip which are not powered-off may not need to be initialized.

Figure 5-62. Power State Transitions Supported

5.8.3.7.2 MPC8315E Low Power Sequencing

The following sections describe the sequence of events that occur when entering or exiting the device's lowest power state (D3Warm).

PMC Power-Down When the MPC8315E is a PCI Host

In this case the PCI device driver runs on the device e300 core and directly controls the PCI interface. The e300 device driver also runs on e300 and controls the PMC module internally.

At initial power-up (POR) the host software running on the e300 is required to determine the power management capabilities of all agents and initialize their PME context. If an agent is able to generate PME its PME status must be cleared and its PME capability initialized.

Steps to power down an external PCI agent, and the device, are as follows:

1. The PCI device driver executes code to save any PCI function context that would not otherwise survive the transition to the new PM state. The OS is required to disable the I/O and memory space as well as bus mastering via the PCI Command register.
2. The PCI device driver enables external PCI agents to generate $\overline{\text{PCI_PME}}$. The driver clears PME_Status in the agent and programs the D3Hot state into the PCI agent's Power_State field.
3. In response to this PowerState field change the external PCI agent transitions itself to the new power state and then updates its own PowerState field in its PCI configuration registers.
4. The device detects the external agent's PowerState change by polling the agent's configuration registers. This process is repeated until all external agents are in their low power states.
5. Once all agents have transitioned appropriately the device may enter the D3Hot state. If the device wants to transition to the D3Warm state (power off a portion of the die) it must set the PMCCR1[POWER_OFF] bit so that the EXT_PWR_CTRL signal will toggle, turning off the VDD externally.
6. As host, the device transitions to D3Warm state by putting the e300 into deep sleep. This action causes the qreq signal to be asserted which causes PMC to sequence in to D3Warm. If the PMCCR1[POWER_OFF] bit is set the EXT_PWR_CTRL signal will transition low causing power to be removed to a portion of the die.

The power-down sequence is shown in [Figure 5-63](#) and [Figure 5-64](#).




Figure 5-63. PMC Power-Down Sequence to D3Warm




Figure 5-64. PMC Power-Down Sequence to D3Warm (continued)

Note the following:

- In host mode, the PMCCR1[PME_EN] bit should be cleared. Also, the PMCCR1[USE_STATE] bit should be cleared to indicate that the next_state and curr_state fields are ignored. The next_state and curr_state fields in that register are NOT used in host mode. The curr_state field should be left as 00.
- When powering down in host mode, the e300 should ensure that the PMCER is cleared. Any pending interrupt will prevent PMC from entering the low power state and asserting EXT_PWR_CTRL.
- Note that the next_state field in PMCCR1 is not writable by the e300. These bits are a reflection of the values in the corresponding PCIPMR1[Power_State] field.

PMC Wake-Up When the MPC8315E is a PCI Host

This sequence assumes the device is a PCI host in D3Warm state. Power is off (0v) on the VDD voltage supply. The PMC is programmed to wake on one of its wake-events: Ethernet magic packet, GPIO, internal timer, external interrupt. The PMCCR1[PME_EN] bit is reset, indicating the device is a host (PCI_PME is an input).

Steps to wake up the device are as follows.

1. A wake-up event occurs and sets the corresponding wake-up event bit in the PM CER[x] register. This event will trigger an interrupt from the PMC to the IPIC controller if not masked. This interrupt will not be serviced by the e300 until it wakes up from D3Warm. The PM CER[x] bit remains set until cleared by e300 software.
2. The wake-up event in the PM CER[x] register triggers the PMC to begin the wake-up process. It asserts EXT_PWR_CTRL to the external power supply switch. This applies power to the VDD power rail.
3. The PMC then waits for the PMC_PWR_OK signal to be asserted. PMC_PWR_OK is an indication that the external VDD supply is stable and within spec. The PMC_PWR_OK signal is an input to the PMC that can be supplied from an external source, or can be tied active internally.
4. When the PMC_PWR_OK signal asserts, the PMC's reset timer begins to count down. If the PMC_PWR_OK signal is programmed to be asserted externally but never asserts, this is a failure condition. The host must detect this failure through time-out by polling the agent's state, which will never change.
5. The PMC asserts a reset to the logic blocks in the powered-off region upon entering the low-power mode (D3Warm). This reset signal is ORed with the hard reset asserted during chip POR. When a wake-up event occurs the reset signal continues to be asserted for the duration of the reset timer count. This is to allow the e300 PLL to lock. During wake-up the reset configuration word (RCW) is not reloaded. The e300 PLL locks with previous RCW settings. The length of the PMC reset is programmable via software and must be initialized by an e300 device driver before power down. As a guideline for determining the required length of the PMC reset, the core PLL takes 100 μ s to lock (as mentioned in Table, "PLL Lock times," in MPC8315E PowerQUICC™ II Pro Processor Hardware Specifications) after PMC_PWR_OK is asserted.

Note that if the PMC_PWR_OK signal is not used, permanently asserted, the value of the reset time should be increased to account for VDD becoming stable.

6. When the PMC reset timer expires, the pmc_reset is de-asserted and the conditioning logic is removed allowing the powered on and powered off regions of the die to operate normally.
7. The e300 begins fetching instructions just as it did when initially reset (POR), starting at the memory address specified by MSRP[IP]. This address is decoded by the local address window logic (eLBC) and directed to a particular boot device.
8. The e300's initialization code checks the PMCCR1[POWER_OFF] bit to see that this is a reset from the D3Warm state, not a full POR. The e300 initializes the DDR controller so that it does not do an initialization of the DDR when coming out of self-refresh (DDR_SDRAM_CFG[B1] = 1).
9. The initialization code initializes the IPIC controller and the MSR[EE] such that the e300 can see pending interrupts including those from the PMC. After responding, the e300 clears the IPIC, PMC, and eTSEC interrupts.
10. The PMCCR1[CURR_STATE] register field is updated to reflect the new active state. This update is reflected in the PCIPMR1[Power_State] field indicating to the PCI host that the device has returned to its active state.

11. Once the e300 boots up and the IPIC has been initialized, the source of the interrupt can be detected from the PMC registers. The e300 will then reset the source of the wake-up. The source of the interrupt may be an on-chip device (eTSEC, or GPIO). An external PCI agent may need to have its power state changed and be initialized.
12. The e300 clears the on-chip IPIC, PMC interrupts, and wake-up events.

The wake-up sequence is shown in [Figure 5-65](#) and [Figure 5-66](#).




Figure 5-65. PMC Wake-Up Sequence from D3Warm




Figure 5-66. PMC Wake-Up Sequence from D3Warm (continued)

5.8.3.7.3 PMC External Power Supply Control

The following diagrams show the assumed scenarios for controlling the power supply to the device. An external means is required to switch off the VDD supply in low-power mode. A commercial power switch can be used but there is usually a switching delay associated with these devices, sometimes around 1 ms. A cheaper and faster solution is to use a FET, as shown in [Figure 5-67](#). The use of the PMCCR2[RCNT] and PMCCR2[PDCNT] timer fields will allow calibration of the device to the switching characteristics of the power switch.

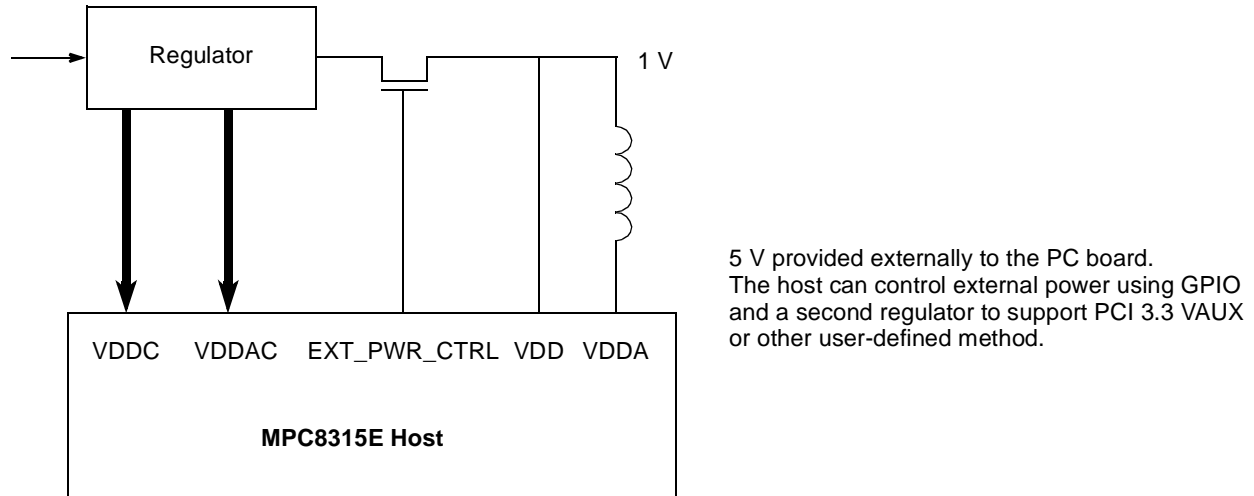


Figure 5-67. Example VDD Control of Device as Host, Using Optional PMC_PWR_OK Signal

5.8.3.7.4 Low-Power Considerations

The following should be considered in a low-power system implementation.

PME Generation

If the host wants $\overline{\text{PCI_PME}}$ to be asserted it must make sure both the PMCCR1[PME_EN] and the PCIPMR1[PME_EN] bits are set to 1.

If the external host does not want $\overline{\text{PCI_PME}}$ to be asserted it can set PCIPMR1[PME_EN] to 0. However if the PMCCR1[PME_EN] is set to 1, the PMC will assume PME signaling. If a wake-up event occurs PMC will assume $\overline{\text{PCI_PME}}$ will be asserted and that the host will respond requesting it to change its power state. If PCIPMR1[PME_EN] = 0, $\overline{\text{PCI_PME}}$ will not be asserted and the host will never know a wake-up event has occurred. In this case the host can still wake-up the device by writing a new power state in the PCIPMR1[Power_State] field, but it will never know about the device wake-up event.

The PCIPMR1[PME_EN] in the PCI configuration space will default to 0 (per spec). The PMCCR1[PME_EN] in the PMC block also defaults to 0. If there is no PCI host in the system, $\overline{\text{PCI_PME}}$ generation will most likely not be required. However if $\overline{\text{PCI_PME}}$ assertion is desired on wake-up for some reason, the e300 can assert the two PME_EN's and still generate $\overline{\text{PCI_PME}}$.

If PMCCR1[PME_EN] = 1 the implication is that the device is an agent and $\overline{\text{PCI_PME}}$ is not a wake-up event. If PMCCR1[PME_EN] = 0 the implication is that the device is the host and will wake up directly without $\overline{\text{PCI_PME}}$ signaling.

The device has the ability to assert $\overline{\text{PCI_PME}}$ under software control. The e300 can write PMCCR1[ASSERT_PME] = 1 and $\overline{\text{PCI_PME}}$ will be asserted (assuming PCIPMR1[PME_EN] = 1). This capability is provided to support PCI protocol when waking up from D1, D2 or D3Hot when the cause of the wake-up is not a defined wake-up event. In this case the e300 will wake-up or return to D0 when an

interrupt is asserted or there is CSB bus activity independent of PMC. Software can manually assert `PCI_PME` in this way to inform the host that a power state change has occurred.

If when waking up from D1–D3Hot the wake-up event is a defined wake-up event, it will be registered in the PMCCER register. When the event occurs, if the interrupt to the e300 is enabled, an interrupt will be asserted and e300 will wake-up directly. If `PMCCR1[PME_EN]` and `PCIPMR1[PME_EN]` are both set, the device will assert `PCI_PME`. In this case the software on the e300 may be written such that when it wakes up due to the PMC interrupt it waits until the host request that it go to D0 before continuing operation so that it follows PCI protocol.

If multiple wake-up events occur before the device is instructed to go to D0, multiple `PCI_PME` assertions will occur. All interrupt sources will be stored in the PMCCER until cleared. As an agent (`PMCCR1[PME_EN] = 1`), a wake-up interrupt will not be asserted to the e300 until the host has instructed the device to wake-up by writing D0 into the `PCIPMR1[Power_State]` field. As a host (`PMCCR1[PME_EN] = 0`) the interrupt will be asserted to the e300 as soon as the wake-up event occurs.

The `PMCCR1[USE_STATE]` bit is there to tell the PMC not to respond to differences between the `next_state` and `current_state` values. This is mainly to prevent PMC from waking up in Host mode if there are inadvertent changes made to the `PCIPMCR1[Power_State]` register.

In general, `EXT_PWR_CTRL` should only be toggled when going to D3Warm. This is according to general MPC8315E definitions for power levels. However, the implementation is flexible and allows for the assertion of `EXT_PWR_CTRL` even in D1 or D2 if desired.

PMC_PWR_OK Signal

`PMC_PWR_OK` is an external input indication that the VDD that was switched off in the device D3Warm mode has returned to specified levels after a wake-up event occurs. If an external power switch device is used (not a FET), it will typically provide such a signal. When a wake-up event occurs and PMC asserts the `EXT_PWR_CTRL` signal to turn on power, it will wait until the `PMC_PWR_OK` signal is asserted before it will proceed to wait for the e300 pll to lock.

PMCCR1[POWER_OFF] Bit

The `PMCCR1[POWER_OFF]` bit is set when the user wants to remove power to a portion of the die in low-power mode (D3Warm). This bit also allows boot code to distinguish between boot from power-on reset and boot from D3Warm. It is referenced relative to the `IMMRBAR` register value. If the `IMMRBAR` register is modified from its default location (the device configuration registers are moved to a different location in memory), boot software must take care to ensure it can still find the `PMCCR1[POWER_OFF]` bit. It may be necessary before entering D3Warm to change the `IMMRBAR` register back to its default location.

Debugger in Low-Power Mode

In D3Warm, the JTAG debug logic is powered up but JTAG debug accesses will not be possible. The conditioning logic in D3Warm causes JTAG inputs (TDI) to be ignored and holds TDO constant. This will cause JTAG accesses to time out.

SerDes

The SerDes PHY is kept in low-power mode when SGMII mode is not selected through the reset configuration word. The SerDes PHY can also be put in low-power mode under software control by clearing the PMCCR1[llpen] bit.

5.8.4 Initialization/Application Information

5.8.4.1 Core Disable in Low-Power Mode

If the device is required to operate with the core permanently disabled, the following steps must be taken:

1. Initialize the device with the core enable.
2. Clear PMCCR[SLPEN] and disable the core time base unit by clearing SPCR[TBEN]. See [Section 5.3.2.4, “System Priority and Configuration Register \(SPCR\),”](#) for more information.
3. The e300 core enters low power state by accessing the HID0 register.
4. Set ACR[COREDIS] in the system arbiter register with an external master (that is, PCI). This steers all device interrupts to the $\overline{\text{PCI_INTA}}$ so the core cannot receive any interrupt requests.

NOTE

Make sure that the core cannot receive any interrupt requests during the time interval between setting HID0 and setting ACR[COREDIS]. This can be achieved if the rest of the system is idle (during the initialization).

By following this flow, the e300 core remains in low power state while the rest of the system is operational, and does not get out of this state as a result of any interrupt or time-based event.

Chapter 6

Arbiter and Bus Monitor

This chapter describes operation theory of the arbiter in the device. In addition, it describes configuration, control, and status registers of the arbiter.

6.1 Arbiter Overview

The arbiter is responsible for providing coherent system bus arbitration. It tracks all address and data tenures and provides all the arbitration signals to masters and slaves. In addition, it monitors the bus and reports on errors and protocol violations.

The arbiter includes the following features:

- Supports a programmable pipeline depth (from 1 to 4)
- Supports four levels of priority for bus arbitration
- Supports repeat-request mode: number of programmable consecutive transactions from the same master (up to eight transactions)
- Supports data streaming operations
- Supports programmable address bus parking mode: disable, park to last bus owner, park to software selected master
- Claims address only, reserved and illegal transaction types, report on it and can raise maskable interrupt
- Provides timers for address tenure time out and data tenure time out detection and can issue maskable interrupt, if any timer expired
- Reports on transfer error and can issue maskable interrupt
- Can issue regular or machine check interrupt for each type of error event (programmable)

6.1.1 Coherent System Bus Overview

The coherent system bus is the central bus of the device. Any data transaction from master to slave in the device passes through the coherent system bus. The device coherent system bus supports pipelined transactions. It has independent address and data tenures. Pipeline depth determines the number of address tenures that can be started before the first data tenure is finished.

Basic burst size is equal to cache line length of the core, which is 32 bytes. Using repeat request mode enables up to eight consecutive bursts to be executed by the same master. Maximum number of consecutive transactions can be limited by programming arbiter configuration register. See [Section 6.2.1, “Arbiter Configuration Register \(ACR\)”](#) for more details.

NOTE

Write accesses to different interfaces are not guaranteed to finish in order.

6.2 Arbiter Memory Map/Register Definition

Table 6-1 shows the memory map for arbiter’s configuration, control and status registers.

Table 6-1. Arbiter Register Map


| System Arbiter—Block Base Address 0x0_0800 | | | | |
|--|--|--------|--|--------------|
| Memory Offset (Hex) | Register | Access | Reset | Section/Page |
| 0x00 | Arbiter configuration register (ACR) | R/W | All zeros/ 0x0010_0000 ¹ | 6.2.1/6-3 |
| 0x04 | Arbiter timers register (ATR) | R/W | FFFF_FFFF | 6.2.2/6-4 |
| 0x08 | Arbiter Event Enable Register (AEER) | R/W | 0x0000_007F | 6.2.3/6-5 |
| 0x0C | Arbiter event register (AER) | w1c | All zeros | 6.2.4/6-6 |
| 0x10 | Arbiter interrupt definition register (AIDR) | R/W | All zeros | 6.2.5/6-7 |
| 0x14 | Arbiter mask register (AMR) | R/W | All zeros | 6.2.6/6-8 |
| 0x18 | Arbiter event attributes register (AEATR) | R | All zeros ² | 6.2.7/6-9 |
| 0x1C | Arbiter event address register (AEADR) | R | All zeros ² | 6.2.8/6-11 |
| 0x20 | Arbiter event response register (AERR) | R/W | All zeros | 6.2.9/6-12 |

¹ Reset value is determined from the core PLL configuration of the reset configuration word. See Chapter 4, “Reset, Clocking, and Initialization,” for details.

² The registers AEATR and AEADR are affected only by the assertion of $\overline{\text{PORESET}}$

6.2.1 Arbiter Configuration Register (ACR)

Arbiter configuration register (ACR) defines the arbiter modes and parked master on the bus. [Figure 6-1](#) shows the fields of ACR.



¹ Note that the reset value of COREDIS and bits 10–11 are determined from reset configuration word. (See [Section 4.3.2](#), “Reset Configuration Words,” for more details on reset configuration word.)

Figure 6-1. Arbiter Configuration Register (ACR)

[Table 6-2](#) describes ACR fields.

Table 6-2. ACR Field Descriptions

| Bits | Name | Description |
|-------|----------|--|
| 0–6 | — | Write reserved, read = 0 |
| 7 | COREDIS | Core disable. Specifies whether CPU is disabled. When CPU is disabled, it cannot be granted on the bus by the arbiter. After reset, this bit receives its value from the reset configuration bit of COREDIS and can be configured by software. Also, if boot source is boot sequencer, COREDIS must be set to 1 at reset and the last transaction of the boot sequencer must set COREDIS to 0, if CPU enable is needed. 0 CPU enabled. 1 CPU disabled. |
| 8–9 | — | Write reserved, read = 0 |
| 10–11 | — | Reserved. Write should preserve reset value. The reset value is a function of the core PLL configuration, which is part of the reset configuration word. When the core is set to operate at 1:1 or 3:2 bus clock, these bits are set to ‘01’ during reset; otherwise, they are set to ‘00’. |
| 12 | — | Write reserved, read = 0 |
| 13–15 | PIPE_DEP | Pipeline depth (number of outstanding transactions). 000 Pipeline depth 1 (1 outstanding transaction) 001 Pipeline depth 2 (2 outstanding transactions) 010 Pipeline depth 3 (3 outstanding transactions) 011 Pipeline depth 4 (4 outstanding transactions) 1xx Reserved |
| 16 | — | Write reserved, read = 0 |

Table 6-2. ACR Field Descriptions (continued)

| Bits | Name | Description |
|-------|------------|---|
| 17–19 | PCI_RPTCNT | <p>PCI repeat count. Specifies the maximum number of consecutive transactions, that PCI master can perform, using REPEAT request mode.</p> <p>000 One consecutive transaction (REPEAT request mode disable) 001 Two consecutive transactions 010 Three consecutive transactions 011 Four consecutive transactions 100 Five consecutive transactions 101 Six consecutive transactions 110 Seven consecutive transactions 111 Eight consecutive transactions</p> |
| 20 | — | Write reserved, read = 0 |
| 21–23 | RPTCNT | <p>Repeat count. Specifies the maximum number of consecutive transactions, that any master (except PCI) can perform, using REPEAT request mode.</p> <p>000 1 consecutive transactions (REPEAT request mode disable) 001 2 consecutive transactions 010 3 consecutive transactions 011 4 consecutive transactions 100 5 consecutive transactions 101 6 consecutive transactions 110 7 consecutive transactions 111 8 consecutive transactions</p> <p>Note: It is recommended not to program this field for more than four consecutive transactions.</p> |
| 24–25 | — | Write reserved, read = 0 |
| 26–27 | APARK | <p>Address parking. Specifies arbiter bus parking mode.</p> <p>00 Park to master. Arbiter parks the address bus to the master, that is selected by numeric value of PARKM field. 01 Park to last owner. Arbiter parks the address bus to last bus owner. 10 Disable. Arbiter does not assert BG to any master, if no BR is present. 11 Reserved</p> |
| 28–31 | PARKM | <p>Parking master.</p> <p>0000 e300 core 0001 PCI, DMA 0010 TSEC1, TSEC2 0011 Encryption core 0100 SATA2, USB, TDM-DMAC 0101 PCI Express 1 0110 SATA1 0111 PCI Express 2 1000–1111 Reserved</p> |

6.2.2 Arbiter Timers Register (ATR)

The arbiter timers register (ATR) defines the arbiter address time out (ATO) and data time out (DTO) values. [Figure 6-2](#) shows the fields of ATR.



Figure 6-2. Arbiter Timers Register (ATR)

Table 6-3 describes ATR fields.

Table 6-3. ATR Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | DTO | Data time out. Specifies the time-out period for the data tenure. The granularity of this field is 128 bus clocks. The maximum value is 8355840 coherent system bus clocks. Data time_out occurs if the data tenure does not end before the specified time-out period. When DTO = n, the timeout cycle is n*128. 0000 Reserved 0001 128 clock cycles 0002 256 clock cycles 0003 384 clock cycles ... FFFF 8355840 clock cycles |
| 16–31 | ATO | Address time out. Specifies the time-out period for the address tenure. The granularity of this field is 128 bus clocks. Maximum value is 8355840 coherent system bus clocks. Address time-out occurs if the address tenure did not end before the specified time-out period. When ATO = n, the timeout cycle is n*128. 0000 Reserved 0001 128 clock cycles 0002 256 clock cycles 0003 384 clock cycles ... FFFF 8355840 clock cycles |

6.2.3 Arbiter Event Enable Register (AEER)

Arbiter event enable register (AEER) specifies, which kind of events are considered as error events. If event is defined as non error event, it also won't be reported neither in event register nor in event attributes and address registers. For transfer types, that are not defined as error events, arbiter also does not end address/data tenures. Figure 6-3 shows the fields of AEER.




Figure 6-3. Arbiter Event Enable Register (AEER)

Table 6-4 defines the bit fields of AEER.

Table 6-6. AIDR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–25 | — | Write reserved, read = 0 |
| 26 | ETEA | Transfer error. Detection of transfer error by one of the slaves interrupt definition. 0 Detection of transfer error by one of the slaves causes regular interrupt. 1 Detection of transfer error by one of the slaves causes MCP interrupt. |
| 27 | RES | Reserved transfer type. Transaction with reserved transfer type interrupt definition. 0 Transaction with reserved transfer type causes regular interrupt. 1 Transaction with reserved transfer type causes MCP interrupt. |
| 28 | ECW | External control word transfer type. Transaction with external control word transfer type interrupt definition. 0 Transaction with external control word transfer type causes regular interrupt. 1 Transaction with external control word transfer type causes MCP interrupt. |
| 29 | AO | Address only transfer type. Transaction with address only transfer type interrupt definition. 0 Transaction with address only transfer type causes regular interrupt. 1 Transaction with address only transfer type causes MCP interrupt. |
| 30 | DTO | Data time out. Data tenure time out interrupt definition. 0 Data tenure time out causes regular interrupt. 1 Data tenure time out causes MCP interrupt. |
| 31 | ATO | Address time out. Address tenure time out interrupt definition. 0 Address tenure time out causes regular interrupt. 1 Address tenure time out causes MCP interrupt. |

6.2.6 Arbiter Mask Register (AMR)

Arbiter mask register (AMR) is used to mask interrupts. Setting a mask bit enables the corresponding interrupt; clearing a bit masks it. Regular interrupts and MCP interrupts can be masked by the AMR register, except for scenarios that results in a transfer error when the master ID is 0. [Figure 6-6](#) shows the fields of AMR.

Offset 0x14

Access: User read/write




Figure 6-6. Arbiter Mask Register (AMR)

Table 6-7 describes AMR fields.

Table 6-7. AMR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–25 | — | Write reserved, read = 0 |
| 26 | ETEA | Transfer error. Detection of transfer error by one of the slaves interrupt mask bit. 0 Detection of transfer error by one of the slaves interrupt disabled. 1 Detection of transfer error by one of the slaves interrupt enabled. Note: Generation of regular or MCP interrupt is not possible for scenarios that results in a transfer error if the core master ID is 0. |
| 27 | RES | Reserved transfer type. Transaction with reserved transfer type interrupt mask bit. 0 Transaction with reserved transfer type interrupt disabled. 1 Transaction with reserved transfer type interrupt enabled. |
| 28 | ECW | External control word transfer type. Transaction with external control word transfer type interrupt mask bit. 0 Transaction with external control word transfer type interrupt disabled. 1 Transaction with external control word transfer type interrupt enabled. |
| 29 | AO | Address only transfer type. Transaction with address only transfer type interrupt mask bit. 0 Transaction with address only transfer type interrupt disabled. 1 Transaction with address only transfer type interrupt enabled. |
| 30 | DTO | Data time out. Data tenure time out interrupt mask bit. 0 Data tenure time out interrupt disabled. 1 Data tenure time out interrupt enabled. |
| 31 | ATO | Address time out. Address tenure time out interrupt mask bit. 0 Address tenure time out interrupt disabled. 1 Address tenure time out interrupt enabled. |

6.2.7 Arbiter Event Attributes Register (AEATR)

Arbiter event attributes register (AEATR) reports the type of transaction that causes error, which is specified in the event register. See [Section 6.2.4, “Arbiter Event Register \(AER\),”](#) for more information. AEATR is cleared only by power-on reset. The attributes of the first error event are stored. Note that this means that AEATR does not change its value when AER is not clear. As AEATR is not affected by soft or hard reset, software can read this register and determine the cause of the bus failure, even if the failure caused a deadlock situation. Refer to [Section 6.4.2, “Error Handling Sequence,”](#) for more information.

Figure 6-7 shows the fields of AEATR.




Figure 6-7. Arbiter Event Attributes Register (AEATR)

6.2.9 Arbiter Event Response Register (AERR)

Arbiter event response register (AERR) determines whether different error conditions cause interrupt or reset request. Setting a bit defines the corresponding error condition to cause reset request; clearing a bit defines the corresponding error condition to cause interrupt. Figure 6-9 shows the fields of AERR.

Offset 0x20

Access: User read/write




Figure 6-9. Arbiter Event Response Register (AERR)

Table 6-10 describes AERR field.

Table 6-10. AERR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–25 | — | Write reserved, read = 0 |
| 26 | ETEA | Transfer error. Detection of transfer error by one of the slaves event response. 0 Detection of transfer error by one of the slaves causes interrupt. 1 Detection of transfer error by one of the slaves causes reset request. |
| 27 | RES | Reserved transfer type. Transaction with reserved transfer type interrupt definition. 0 Transaction with reserved transfer type causes interrupt. 1 Transaction with reserved transfer type causes reset request. |
| 28 | ECW | External control word transfer type. Transaction with external control word transfer type interrupt definition. 0 Transaction with external control word transfer type causes interrupt. 1 Transaction with external control word transfer type causes reset request. |
| 29 | AO | Address only transfer type. Transaction with address only transfer type interrupt definition. 0 Transaction with address only transfer type causes interrupt. 1 Transaction with address only transfer type causes reset request. |
| 30 | DTO | Data time out. Data tenure time out interrupt definition. 0 Data tenure time out causes interrupt. 1 Data tenure time out causes reset request. |
| 31 | ATO | Address time out. Address tenure time out interrupt definition. 0 Address tenure time out causes interrupt. 1 Address tenure time out causes reset request. |

6.3 Functional Description

The following sections describe arbiter functionality: arbitration policy and bus error detection.

6.3.1 Arbitration Policy

The arbitration process involves the masters and the arbiter. Masters arbitrate on the privilege to own an address tenure. For data tenures, the arbiter uses the same order of transactions as address tenures.

Figure 6-10 shows the interface signals between the arbiter and masters that are involved in address bus arbitration.

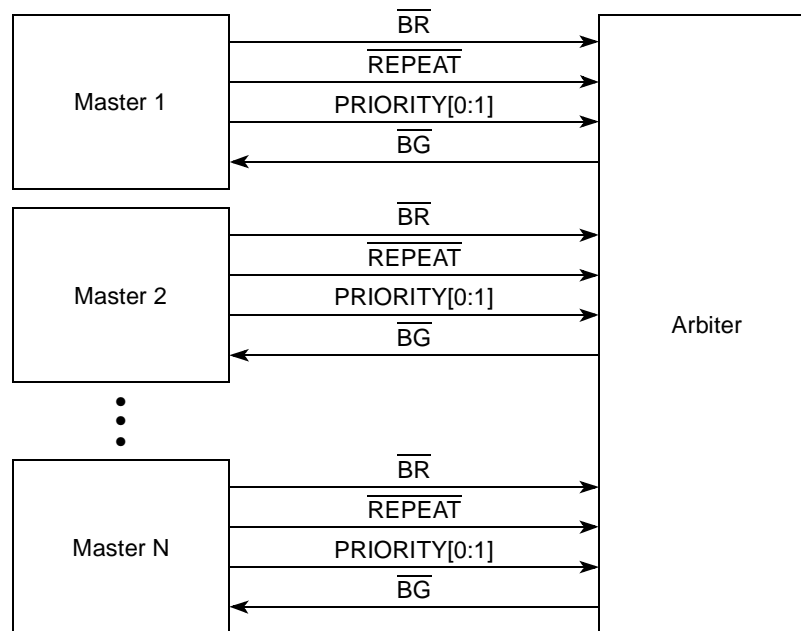


Figure 6-10. Address Bus Arbitration

A master has to acquire address bus ownership before it starts any transaction. The master asserts its own bus request signal along with the arbitration attribute signals \overline{REPEAT} & $PRIORITY[0:1]$. The arbiter later asserts the corresponding address bus grant signal to the requesting master depending on the system states and arbitration scheme. See Section 6.3.1.1, “Address Bus Arbitration with $PRIORITY[0:1]$,” for details on arbitration scheme. When address bus grant is received the master can start the address tenure.

6.3.1.1 Address Bus Arbitration with $PRIORITY[0:1]$

Whenever a master asserts its bus request to acquire address bus ownership, it can drive its $PRIORITY[0:1]$ signals to indicate request priority. The master would be served sooner because of its higher priority level. The arbiter takes this extra information into consideration in order to yield better service for a higher priority request than a lower priority request. Therefore, the arbiter operates according to the following priority based arbitration scheme:

1. For every priority level a fair arbitration scheme is used (a simple round robin scheme)
2. For every priority level other than 0, one place is reserved as a place holder for lower level arbitration rings.
3. Each master can change its priority level at any time.

Figure 6-11 shows an example of priority-based arbitration algorithm with four priority levels. In this example, if all masters request the bus continuously, the following order of bus grants occurs with the specific bandwidth:

- M6 gets 1/2 of the bus bandwidth
- M4 and M5 each gets 1/6 of the bus bandwidth

- M0 and M3 each gets 1/18 of the bus bandwidth
- M1 and M2 each gets 1/36 of the bus bandwidth

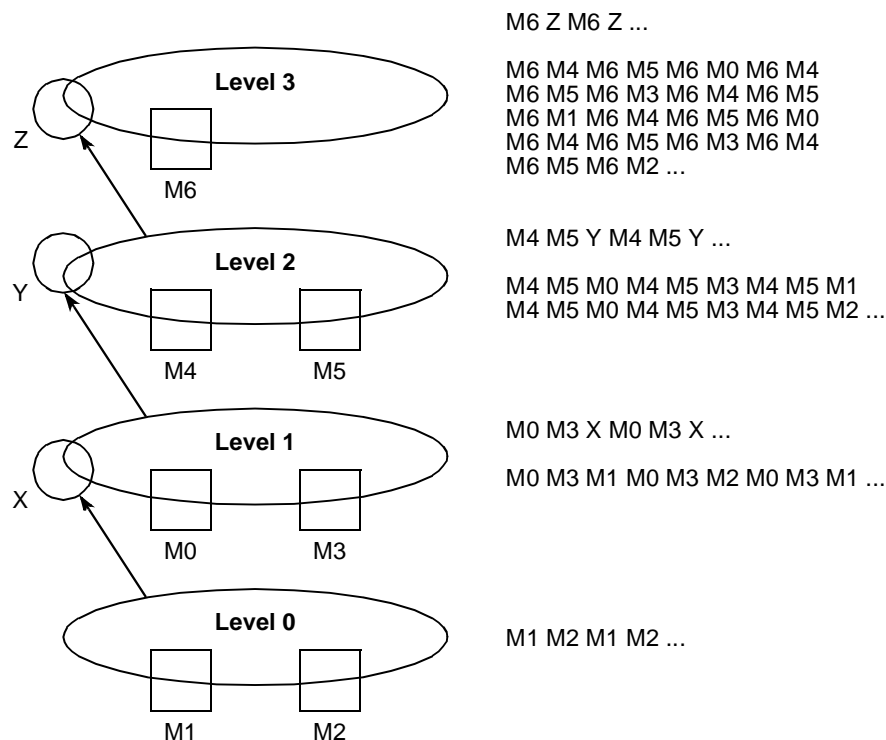


Figure 6-11. An Example of Priority-Based Arbitration Algorithm

NOTE

See each bus master’s chapter and [Section 5.3.2.4, “System Priority and Configuration Register \(SPCR\),”](#) for more details about priority programming.

6.3.1.2 Address Bus Arbitration with REPEAT

When a master owns the address bus and wants to perform another transaction, it can assert bus request along with $\overline{\text{REPEAT}}$, to make a repeat request to the arbiter. Consequently, the arbiter asserts bus grant to the same master if the current address tenure is not being $\overline{\text{ARTRY}}$ ed. This happens regardless of the priority level of bus request from other masters. In another word, “repeat request” overrides the priority scheme.

Even though repeat request can improve the page hit ratio and the overall memory bandwidth efficiency, it can increase the worst case latency of individual master. Therefore, the arbiter has programmable counter to limit the maximum number of consecutive transactions that are performed by masters. Whenever the counter expires, arbiter ignores the $\overline{\text{REPEAT}}$ signal and falls back to the regular arbitration scheme. PCI master has a dedicated repeat counter as it might need more repeated transactions before accepting read requests. PCI ordering rules require that the PCI bridge should empty all queued write operations before any new read operation can begin. See Section 3.2.5, “Transaction Ordering and Posting,” of the *PCI Local Bus Specifications Rev 2.2*, for more information.

See [Section 6.2.1, “Arbiter Configuration Register \(ACR\),”](#) for more details about programming ACR[RPTCNT]PCI.

6.3.1.3 Address Bus Arbitration After $\overline{\text{ARTRY}}$

The $\overline{\text{ARTRY}}$ protocol is used primarily by the CPU to interrupt a transaction that hits to a modified line in its D-cache, so that it can maintain data coherency by performing the snoop copyback. When CPU asserts $\overline{\text{ARTRY}}$, the bus is immediately granted to the CPU to perform snoop copyback. After the completion of snoop copyback, the arbiter grants the bus to the most ahead master among those masters which have an active bus request signal at that time, which may or may not be the same master that had its transaction $\overline{\text{ARTRY}}$ ed. Only when a transaction address phase is completed with no $\overline{\text{ARTRY}}$ (and no repeat conditions), the master moves to the end of the line.

6.3.1.4 Address Bus Parking

The arbiter supports address bus parking. This feature implies that when no master is requesting the bus (all bus requests are negated), the arbiter can choose to park the address bus (or assert the address bus grant) to a master. The parked master can skip the bus request and assume the bus mastership directly. This reduces the access latency for parked master.

See [Section 6.2.1, “Arbiter Configuration Register \(ACR\),”](#) for more details about ACR[APARK] and ACR[PARKM].

6.3.1.5 Data Bus Arbitration

For every committed address tenure a data tenure is required to complete the transaction.

In the device system, the arbiter controls the issuing of data bus grants to a master and a slave, which are involved in a data tenure of a previously performed address tenure.

6.3.2 Bus Error Detection

The arbiter is responsible for tracking the following cases on the bus:

- Address time out
- Data time out
- Transfer error
- Address only transaction type
- Reserved transaction type
- Illegal (**eciwx/ecowx**) transaction type

6.3.2.1 Address Time Out

Address time out occurs, if the address tenure was not ended before the specified time-out period (programmed by ATR[ATO]). In this case, the arbiter performs as follows:

1. Ends the address tenure.

2. Starts data tenure and ends it by asserting transfer error.
3. Reports on the event to AER[ATO].
4. Issues reset request, MCP or regular interrupt according to AERR[ATO] and AIDR[ATO], if enabled by AMR[ATO].
5. Updates transaction attributes and address of AEATR and AEADR for the first error event.

6.3.2.2 Data Time Out

Data time out occurs, if the data tenure was not ended before the specified time-out period (programmed by ATR[DTO]). In this case, the arbiter performs as follows:

1. Ends the data tenure by asserting transfer error.
2. Reports on this event in AER[DTO].
3. Issues reset request, MCP or regular interrupt according to AERR[DTO] and AIDR[DTO], if enabled by AMR[DTO].
4. Updates transaction attributes and address of AEATR and AEADR for the first error event.

6.3.2.3 Transfer Error

The arbiter tracks the transfer error asserted by one of the slaves. In this case, the arbiter performs as follows:

1. Reports on the event to AER[ETEA].
2. Issues reset request, MCP or regular interrupt according to AERR[ETEA] and AIDR[ETEA] if enabled by AMR[ETEA].
3. Updates transaction attributes and address of AEATR and AEADR for the first error event.

6.3.2.4 Address Only Transaction Type

Table 6-11 shows transaction types, which are defined as address only:

Table 6-11. Address Only Transaction Type Encoding

| ttype[0:4] | Bus Commands |
|------------|-----------------------|
| 00000 | Clean block |
| 00100 | Flush block |
| 01000 | Sync |
| 01100 | Kill block |
| 10000 | eieio |
| 11000 | TLB Invalidate |
| 00001 | lwarx reservation set |
| 01001 | tlbsync |
| 01101 | icbi |

The arbiter allows address-only (AO) transactions on the bus and the G2 core has ability to issue address-only (AO) transactions (see HID0 [ABE] in the *G2 PowerPC Core Reference Manual*, Rev 1). As

there is no advantage in using AO transaction in this system, the bus monitor allows the detection of AO transactions and treats them as an error.

The transaction with an address only transfer type, the arbiter performs as follows:

1. Ends the address tenure by asserting $\overline{\text{AACK}}$.
2. Reports on the event to AER[AO].
3. Issues reset request, MCP or regular interrupt according to AERR[AO] and AIDR[AO] if enabled by AMR[AO].
4. Updates transaction attributes and address of AEATR and AEADR for the first error event.

6.3.2.5 Reserved Transaction Type

Table 6-12 shows transaction types defined as reserved.

Table 6-12. Reserved Transaction Type Encoding

| ttype[0:4] | Bus Commands |
|------------|-----------------------|
| 00101 | Reserved |
| 1xx01 | Reserved for customer |
| 10110 | Reserved |
| 00011 | Reserved |
| 00111 | Reserved |
| 01111 | Reserved |
| 1xx11 | Reserved for customer |

The transaction with a reserved transfer type, the arbiter performs as follows:

1. Ends the address tenure by asserting $\overline{\text{AACK}}$.
2. Reports on the event to AER[RES].
3. Issues reset request, MCP or regular interrupt according to AERR[RES] and AIDR[RES], if enabled by AMR[RES].
4. Updates transaction attributes and address of AEATR and AEADR for the first error event.

6.3.2.6 Illegal (eciwx/ecowx) Transaction Type

Table 6-13 shows transaction types defined as illegal.

Table 6-13. Illegal Transaction Type Encoding

| ttype[0:4] | Bus command |
|------------|--|
| 10100 | External control word write (ecowx) |
| 11100 | External control word read (eciwx) |

The transaction with an illegal (eciwx, ecowx) transfer type, the arbiter performs as follows:

1. Ends the address tenure by asserting $\overline{\text{AACK}}$.
2. Starts data tenure and ends data tenure by asserting $\overline{\text{TEA}}$.
3. Reports on the event in AER[ECW].

4. Issues reset request, MCP or regular interrupt according to AERR[ECW] and AIDR[ECW], if enabled by AMR[ECW].
5. Updates transaction attributes and address of AEATR and AEADR for the first error event.

See Section 6.2.4, “Arbiter Event Register (AER),” Section 6.2.5, “Arbiter Interrupt Definition Register (AIDR),” Section 6.2.6, “Arbiter Mask Register (AMR),” Section 6.2.7, “Arbiter Event Attributes Register (AEATR),” Section 6.2.8, “Arbiter Event Address Register (AEADR),” and Section 6.2.9, “Arbiter Event Response Register (AERR),” for more information.

6.4 Initialization/Applications Information

The following sections describe the initialization and error handling sequences for the arbiter.

6.4.1 Initialization Sequence

The following initialization sequence is recommended:

1. Write to ACR to configure pipeline depth, address bus parking mode, global maximum repeat count, PCI maximum.
2. Write to AERR defines whether different error events cause a reset request or an interrupt.
3. Write to AIDR defines the kind of interrupt (regular or MCP) caused by each error event. Note that this is necessary only if interrupts are enabled and AERR defines error events to cause interrupt.
4. Write to AMR to enable interrupts.
5. Write to ATR to set the ATO and DTO timers. Note that this is only necessary if the required timers are less than the maximum value (which is default).

6.4.2 Error Handling Sequence

The following error handling sequence is recommended:

1. Read to AER to find out about the error that occurred in the system. Also, read the values of AEATR and AEADR to check on the first error event in the system.
2. If those registers are not accessible because of a bus deadlock situation, reset the chip and read the values of the AEATR and AEADR registers to check on the event that causes this problem to the system. Use HRESET to reset the chip to guarantee that the information stored in AEATR and AEADR is not lost.
3. Clear all the previous events by writing ones to the AER. This register is also cleared after reset.

Chapter 7

e300 Processor Core Overview

This chapter provides an overview of features for the embedded microprocessors in the e300 core family, which are PowerPC microprocessors built on Power Architecture technology. Throughout this chapter, the terms ‘e300 core’, ‘core’, and ‘processor’ are used interchangeably. The term, ‘e300c3’ is used when describing an implementation-specific feature or when a difference exists between different configurations. The term ‘e300’ is used when describing a feature that pertains to the family of e300 processors. The MPC8315E uses an e300c3 core.

7.1 e300c3 Overview

This section describes the details of the e300 core, provides a block diagram showing the major functional units, and briefly describes how these units interact. All differences between the e300 and previous PowerPC implementations derived from the MPC603e processor are noted. For additional information, refer to the *e300 PowerPC Core Family Reference Manual*.

The e300 core is a low-power implementation of this microprocessor family of reduced instruction set computing (RISC) microprocessors. The core implements the 32-bit portion of the PowerPC architecture, which defines 32-bit effective addresses, integer data types of 8, 16, and 32 bits, and floating-point data types of 32 and 64 bits.

The core is a superscalar processor that can issue and retire as many as three instructions per clock cycle. Instructions can execute out of program order for increased performance; however, the core makes completion appear sequential.

The e300 core integrates independent execution units including: an integer unit (IU) a floating-point unit (FPU), a branch processing unit (BPU), a load/store unit (LSU), and a system register unit (SRU). The e300c3 integrate an additional integer unit for a total of two IUs. The ability to execute instructions in parallel and the use of simple instructions with rapid execution times yield high efficiency and throughput for e300-core-based systems. Most integer instructions execute in one clock cycle. The additional IUs along with enhanced multipliers in the e300c3 improve multiply instructions to a maximum two-cycle latency, a significant improvement from previous processors. In the e300c3 core, the FPU is pipelined so a single-precision multiply-add instruction can be issued and completed every clock cycle. The e300c3 core provide hardware support for all single- and double-precision floating-point operations for most value representations and all rounding modes.

[Figure 7-1](#) shows a block diagram of the e300c3 core. Note that the e300c3 supports floating-point operations and includes two integer units.




Figure 7-1. e300c3 Core Block Diagram

The e300c3 includes 16-Kbyte, four way set-associative instruction and data caches. The MMUs contain 64-entry, two-way, set-associative, data and instruction translation lookaside buffers (DTLB and ITLB) that provide support for demand-paged, virtual-memory, address translation, and variable-sized block translation. The TLBs use a least recently used (LRU) replacement algorithm and the caches use a pseudo least recently used algorithm (PLRU).

The core also supports block address translation through the use of two independent instruction and data block address translation (IBAT and DBAT) arrays, each containing eight pairs of BATs, an increase from four pairs of each type of BATs in the G2 core. This increase provides more flexibility in protecting accesses and providing translation on a segment, block, or page basis for memory accesses and I/O accesses. Effective addresses are compared simultaneously with all eight entries in the BAT array during block translation. In accordance with the PowerPC architecture, if an effective address hits in both the TLB and BAT array, the BAT translation takes priority.

As part of the coherent system bus (CSB), the e300 core has a 64-bit data bus and a 32-bit address bus. During normal operation, the e300 core provides a three-state (modified, exclusive, and invalid) coherency protocol which is a compatible subset of a four-state (modified/exclusive/shared/invalid) MESI protocol. However, the e300 data cache contains a programmable MESI extension that supports the shared cache coherency state (similar to other PowerPC processors). Both protocols operate coherently in systems that contain four-state caches. Although MESI is supported by the e300 core, it is not implemented on the MPC8315E. The core also supports single-beat and burst data transfers for memory accesses and supports memory-mapped I/O operations.

The true little-endian mode is another enhanced capability of the e300 core. Unlike the PowerPC little-endian mode (which manipulates only the address bits), no longer supported on the e300, the true little-endian mode actually operates on true little-endian instructions and data from memory.

The critical interrupt is an additional interrupt in the e300 core and has higher priority order than the system management interrupt. Also, debug features are improved in the e300. Additional SPRG interrupt handling registers are provided for enhancing flexibility for the operating system.

The e300c3 include a performance monitor facility that provides the ability to monitor and count predefined events such as core clocks, misses in the instruction cache, data cache, or L2 cache, types of instructions dispatched, mispredicted branches, and other occurrences. The count of such events (which may be an approximation) can be used to trigger the performance monitor interrupt. [Section 7.1.7.5, “Core Performance Monitor,”](#) describes the operation of the performance monitor diagnostic tool.

7.1.1 e300c3 Features

This section describes the major features of the e300 core:

- High-performance, superscalar microprocessor core
 - As many as three instructions issued and retired per clock (two instructions plus one branch instruction)
 - As many as five instructions in execution per clock
 - Single-cycle execution for most instructions
 - Pipelined floating-point unit (FPU) for all single- and double-precision operations

- Independent execution units and two register files
 - Branch processing unit (BPU) featuring static branch prediction
 - Two 32-bit integer units (IU) in the e300c3.
 - FPU based on the IEEE Std 754™ for both single- and double-precision operations
 - Load/store unit (LSU) for data transfer between data-cache and general-purpose registers (GPRs) and floating-point registers (FPRs)
 - System register unit (SRU) that executes condition register (CR), special-purpose register (SPR), and integer add/compare instructions. Add/compare instructions are also executed in the IUs.
 - Thirty-two 32-bit GPRs for integer operands
 - Thirty-two 64-bit FPRs for single- or double-precision operands
- High instruction and data throughput
 - Zero-cycle branch capability (branch folding)
 - Programmable static branch prediction on unresolved conditional branches
 - Two integer units with enhanced multipliers in the e300c3 for increased integer instruction throughput and a maximum two-cycle latency for multiply instructions
 - Instruction fetch unit capable of fetching two instructions per clock from the instruction cache
 - A six-entry instruction queue (IQ) that provides lookahead capability
 - Independent pipelines with feed-forwarding that reduces data dependencies in hardware
 - 16-Kbyte, four-way set-associative instruction and data caches on the e300c3.
 - Cache write-back or write-through operation programmable on a per-page or per-block basis
 - Features for instruction and data cache locking and protection
 - BPU that performs CR lookahead operations
 - Address translation facilities for 4-Kbyte page size, variable block size, and 256-Mbyte segment size
 - A 64-entry, two-way, set-associative ITLB and DTLB
 - Eight-entry data and instruction BAT arrays providing 128-Kbyte to 256-Mbyte blocks
 - Software table search operations and updates supported through fast trap mechanism
 - 52-bit virtual address; 32-bit physical address
- Facilities for enhanced system performance
 - A 64-bit split-transaction internal data bus interface to the coherent system bus (CSB) with burst transfers
 - Support for one-level address pipelining on the CSB interface
 - True little-endian mode for compatibility with other true little-endian devices
 - Critical interrupt support
 - Hardware support for misaligned little-endian accesses
 - Configurable processor bus frequency multipliers as defined in the *MPC8315E Integrated Processor Hardware Specifications*

- Integrated power management
 - Internal processor/bus clock multiplier ratios
 - Three power-saving modes: doze, nap, and sleep
 - Automatic dynamic power reduction when internal functional units are idle
- In-system testability and debugging features through JTAG boundary-scan capability

Features specific to the e300 core not present on the G2 processors follow:

- Enhancements to the register set
 - The e300 core has one more HID0 bit than the G2:
 - The enable cache parity checking (ECPE) bit, HID0[1], gives the e300 core the ability to enable the taking of a machine check interrupt based on the detection of a cache parity error
- Enhancements to cache implementation
 - 16-Kbyte, four-way set-associative instruction and data caches on the e300c3.
 - Full parity checking is performed on both instruction and data cache memory arrays
 - Lockable L1 instruction and data caches—entire cache or on a per-way basis up to 3 of 4 ways on the e300c3
 - New **icbt** instruction supports initialization of instruction cache
 - Data cache supports four-state MESI coherency protocol (not implemented on MPC8315E)
 - The instruction cache is blocked only until the critical load completes (hit under reloads allowed)
 - Instruction cancel mechanism improves utilization of instruction cache by supporting hits-under-cancels and misses-under-cancels.
 - The critical double word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays.
 - Data cache queue sharing makes cast-outs and snoop pushes more efficient
 - Provides for an optional data cache operation broadcast feature (enabled by HID0[ABE]) that allows for coherent system management. All of the data cache control instructions, except **dcbz** (**dcbi**, **dcbf**, and **dcbst**) require that HID0[ABE] be enabled to broadcast.
 - Instruction fetch burst feature allows all instruction fetches from caching-inhibited space to be performed on the bus as burst transactions
- Interrupts
 - The e300 core offers hardware support for misaligned little-endian accesses. Little-endian load/store accesses that are not on a word boundary, except for strings and multiples, generate interrupts under the same circumstances as big-endian accesses.
 - The e300 core supports true little-endian mode to minimize the impact on software porting from true little-endian systems.
 - An input interrupt signal, \overline{cint} , is provided to trigger the critical interrupt exception on the e300 core. The pm_event_in input signal can be used by the performance monitor counters to trigger an interrupt upon overflow on the e300c3.
- Bus clock—PLL configuration signals include seven signals for settings and control: *pll_cfg[0:6]*.

- Debug features
 - Breakpoint status recorded in DBCR and IBCR control registers
 - Two signals for the debug interface: *stopped* and *ext_halt*
 - Performance monitor registers for system analysis in the e300c3

Figure 7-1 provides a block diagram of the e300 core that shows how the execution units—IU, FPU, BPU, LSU, and SRU—operate independently and in parallel. It should be noted that this is a conceptual diagram and does not attempt to show how these features are physically implemented on the device.

The e300 core provides address translation and protection facilities, including an ITLB, DTLB, and instruction and data BAT arrays. Instruction fetching and issuing are handled in the instruction unit. Translation of addresses for cache or external memory accesses are handled by the MMUs. Both units are discussed in more detail in Section 7.1.2, “Instruction Unit,” and Section 7.1.5.1, “Memory Management Units (MMUs).”

7.1.2 Instruction Unit

As shown in Figure 7-1, the e300 core instruction unit, containing a fetch unit, instruction queue, dispatch unit, and BPU, provides centralized control of instruction flow to the execution units. The instruction unit determines the address of the next instruction to be fetched based on information from the sequential fetcher and from the BPU.

The instruction unit fetches the instructions from the instruction cache into the instruction queue. The BPU receives branch instructions from the fetcher and uses static branch prediction to allow fetching from a predicted instruction stream while a conditional branch is evaluated. The BPU folds out for unconditional branch instructions and conditional branch instructions unaffected by instructions in the execution pipeline.

Instructions issued beyond a predicted branch cannot complete execution until the branch is resolved, preserving the programming model of sequential execution. If any of these are branch instructions, they are decoded but not issued. Instructions to be executed by the FPU, IU, LSU, and SRU are issued and allowed to progress up to the register write-back stage. Write-back is allowed when a correctly predicted branch is resolved, and execution continues along the predicted path.

If branch prediction is incorrect, the instruction unit flushes all predicted path instructions, and instructions are issued from the correct path.

7.1.2.1 Instruction Queue and Dispatch Unit

The instruction queue (IQ), shown in Figure 7-1, holds as many as six instructions and loads up to two instructions from the instruction unit during a single cycle. The instruction fetch unit continuously loads as many instructions as space in the IQ allows. Instructions are dispatched to their respective execution units from the dispatch unit at a maximum rate of two instructions per cycle. Dispatching is facilitated to the IUs, FPU, LSU, and SRU by the provision of a reservation station at each unit. The dispatch unit performs source and destination register dependency checking, determines dispatch serializations, and inhibits subsequent instruction dispatching as required.

For a more detailed overview of instruction dispatch, see Section 7.4.6, “Instruction Timing.”

7.1.2.2 Branch Processing Unit (BPU)

The BPU receives branch instructions from the fetch unit and performs CR lookahead operations on conditional branches to resolve them early, achieving the effect of a zero-cycle branch in many cases.

The BPU uses a bit in the instruction encoding to predict the direction of the conditional branch. Therefore, when an unresolved conditional branch instruction is encountered, the core fetches instructions from the predicted target stream until the conditional branch is resolved.

The BPU contains an adder to compute branch target addresses and three user-control registers: the link register (LR), the count register (CTR), and the conditional register (CR). The BPU calculates the return pointer for sub-routine calls and saves it into the LR for certain types of branch instructions. The LR also contains the branch target address for the Branch Conditional to Link Register (**bclrx**) instruction. The CTR contains the branch target address for the Branch Conditional to Count Register (**bcctrx**) instruction. The contents of the LR and CTR can be copied to or from any GPR. Because the BPU uses dedicated registers rather than GPRs or FPRs, execution of branch instructions is largely independent from execution of integer and floating-point instructions.

7.1.3 Independent Execution Units

The PowerPC architecture's support for independent execution units allows implementation of processors with out-of-order instruction execution. For example, because branch instructions do not depend on GPRs or FPRs, branches can often be resolved early, eliminating stalls caused by taken branches.

The four other execution units and the completion unit are described in the following sections.

7.1.3.1 Integer Unit (IU)

The IU executes all integer instructions. The IU executes one integer instruction at a time, performing computations with its arithmetic logic unit (ALU), multiplier, divider, and XER register. Most integer instructions are single-cycle instructions. The 32 GPRs hold integer operands. Stalls due to contention for GPRs are minimized by the automatic allocation of rename registers. The core writes the contents of the rename registers to the appropriate GPR when integer instructions are retired by the completion unit. The e300c3 provides two integer units for greater integer instruction throughput along with enhanced multipliers in each IU for faster multiply instruction execution.

7.1.3.2 Floating-Point Unit (FPU)

The FPU contains a single-precision multiply-add array and the floating-point status and control register (FPSCR). The multiply-add array allows the core to efficiently implement multiply and multiply-add operations. The FPU is pipelined so that single- and double-precision instructions can be issued back-to-back. The 32 FPRs are provided to support floating-point operations. Stalls due to contention for FPRs are minimized by the automatic allocation of rename registers. The core writes the contents of the rename registers to the appropriate FPR when floating-point instructions are retired by the completion unit.

The e300c3 core supports all floating-point data types based on the IEEE 754 standard (normalized, denormalized, NaN, zero, and infinity) in hardware, eliminating the latency incurred by software interrupt routines.

7.1.3.3 Load/Store Unit (LSU)

The LSU executes all load and store instructions and provides the data transfer interface between the GPRs, FPRs, and the cache/memory subsystem. The LSU calculates effective addresses, performs data alignment, and provides sequencing for load/store string and multiple instructions.

Load and store instructions are issued and executed in program order; however, the memory accesses can occur out of order. Synchronizing instructions are provided to enforce strict ordering.

Cacheable loads, when free of data bus dependencies, can execute out of order with a maximum throughput of one per cycle and with a two-cycle total latency. Data returned from the cache is held in a rename register until the completion logic commits the value to a GPR or FPR. Stores cannot be executed in a predicted manner and are held in the store queue until the completion logic signals that the store operation is to be completed to memory. The core executes store instructions with a maximum throughput of one per cycle and with a three-cycle total latency. The time required to perform the actual load or store depends on whether the operation involves the cache, system memory, or an I/O device.

7.1.3.4 System Register Unit (SRU)

The SRU executes various system-level instructions, including condition register logical operations and move to/from special-purpose register instructions. It also executes integer add/compare instructions. In order to maintain system state, most instructions executed by the SRU are completion-serialized; that is, the instruction is held for execution in the SRU until all prior instructions issued have completed. Results from completion-serialized instructions executed by the SRU are not available or forwarded for subsequent instructions until they complete.

7.1.4 Completion Unit

The completion unit tracks instructions in program order from dispatch through execution and then completes. Completing an instruction commits the core to any architectural register changes caused by that instruction. In-order completion ensures the correct architectural state when the core must recover from a mispredicted branch or an interrupt.

Instruction state and other information required for completion is kept in a five-entry FIFO completion queue. A single completion queue entry is allocated for each instruction once it enters the execution unit from the dispatch unit. An available completion queue entry is a required resource for dispatch; if no completion entry is available, dispatch stalls. A maximum of two instructions per cycle are completed in order from the queue.

7.1.5 Memory Subsystem Support

The core provides separate instruction and data caches and MMUs. The core also provides an efficient processor bus interface to facilitate access to main memory and other bus subsystems. The memory subsystem support functions are described in the following sections.

7.1.5.1 Memory Management Units (MMUs)

The core MMUs support up to 4 Petabytes (2^{52}) of virtual memory and 4 Gigabytes (2^{32}) of physical memory (referred to as real memory in the architecture specification) for instruction and data. The MMUs also control access privileges for these spaces on block and page granularities. Referenced and changed status is maintained by the processor for each page to assist implementation of a demand-paged virtual memory system. Note that software assistance is required for the device to maintain reference and changed status. A key bit is implemented to provide information about memory protection violations prior to page table search operations.

The LSU calculates effective addresses for data loads and stores, performs data alignment to and from cache memory, and provides the sequencing for load and store string and multiple word instructions. The instruction unit calculates effective addresses for instruction fetching.

After an EA is generated, its higher-order bits are translated by the appropriate MMU into physical address bits. The lower-order EA bits are the same on the physical address which are directed to the on-chip cache and formed the index into a four-way set-associative tag array. After translating the address, the MMU passes the higher-order physical address bits to the cache and the cache lookup completes. For caching-inhibited accesses or accesses that miss in the cache, the untranslated lower-order address bits are concatenated with the translated higher-order address bits; the resulting 32-bit physical address is then used by the memory unit and the core interface to access external memory.

The MMU also directs the address translation and enforces the protection hierarchy programmed by the operating system in relation to the supervisor/user privilege level of the access and in relation to whether the access is a load or store.

For instruction fetches, the IMMU looks for the address in the ITLB and in the IBAT array. If an address hits both, the IBAT array translation is used. Data accesses cause a lookup in the DTLB and DBAT array. In most cases, the translation is in a TLB and the physical address bits are available to the on-chip cache.

The e300 core implements four more IBAT and four more DBAT entries than the G2.

When the EA misses in the TLBs, the core provides hardware assistance for software to perform a search of the translation tables in memory. The hardware assist consists of the following features:

- Automatic storage of the missed effective address in IMISS and DMISS
- Automatic generation of the primary and secondary hashed real addresses of the page-table entry group (PTEG), which are readable from the HASH1 and HASH2 register locations.
The HASH data is generated from the contents of the IMISS or DMISS register. The register that is selected depends on the miss (instruction or data) that was last acknowledged.
- Automatic generation of the first word of the page table entry (PTE) of the tables being searched
- A real page address (RPA) register that matches the format of the lower word of the PTE
- TLB access instructions (**tlbli** and **tlbld**) that are used to load an address translation into the instruction or data TLBs
- Shadow registers for GPR0–GPR3 that allow miss code to execute without corrupting the state of any of the existing GPRs. Shadow registers are used only for servicing a TLB miss.

See [Section 7.4.5.2, “Implementation-Specific Memory Management,”](#) for more information about memory management for the core.

7.1.5.2 Cache Units

The e300c3 provide 16-Kbyte, four-way set-associative instruction and data caches. The cache block is 32 bytes long. The caches adhere to a write-back policy, but the e300 core allows control of cacheability, write policy, and memory coherency at the page and block levels. The caches use a pseudo LRU replacement policy.

As shown in [Figure 7-1](#) the caches provide a 64-bit interface to the instruction fetch unit and LSU. The surrounding logic selects, organizes, and forwards the requested information to the requesting unit. Write operations to the cache can be performed on a byte basis, and a complete read-modify-write operation to the cache can occur in each cycle.

The load/store and instruction fetch units provide the caches with the address of the data or instruction to be fetched. In the case of a cache hit, the cache returns two words to the requesting unit.

Because the data cache tags are single-ported, simultaneous load/store and snoop accesses cause resource contention. Snoop accesses have the highest priority and are given first access to the tags, unless the snoop access coincides with a tag write; in this case the snoop is retried and must re-arbitrate for cache access. Loads or stores deferred due to snoop accesses are performed on the clock cycle following the snoop.

The e300 core includes a new instruction cancel extension. The instruction cancel extension improves utilization of the instruction cache during cancel operations. It allows a new instruction fetch to be issued to the cache or to the bus if a canceled instruction fetch is pending or active on the bus. This supports hit-under-cancel and miss-under-cancel instruction fetch operations.

7.1.6 Bus Interface Unit (BIU)

Because the caches are on-chip, write-back caches, the most common transactions are burst-read memory operations, burst-write memory operations, and single-beat (noncacheable or write-through) memory read and write operations. There can also be address-only operations, variants of the burst and single-beat operations, (for example, global memory operations that are snooped and atomic memory operations), and address retry activity (for example, when a snooped read access hits a modified cache block).

Memory accesses can occur in single-beat (1–8 bytes) and four-beat burst (32 bytes) data transfers on the 64-bit data bus. The address and data buses operate independently to support pipelining and split transactions during memory accesses.

The e300 bus interface unit (BIU) has been enhanced to allow a pipeline slot to become available once a previous transaction has been granted the data bus (that is, as early as when the data tenure starts rather than after the data tenure completes), thus allowing for greater bus utilization in systems that support it. This is sometimes referred to as 1 1/2-level pipelining.

Typically, memory accesses are weakly ordered, meaning that sequences of operations, including load/store string and multiple instructions, do not necessarily complete in the order they begin. This weak ordering maximizes the efficiency of the bus without sacrificing coherency of the data. The core allows read operations to precede store operations (except when a dependency exists, or in cases where a

noncacheable access is performed), and provides support for a write operation to proceed a previously queued read data tenure (for example, allowing a snoop push to be enveloped by the address and data tenures of a read operation). Because the processor can dynamically optimize run-time ordering of load/store traffic, overall performance is improved.

7.1.7 System Support Functions

The e300 core implements several support functions that include power management, time base/decrementer registers for system timing tasks, a JTAG (based on IEEE Std 1149.1™) interface, hardware debug, and a phase-locked loop (PLL) clock multiplier. These system support functions are described in the following sections.

7.1.7.1 Power Management

The e300 core provides four power modes, selectable by setting the appropriate control bits in the machine state register (MSR) and the hardware implementation register 0 (HID0). When entering into a power mode other than full-power, the core will request entry via a *qreq* signal and will only enter another power mode after an acknowledge (*qack*) is received. The four power modes are as follows:

- Full-power—This is the default power state of the e300 core. The e300 core is fully powered and the internal functional units are operating at the full processor clock speed. If the dynamic power management mode is enabled, functional units that are idle will automatically enter a low-power state without affecting performance, software execution, or external hardware.
- Doze—All the functional units of the e300 core are disabled except for the time base/decrementer registers and the bus snooping logic. When the processor is in doze mode, an external asynchronous interrupt, system management interrupt, decremter interrupt, hard or soft reset, or machine check brings the e300 core into the full-power state. The core in doze mode maintains the PLL in a fully powered state and locked to the system external clock input (*sysclk*), so a transition to the full-power state takes only a few processor clock cycles.
- Nap—The nap mode further reduces power consumption by disabling bus snooping, leaving only the time base register and the PLL in a powered state. The core returns to the full-power state on receipt of an external asynchronous interrupt, system management interrupt, decremter interrupt, hard or soft reset, or machine check input (*mcp*) signal. A return to full-power state from a nap state takes only a few processor clock cycles.
- Sleep—Sleep mode reduces power consumption to a minimum by disabling all internal functional units; then external system logic may disable the PLL and *sysclk*. Returning the core to the full-power state requires the enabling of the PLL and *sysclk*, followed by the assertion of an external asynchronous interrupt, system management interrupt, hard or soft reset, or *mcp* signal after the time required to relock the PLL.

7.1.7.2 Time Base/Decrementer

The time base is a 64-bit register (accessed as two 32-bit registers) that is incremented once every four bus clock cycles; external control of the time base is provided through the time base/decrementer clock base enable (*tben*) signal. The decremter is a 32-bit register that generates a decremter interrupt after a

programmable delay. The contents of the decremter register are decremented once every four bus clock cycles, and the decremter interrupt is generated as the count passes through zero.

7.1.7.3 JTAG Test and Debug Interface

The core provides JTAG and hardware debug functions for facilitating board testing and chip debugging. The JTAG test interface (based on IEEE 1149.1) provides a means for boundary-scan testing of the core and the attached system logic. The hardware debug function accesses the JTAG test port, providing a means for executing test routines and facilitating chip and software debugging.

All instruction and data address breakpoints are accessible in the IBCR and DBCR. See [Section 7.4.8, “Debug Features,”](#) for more information.

7.1.7.4 Clock Multiplier

The internal clocking of the e300 core is generated from and synchronized to the external clock signal, *sysclk*, by means of a voltage-controlled, oscillator-based PLL. The PLL provides programmable internal processor clock multiplier ratios which multiply the externally supplied clock frequency. The bus clock is the same frequency and is synchronous with *sysclk*. The configuration of the PLL can be read by software from the hardware implementation register 1 (HID1).

7.1.7.5 Core Performance Monitor

The performance monitor provides the ability to count predefined events and processor clocks associated with particular operations, such as cache misses, mispredicted branches, or the number of cycles an execution unit stalls. The count of such events can be used to trigger the performance monitor interrupt.

The performance monitor can be used to do the following:

- Improve system performance by monitoring software execution and then recoding algorithms for more efficiency. For example, memory hierarchy behavior can be monitored and analyzed to optimize task scheduling or data distribution algorithms.
- Characterize processors in environments not easily characterized by benchmarking.
- Help system developers bring up and debug their systems.

The performance monitor uses the following resources:

- The performance monitor mark bit in the MSR (MSR[PMM]). This bit controls which programs are monitored.
- The move to/from performance monitor registers (PMR) instructions, **mtpmr** and **mfpmr**.
- The external core input, *pm_event_in*.
- PMRs:
 - The performance monitor counter registers (PMC0–PMC3) are 32-bit counters used to count software-selectable events. Each counter counts up to 128 events. UPMC0–UPMC3 provide user-level read access to these registers. They are identified in [Table 7-2](#).

- The performance monitor global control register (PMGC0) controls the counting of performance monitor events. It takes priority over all other performance monitor control registers. UPMGC0 provides user-level read access to PMGC0.
- The performance monitor local control registers (PMLCa0–PMLCa3) control each individual performance monitor counter. Each counter has a corresponding PMLCa register. UPMLCa0–UPMLCa3 provide user-level read access to PMLCa0–PMLCa3).
- The performance monitor interrupt is assigned to interrupt vector 0x0F00.

Software communication with the performance monitor is achieved through PMRs rather than SPRs. The PMRs are used for enabling conditions that can trigger the performance monitor interrupt.

7.2 e300 Processor and System Version Numbers

Table 7-1 lists the revision codes in the processor version register (PVR) and the system version register (SVR) to the revision level marked on the device. These registers can be accessed as SPRs through the e300 core (see Figure 7-2).

Table 7-1. Device Revision Level Cross-Reference

| MPC8315E Revision | Processor Version Register (PVR) | System Version Register (SVR) |
|-------------------|----------------------------------|--|
| 1.0 | 8085_0020 | MPC8315E (with security): 0x80B4_0010 MPC8315 (without security): 0x80B5_0010 MPC8314E (with security): 0x80B6_0010 MPC8314 (without security): 0x80B7_0010 |
| 1.1 | 8085_0020 | MPC8315E (with security): 0x80B4_0011 MPC8315 (without security): 0x80B5_0011 MPC8314E (with security): 0x80B6_0011 MPC8314 (without security): 0x80B7_0011 |

7.3 PowerPC Architecture Implementation

The PowerPC architecture consists of the following layers, and adherence to the PowerPC architecture can be measured in terms of which of the following levels of the architecture is implemented:

- User instruction set architecture (UISA)—Defines the base user-level instruction set, user-level registers, data types, floating-point interrupt model, memory models for a uniprocessor environment, and programming model for a uniprocessor environment.
- Virtual environment architecture (VEA)—Describes the memory model for a multiprocessor environment, defines cache control instructions, and describes other aspects of virtual environments. Implementations that conform to the VEA also adhere to the UISA but may not necessarily adhere to the OEA.
- Operating environment architecture (OEA)—Defines the memory management model, supervisor-level registers, synchronization requirements, and interrupt model. Implementations that conform to the OEA also adhere to the UISA and VEA.

The PowerPC architecture allows a wide range of designs for such features as cache and core interface implementations.

7.4 Implementation-Specific Information

This section describes the PowerPC architecture in general and specific details about the implementation of the e300 core as a low-power, 32-bit member of this PowerPC core family. The main topics addressed are as follows:

- [Section 7.4.1, “Register Model,”](#) describes the registers for the operating environment architecture common among e300 cores that implement the PowerPC architecture and describes the programming model. It also describes the additional registers that are unique to the core.
- [Section 7.4.2, “Instruction Set and Addressing Modes,”](#) describes the PowerPC instruction set and addressing modes for the OEA, and defines and describes the instructions implemented in the core.
- [Section 7.4.3, “Cache Implementation,”](#) describes the cache model that is defined generally for cores that implement the PowerPC architecture by the VEA. It also provides specific details about the e300 core cache implementation.
- [Section 7.4.4, “Interrupt Model,”](#) describes the interrupt model of the OEA and the differences in the core interrupt model.
- [Section 7.4.5, “Memory Management,”](#) describes generally the conventions for memory management among these cores. This section also describes the core implementation of the 32-bit PowerPC memory management specification.
- [Section 7.4.6, “Instruction Timing,”](#) provides a general description of the instruction timing provided by the superscalar, parallel execution supported by the PowerPC architecture and the e300 core.
- [Section 7.1.6, “Bus Interface Unit \(BIU\),”](#) describes the signals implemented on the core.

The e300 core is a high-performance, superscalar processor core. The PowerPC architecture allows optimizing compilers to schedule instructions to maximize performance through efficient use of the PowerPC instruction set and register model. The multiple, independent execution units allow compilers to optimize instruction throughput. Compilers that take advantage of the flexibility of the PowerPC architecture can additionally optimize system performance.

The following sections summarize the features of the core, including both those that are defined by the architecture and those that are unique to the various core implementations.

Specific features of the core are listed in [Section 7.1.1, “e300c3 Features.”](#)

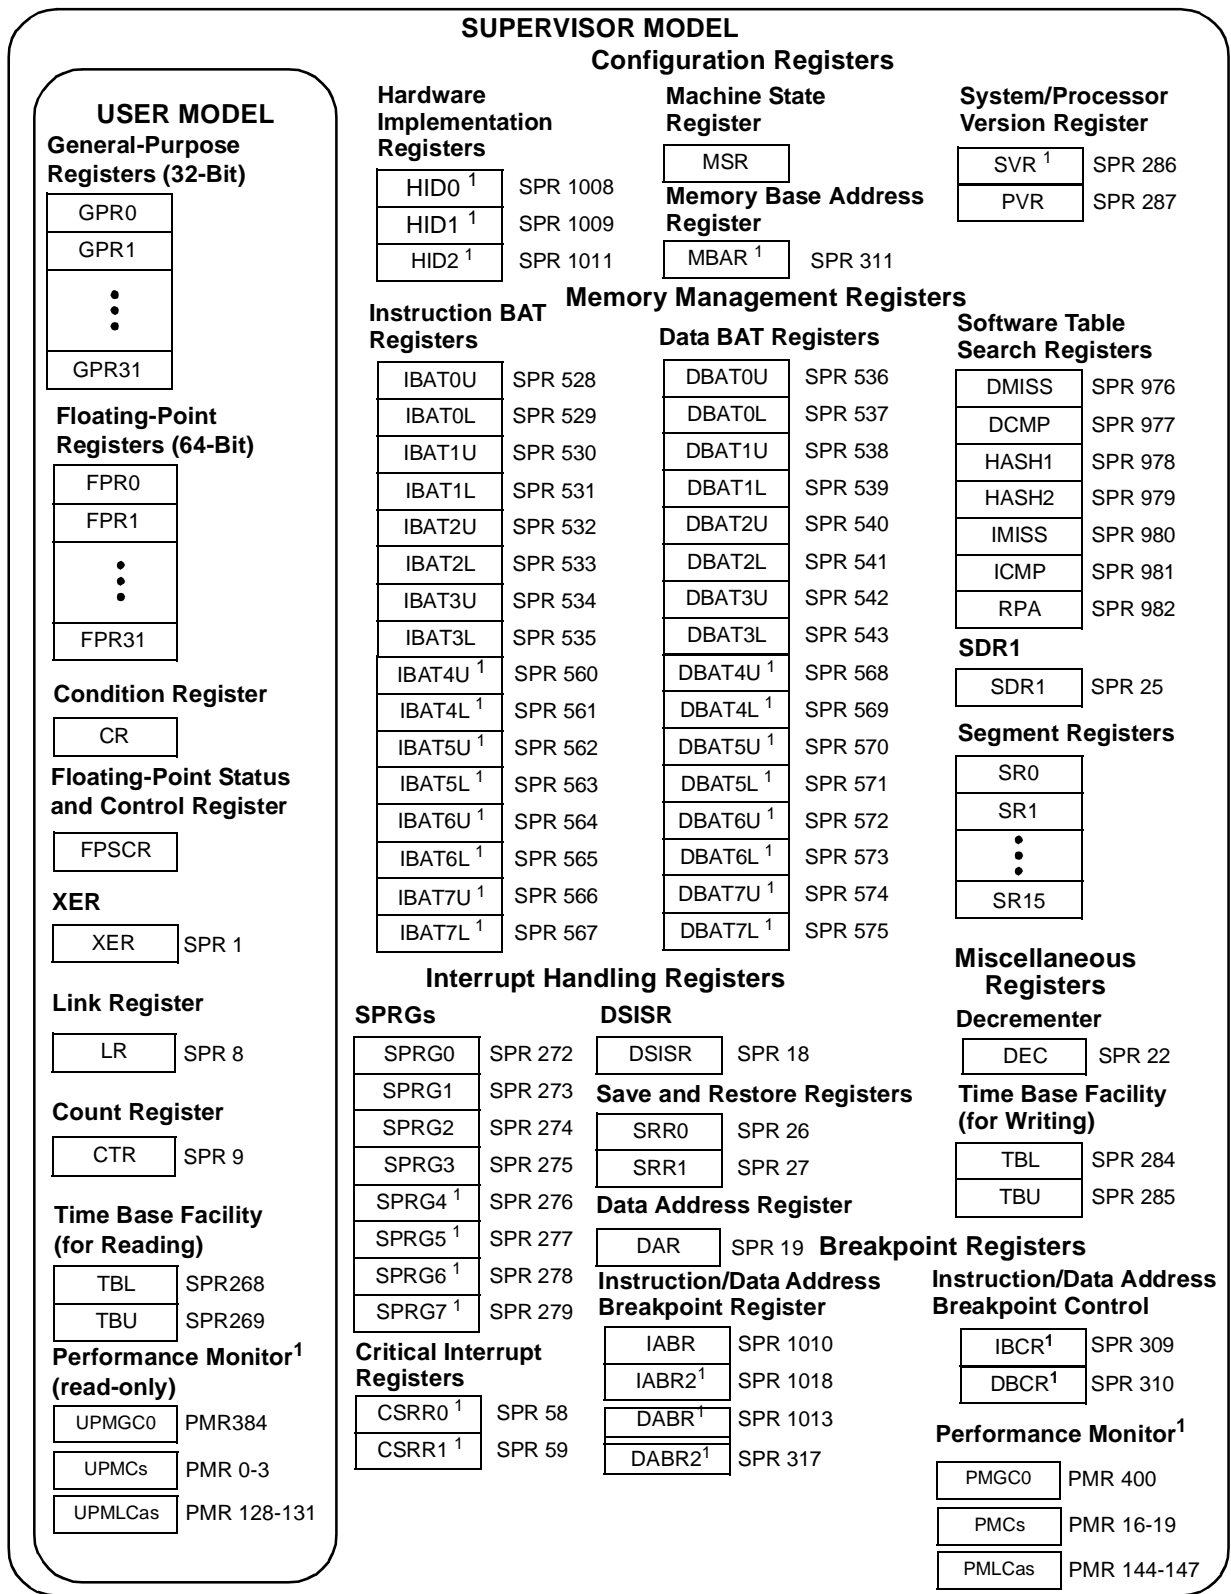
7.4.1 Register Model

The PowerPC architecture defines register-to-register operations for most computational instructions. Source operands for these instructions are accessed from the registers or are provided as immediate values embedded in the instruction opcode. The three-register instruction format allows specification of a target register distinct from the two-source operands. Load and store instructions transfer data between registers and memory.

The e300 core has two levels of privilege: supervisor mode of operation (typically used by the operating system) and user mode of operation (used by the application software). The programming models incorporate 32 GPRs, 32 FPRs, special-purpose registers (SPRs), and several miscellaneous registers. Each core also has its own unique set of hardware implementation (HID) registers.

Having access to privileged instructions, registers, and other resources allows the operating system to control the application environment (providing virtual memory and protecting operating system and critical machine resources). Instructions that control the state of the e300 core, the address translation mechanism, and supervisor registers can be executed only when the core is operating in supervisor mode.

Figure 7-2 shows all the core registers available at the user and supervisor level. The numbers to the right of the SPRs indicate the number that is used in the syntax of the instruction operands for the move to/from SPR instructions.



¹ These registers are e300 core implementation-specific (not defined by the PowerPC architecture).

Figure 7-2. e300 Programming Model—Registers

The following sections describe the e300-core-implementation-specific features as they apply to registers.

7.4.1.1 UISA Registers

UISA registers are user-level registers that include the following.

7.4.1.1.1 General-Purpose Registers (GPRs)

The PowerPC architecture defines 32 user-level GPRs that are 32-bit wide in 32-bit cores. The GPRs serve as the data source or destination for all integer instructions.

7.4.1.1.2 Floating-Point Registers (FPRs)

The PowerPC architecture also defines 32 user-level, 64-bit FPRs. The FPRs serve as the data source or destination for floating-point instructions. These registers can contain data objects of either single- or double-precision floating-point formats.

7.4.1.1.3 Condition Register (CR)

The CR is a 32-bit user-level register that provides a mechanism for testing and branching. It consists of eight 4-bit fields that reflect the results of certain operations, such as move, integer and floating-point comparisons, arithmetic, and logical operations.

7.4.1.1.4 Floating-Point Status and Control Register (FPSCR)

The user-level FPSCR contains all floating-point exception signal bits, exception summary bits, exception enable bits, and rounding control bits needed for compliance with the IEEE 754 standard.

7.4.1.1.5 User-Level SPRs

The PowerPC architecture defines numerous special purpose registers that serve a variety of functions, such as providing controls, indicating status, configuring the core, and performing special operations. During normal execution, a program can access the registers, as shown in [Figure 7-2](#), depending on the program's access privilege (supervisor or user, determined by the privilege-level bit, MSR[PR]). Note that GPRs and FPRs are accessed through operands that are part of the instructions. Access to registers can be explicit (that is, through the use of specific instructions for that purpose such as Move to Special-Purpose Register (**mtspr**) and Move from Special-Purpose Register (**mfspir**) instructions) or implicit, as the part of the execution of an instruction. Some registers are accessed both explicitly and implicitly. In the e300 core, all SPRs are 32 bits wide.

The following SPRs are accessible by user-level software:

- Link register (LR)—The LR can be used to provide the branch target address and to hold the return address after branch and link instructions. The LR is 32 bits wide in 32-bit implementations.
- Count register (CTR)—The CTR is decremented and tested automatically as a result of branch-and-count instructions. The CTR is 32 bits wide in 32-bit implementations.
- XER register—The 32-bit XER contains the summary overflow bit, integer carry bit, overflow bit, and a field specifying the number of bytes to be transferred by a Load String Word Indexed (**lswx**) or Store String Word Indexed (**stswx**) instruction.

7.4.1.2 VEA Registers

The VEA introduces the time base facility (TB) for reading. The TB is a 64-bit register pair whose contents are incremented once every four core input clock cycles. The TB consists of two 32-bit registers—time base upper (TBU) and time base lower (TBL). Note that the time base registers are read-only in user state.

7.4.1.3 OEA Registers

OEA registers are supervisor-level registers that include the following.

7.4.1.3.1 Machine State Register (MSR)

The MSR, shown in [Figure 7-3](#), is a supervisor-level register that defines the state of the core. The contents of this register are saved when an interrupt is taken, and restored when the interrupt handling completes. A critical interrupt is taken in the e300 core when the *cint* signal is asserted and MSR[CE] is set. The e300 core implements the MSR as a 32-bit register.




Figure 7-3. Machine State Register (MSR)

Table 7-2 shows the bit definitions for MSR.

Table 7-2. MSR Bit Descriptions

| Bits | Name | Description |
|--------------------|------|---|
| 0 ¹ | — | Reserved. Full function. |
| 1–4 ¹ | — | Reserved. Partial function. |
| 5–9 ¹ | — | Reserved. Full function. |
| 10–12 ¹ | — | Reserved. Partial function. |
| 13 | POW | Power management enable (implementation-specific) 0 Disables programmable power modes (normal operation mode) 1 Enables programmable power modes (nap, doze, or sleep mode). This bit controls the programmable power modes only; it has no effect on dynamic power management (DPM). MSR[POW] may be altered with an mtmsr instruction only. Also, when altering the POW bit, software may alter only this bit in the MSR and no others. The mtmsr instruction must be followed by a context-synchronizing instruction. |
| 14 | TGPR | Temporary GPR remapping (implementation-specific) 0 Normal operation 1 TGPR mode. GPR0–GPR3 are remapped to TGPR0–TGPR3 for use by TLB miss routines. The contents of GPR0–GPR3 remain unchanged while MSR[TGPR] = 1. Attempts to use GPR4–GPR31 with MSR[TGPR] = 1 yield undefined results. Temporarily replaces TGPR0–TGPR3 with GPR0–GPR3 for use by TLB miss routines. The TGPR bit is set when either an instruction TLB miss, data read miss, or data write miss interrupt is taken. The TGPR bit is cleared by an rfi instruction. |

Table 7-2. MSR Bit Descriptions (continued)

| Bits | Name | Description |
|--------------------|------|---|
| 15 | ILE | Interrupt little-endian mode. When an interrupt occurs, this bit is copied into MSR[LE] to select the endian mode for the context established by the interrupt. |
| 16 | EE | External interrupt enable 0 The processor ignores external interrupts, system management interrupts, and decremter interrupts. 1 The processor is enabled to take an external interrupt, system management interrupt, or decremter interrupt. |
| 17 | PR | Privilege level 0 The processor can execute both user- and supervisor-level instructions 1 The processor can only execute user-level instructions |
| 18 | FP | Floating-point available 0 The processor prevents dispatch of floating-point instructions, including floating-point loads, stores, and moves. 1 The processor can execute floating-point instructions and can take floating-point enabled exception type program interrupts. |
| 19 | ME | Machine check enable 0 Machine check interrupts are disabled 1 Machine check interrupts are enabled |
| 20 | FE0 | Floating-point exception mode 0 |
| 21 | SE | Single-step trace enable 0 The processor executes instructions normally 1 The processor generates a trace interrupt upon the successful completion of the next instruction |
| 22 | BE | Branch trace enable 0 The processor executes branch instructions normally 1 The processor generates a trace interrupt upon the successful completion of a branch instruction |
| 23 | FE1 | Floating-point exception mode 1 |
| 24 | CE | Critical interrupt enable 0 Critical interrupts disabled 1 Critical interrupts enabled; critical interrupt and rfci instruction enabled The critical interrupt is an asynchronous implementation-specific interrupt. The critical interrupt vector offset is 0x00A00. The rfci instruction is implemented to return from these interrupt handlers. Also, CSRR0 and CSRR1 are used to save and restore the processor state for critical interrupts. |
| 25 | IP | Interrupt prefix. The setting of this bit specifies whether an interrupt vector offset is prepended with Fs or 0s. In the following description, <i>nnnn</i> is the offset of the interrupt. 0 Interrupts are vectored to the physical address 0x000n_nnnn 1 Interrupts are vectored to the physical address 0xFFFFn_nnnn |
| 26 | IR | Instruction address translation 0 Instruction address translation is disabled 1 Instruction address translation is enabled |
| 27 | DR | Data address translation 0 Data address translation is disabled 1 Data address translation is enabled |
| 28–29 ¹ | — | Reserved. Full function. Bit 29 not reserved on e300c3. |

Table 7-2. MSR Bit Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 29 | PMM | Performance monitor mark bit (e300c3). System software can set PMM when a marked process is running to enable statistics to be gathered only during the execution of the marked process. MSR[PR] and MSR[PMM] together define a state that the processor (supervisor or user) and the process (marked or unmarked) may be in at any time. If this state matches an individual state specified in the PMLCan, the state for which monitoring is enabled, counting is enabled. |
| 30 | RI | Recoverable interrupt (for system reset and machine check interrupts) 0 Interrupt is not recoverable 1 Interrupt is recoverable |
| 31 | LE | Little-endian mode enable 0 The processor runs in big-endian mode 1 The processor runs in little-endian mode. |

¹ All reserved bits should be set to zero for future compatibility.

7.4.1.3.2 Segment Registers (SRs)

For memory management, 32-bit processors implement sixteen 32-bit SRs. To speed access, the core implements the SRs as two arrays: a main array, for data memory accesses, and a shadow array, for instruction memory accesses. Loading a segment entry with the Move to Segment Register (**mtsr**) instruction loads both arrays.

7.4.1.3.3 Supervisor-Level SPRs

The e300 core, like the G2_LE core, has additional supervisor-level SPRs, which are shown in [Figure 7-2](#). Two critical interrupt SPRs (CSRR0 and CSRR1), eight SPRGs (SPRG0–SPRG7), eight pairs of instruction BATs (IBAT0–IBAT7), eight pairs of data BATs (DBAT0–DBAT7), one system version register (SVR), one system memory base address (MBAR), one instruction address breakpoint control (IBCR), one data address breakpoint control (DBCR), a new instruction breakpoint register (IABR2), and two data address breakpoint registers (DABR and DABR2) are integrated into the core.

Supervisor-level SPRs include the following:

- The DSISR defines the cause of data access and alignment interrupts. The cause of a DSI interrupt for a data breakpoint (match with DABR and DABR2) can be determined by the value of the DSISR[DABR] bit (bit 9).
- The data address register (DAR) holds the address of an access after an alignment or DSI interrupt. For example, it contains the address of the breakpoint match condition.
- The decremter register (DEC) is a 32-bit decremting counter that provides a mechanism for causing a decremter interrupt after a programmable delay.
- SDR1 specifies the page table format used in virtual-to-physical address translation for pages. (Note that physical address is referred to as ‘real address’ in the architecture specification.)
- The machine status save/restore register 0 (SRR0) is used for saving the address of the instruction that caused the interrupt, and the address to return to when a Return from Interrupt (**rfi**) instruction is executed.

- The machine status save/restore register 1 (SRR1) is used to save machine status on interrupts and to restore machine status when an **rfi** instruction is executed.
- The SPRG0–SPRG7 registers are provided for operating system use. They reduce the latency that may be incurred in the saving of registers to memory while in a handler. Note that the e300 implements four more SPRGs than the G2 (SPRG0–SPRG3).
- The time base register (TB) is a 64-bit register that maintains the time of day and operates interval timers. It consists of two 32-bit fields: time base upper (TBU) and time base lower (TBL).
- The processor version register (PVR) is a read-only register that identifies the version (model) and revision level of the processor. See [Table 7-9](#) for the version and revision level of the PVR for the e300 processor core.
- Block address translation (BAT) arrays—The PowerPC architecture defines 16 BAT registers. The e300 core includes a total of eight pairs of DBAT and eight pairs of IBAT registers. See [Figure 7-2](#) for a list of the SPR numbers for the BAT arrays.

The following supervisor-level SPRs are implementation-specific (not defined in the PowerPC architecture):

- DMISS and IMISS are read-only registers that are loaded automatically on an instruction or data TLB miss.
- HASH1 and HASH2 contain the physical addresses of the primary and secondary page table entry groups (PTEGs).
- ICMP and DCMP contain a duplicate of the first word in the page table entry (PTE) for which the table search is looking.
- The required physical address (RPA) register is loaded by the core with the second word of the correct PTE during a page table search.
- The system version register (SVR) is available on the e300 core, which identifies the specific version (model) and revision level of the system-on-a-chip (SOC) integration.
- System memory base address (MBAR) is an implementation-specific register available on the e300 core. It supports a temporary storage for the system-level memory map.
- The instruction and data address breakpoint registers (IABR, IABR2, DABR, DABR2) are loaded with an instruction or data address, respectively, that is compared to instruction addresses in the dispatch queue or to the data address in the LSU. When an address match occurs, a breakpoint interrupt is generated.
- One instruction breakpoint control register (IBCR) and one data breakpoint control register (DBCR) are implemented in the e300 core.
- To support critical interrupts, two registers (CSRR0 and CSRR1) are included in the e300 core.
- Eight SPRG registers (SPRG0–SPRG7) are in the e300 core.
- Block address translation (BAT) arrays—The e300 core has eight instruction and eight data BAT registers.
- The hardware implementation (HID0 and HID1) registers provide the means for enabling core checkstops and features and allow software to read the configuration of the PLL configuration signals. The HID2 register enables the true little-endian mode, cache way-locking, and the additional BAT registers.

Table 7-3 shows the bit definitions for HID0.

Table 7-3. e300 HID0 Bit Descriptions

| Bits | Name | Function |
|------|-------|--|
| 0 | EMCP | Enable \overline{mcp} . The purpose of this bit is to mask out machine check interrupts caused by assertion of mcp , similar to how MSR[EE] can mask external interrupts. 0 Masks mcp . Asserting mcp does not generate a machine check interrupt or a checkstop. 1 Asserting mcp causes checkstop if MSR[ME] = 0 or a machine check interrupt if ME = 1 |
| 1 | ECPE | Enable cache parity errors. 0 Disables instruction and data cache parity error reporting 1 Allows a detected cache parity error to cause a machine check interrupt if MSR[ME] = 1 or a checkstop if MSR[ME] = 0 |
| 2 | EBA | Enable $ap_in[0:3]$ and \overline{ape} for address parity checking. 0 Disables address parity checking during a snoop operation 1 Allows an address parity error during snoop operations to cause a checkstop if MSR[ME] = 0 or a machine check interrupt if MSR[ME] = 1 Note: Do not set this bit; the CSB does not have parity signals. |
| 3 | EBD | Enable \overline{dpe} for data parity checking. 0 Disables data parity checking 1 Allows a data parity error during reads to cause a checkstop if MSR[ME] = 0 or a machine check interrupt if MSR[ME] = 1 Note: Do not set this bit; the CSB does not have parity signals. |
| 4 | SBCLK | clk_out output enable. Used in conjunction with HID0[ECLK] and \overline{hreset} to configure clk_out . See Table 7-4 for settings. |
| 5 | — | Reserved |
| 6 | ECLK | clk_out output enable. Used in conjunction with HID0[SBCLK] and the \overline{hreset} signal to configure clk_out . See Table 7-4 for settings. |
| 7 | PAR | Disable precharge of $\overline{artry_out}$ 0 Precharge of $\overline{artry_out}$ enabled 1 Alters bus protocol slightly by preventing the processor from driving $\overline{artry_out}$ to high (negated) state. If this is done, the integrated device must restore the signals to the high state. |
| 8 | DOZE | Doze mode enable. Operates in conjunction with MSR[POW]. 0 Doze mode disabled 1 Doze mode enabled. Doze mode is invoked by setting MSR[POW] while this bit is set. In doze mode, the PLL, time base, and snooping remain active. |
| 9 | NAP | Nap mode enable. Operates in conjunction with MSR[POW]. 0 Nap mode disabled 1 Nap mode enabled. Nap mode is invoked by setting MSR[POW] while this bit is set. In nap mode, the PLL and time base remain active. |
| 10 | SLEEP | Sleep mode enable. Operates in conjunction with MSR[POW]. 0 Sleep mode disabled 1 Sleep mode enabled. Sleep mode is invoked by setting MSR[POW] while this bit is set. \overline{qreq} is asserted to indicate that the processor is ready to enter sleep mode. If the system logic determines that the processor may enter sleep mode, the quiesce acknowledge signal, \overline{qack} , is asserted back to the processor. Once \overline{qack} assertion is detected, the processor enters sleep mode after several processor clocks. At this point, the system logic may turn off the PLL by first configuring $pll_cfg[0:6]$ to PLL bypass mode, then disabling $sysclk$. |

Table 7-3. e300 HID0 Bit Descriptions (continued)

| Bits | Name | Function |
|-------|-------|--|
| 11 | DPM | Dynamic power management enable 0 Dynamic power management is disabled 1 Functional units enter a low-power mode automatically if the unit is idle. This does not affect operational performance and is transparent to software or any external hardware. |
| 12–15 | — | Reserved |
| 16 | ICE | Instruction cache enable 0 The instruction cache is neither accessed nor updated. All pages are accessed as if they were marked cache-inhibited (WIM = x1x). Potential cache accesses from the bus (snoop and cache operations) are ignored. In the disabled state for the L1 caches, the cache tag state bits are ignored and all instruction fetches are propagated to the coherent system bus (CSB) as single-beat transactions. For those transactions, however, \overline{ci} reflects the state of the I bit in the MMU for that page regardless of cache disabled status. ICE is zero at power-up. 1 The instruction cache is enabled |
| 17 | DCE | Data cache enable 0 The data cache is neither accessed nor updated. All pages are accessed as if they were marked cache-inhibited (WIM = x1x). Potential cache accesses from the bus (snoop and cache operations) are ignored. In the disabled state for the L1 caches, the cache tag state bits are ignored and all data read and write accesses are propagated to the CSB as single-beat transactions. For those transactions, however, \overline{ci} reflects the state of the I bit in the MMU for that page regardless of cache disabled status. DCE is zero at power-up. 1 The data cache is enabled |
| 18 | ILOCK | Instruction cache lock 0 Normal operation 1 The entire instruction cache is locked (that is, all eight ways of the cache are locked). A locked cache supplies data normally on a hit, but the access is treated as a cache-inhibited transaction on a miss. On a miss, the transaction to the bus is single-beat; however, \overline{ci} still reflects the state of the I bit in the MMU for that page independent of cache locked or disabled status. To prevent locking during a cache access, an isync instruction must precede the setting of ILOCK. |
| 19 | DLOCK | Data cache lock 0 Normal operation 1 The entire data cache is locked (that is, all eight ways of the cache are locked). A locked cache supplies data normally on a hit, but is treated as a cache-inhibited transaction on a miss. On a miss, the transaction to the bus is single-beat; however, \overline{ci} still reflects the state of the I bit in the MMU for that page independent of cache locked or disabled status. A snoop hit to a locked L1 data cache performs as if the cache were not locked. A cache block invalidated by a snoop remains invalid until the cache is unlocked. To prevent locking during a cache access, a sync instruction must precede the setting of DLOCK. |
| 20 | ICFI | Instruction cache Flash invalidate 0 The instruction cache is not invalidated. The bit is cleared when the invalidation operation begins (usually the next cycle after the write operation to the register). The instruction cache must be enabled for the invalidation to occur. 1 An invalidate operation is issued that marks the state of each instruction cache block as invalid. Cache access is blocked during this time. Setting ICFI clears all the valid bits of the blocks and the PLRU bits to point to way L0 of each set. For the e300 core, the proper use of the ICFI and DCFI bits is to set and clear them with two consecutive mtspr operations. |

Table 7-3. e300 HID0 Bit Descriptions (continued)

| Bits | Name | Function |
|-------|---------|---|
| 21 | DCFI | Data cache Flash invalidate 0 The data cache is not invalidated. The bit is cleared when the invalidation operation begins (usually the next cycle after the write operation to the register). The data cache must be enabled for the invalidation to occur. 1 An invalidate operation is issued that marks the state of each data cache block as invalid without writing back modified cache blocks to memory. Cache access is blocked during this time. Bus accesses to the cache are signaled as a miss during invalidate-all operations. Setting DCFI clears all the valid bits of the blocks and the PLRU bits to point to way L0 of each set. For the e300 core, the proper use of the ICFI and DCFI bits is to set and clear them with two consecutive mtspr operations. |
| 22–23 | — | Reserved, should be cleared. |
| 24 | IFEM | Enable M bit on bus for instruction fetches 0 M bit not reflected on bus for instruction fetches. Instruction fetches are treated as nonglobal on the bus. 1 Instruction fetches reflect the M bit from the WIM settings |
| 25 | DECAREN | Decrementer auto reload 0 Normal operation. 1 Decrementer loads last mtdec value for precise periodic interrupt. |
| 26 | — | Reserved, should be cleared. |
| 27 | FBIOB | Force branch indirect on the bus 0 Register indirect branch targets are fetched normally 1 Forces register indirect branch targets to be fetched externally |
| 28 | ABE | Address broadcast enable. Controls whether certain address-only operations (such as cache operations) are broadcast on the bus. 0 Address-only operations affect only local caches and are not broadcast 1 Address-only operations are broadcast on the bus Affected instructions are dcbi , dcbf , and dcbst . Note that these cache control instruction broadcasts are not snooped by the e300 core. Refer to Section 4.3.3, “Data Cache Control,” for more information. |
| 29–30 | — | Reserved |
| 31 | NOOPTI | No-op the data cache touch instructions 0 The dcbt and dcbst instructions are enabled 1 The dcbt and dcbst instructions are no-oped internal to the e300 core |

Table 7-4 shows how HID0[ECLK] and HID0[SBCLK] are used to configure the *clk_out* signal.

Table 7-4. Using HID0[ECLK] and HID0[SBCLK] to Configure *clk_out*

| \overline{hreset} | ECLK | SBCLK | <i>clk_out</i> |
|---------------------|------|-------|---|
| Asserted | x | x | Bus clock (small pulse for every rising edge of sysclk) |
| Negated | 0 | 0 | Clock output off |
| | 0 | 1 | Core clock/2 |
| | 1 | 0 | Core clock |
| | 1 | 1 | Bus clock |

Table 7-5 shows the bit definitions for HID1

Table 7-5. HID1 Bit Descriptions

| Bits | Name | Description |
|------|------|-------------------------------------|
| 0 | PC0 | PLL configuration bit 0 (read-only) |
| 1 | PC1 | PLL configuration bit 1 (read-only) |
| 2 | PC2 | PLL configuration bit 2 (read-only) |
| 3 | PC3 | PLL configuration bit 3 (read-only) |
| 4 | PC4 | PLL configuration bit 4 (read-only) |
| 5 | PC5 | PLL configuration bit 5 (read-only) |
| 6 | PC6 | PLL configuration bit 6 (read-only) |
| 7–31 | — | Reserved, should be cleared |

Note: The clock configuration bits reflect the state of the *pll_cfg[0:6]* signals.

Table 7-6 shows the bit definitions for HID2.

Table 7-6. e300HID2 Bit Descriptions

| Bits | Name | Description |
|------|-----------|---|
| 0–3 | — | Reserved, should be cleared. |
| 4 | LET | True little-endian. This bit enables true little-endian mode operation for instruction and data accesses. This bit is set to reflect the state of the <i>tle</i> signal at the negation of <i>hreset</i> . This bit is used in conjunction with MSR[LE] to determine the endian mode of operation. 0 No function 1 True little-endian mode, when MSR[LE] = 1 Changing the value of this bit during normal operation is not recommended |
| 5 | IFEB | Instruction fetch burst extension. This bit enables the instruction fetch burst extension. 0 Instruction fetch burst extension disabled 1 Instruction fetch burst extension enabled |
| 6 | — | Reserved, should be cleared. |
| 7 | MESISTATE | MESI state enable. This bit enables the four-state MESI cache coherency protocol. 0 MESI disabled. The data cache uses a three-state MEI coherency protocol. 1 MESI enabled. The data cache uses a four-state MESI protocol. |
| 8 | IFEC | Instruction fetch cancel extension. This bit enables the instruction fetch cancel extension. 0 Instruction fetch cancel extension disabled 1 Instruction fetch cancel extension enabled |
| 9 | EBQS | Enable BIU queue sharing. This bit enables data cache queue sharing. 0 Data cache queue sharing disabled 1 Data cache queue sharing enabled |
| 10 | EBPX | Enable BIU pipeline extension. This bit enables the bus interface unit pipeline extension. 0 BIU pipeline extension disabled; 1 level pipeline 1 BIU pipeline extension enabled; 1-1/2 level pipeline |

Table 7-6. e300HID2 Bit Descriptions (continued)

| Bits | Name | Description |
|-------|------------|---|
| 11 | ELRW | Enable weighted LRU. This bit enables the use of an adjusted (weighted) LRU. 0 Normal operation. 1 The dcbt, dcbst, and dcbz instructions use an adjusted (weighted) LRU such that they always select and replace the lowest unlocked way in the data cache. |
| 12 | NOKS | No kill for snoop. This bit enables the forcing of kill-type snoops to flush data instead of killing it. 0 Normal operation. 1 Forces write-with-kill snoops to flush instead of kill (snoop can never kill data). |
| 13 | HBE | High BAT enable. Regardless of the setting of HID2[HBE], these BATs are accessible by mf spr and mt spr . 0 IBAT[4–7] and DBAT[4–7] are disabled 1 IBAT[4–7] and DBAT[4–7] are enabled |
| 14–15 | — | Reserved, should be cleared. |
| 16–18 | IWLCK[0–2] | Instruction cache way-lock. Useful for locking blocks of instructions into the instruction cache for time-critical applications that require deterministic behavior. 000 no ways locked 001 way 0 locked 010 way 0 through way 1 locked 011 way 0 through way 2 locked 100 way 0 through way 2 locked in e300c3. 101 way 0 through way 2 locked in e300c3. 110 way 0 through way 2 locked in e300c3. 111 way 0 through way 2 locked in e300c3. Setting HID0[ILOCK] will lock all ways. |
| 19 | ICWP | Instruction cache way protection. Used to protect locked ways in the instruction cache from being invalidated. 0 Instruction cache way protection disabled 1 Instruction cache way protection enabled |
| 20–23 | — | Reserved, should be cleared. |
| 24–26 | DWLCK[0–2] | Data cache way-lock. Useful for locking blocks of data into the data cache for time-critical applications where deterministic behavior is required. 000 no ways locked 001 way 0 locked 010 way 0 through way 1 locked 011 way 0 through way 2 locked 100 way 0 through way 2 locked in e300c3. 101 way 0 through way 2 locked in e300c3. 110 way 0 through way 2 locked in e300c3. 111 way 0 through way 2 locked in e300c3. Setting HID0[DLOCK] will lock all ways. |
| 27–31 | — | Reserved, should be cleared. |

7.4.2 Instruction Set and Addressing Modes

The following sections describe the PowerPC instruction set and addressing modes in general.

7.4.2.1 PowerPC Instruction Set and Addressing Modes

All PowerPC instructions are encoded as single-word (32-bit) opcodes. Instruction formats are consistent among all instruction types, permitting efficient decoding to occur in parallel with operand accesses. This fixed instruction length and consistent format simplifies instruction pipelining.

The PowerPC instructions are divided into the following categories:

- Integer instructions—These include computational and logical instructions.
 - Integer arithmetic instructions
 - Integer compare instructions
 - Integer logical instructions
 - Integer rotate and shift instructions
- Floating-point instructions—These include floating-point computational instructions, as well as instructions that affect the FPSCR.
 - Floating-point arithmetic instructions
 - Floating-point multiply/add instructions
 - Floating-point rounding and conversion instructions
 - Floating-point compare instructions
 - Floating-point status and control instructions
- Load/store instructions—These include integer and floating-point load and store instructions.
 - Integer load and store instructions
 - Integer load and store multiple instructions
 - Floating-point load and store
 - Primitives used to construct atomic memory operations (**lwarx** and **stwcx.** instructions)
- Flow control instructions—These include branching instructions, condition register logical instructions, trap instructions, and other instructions that affect the instruction flow.
 - Branch and trap instructions
 - Condition register logical instructions
- Processor control instructions—These instructions are used for synchronizing memory accesses and management of caches, TLBs, and the segment registers.
 - Move to/from SPR instructions
 - Move to/from MSR
 - Move to/from PMR
 - Synchronize
 - Instruction synchronize
- Memory control instructions—These instructions provide control of caches, TLBs, and segment registers.
 - Supervisor-level cache management instructions
 - Translation lookaside buffer management instructions. Note that there are additional implementation-specific instructions.

- User-level cache instructions
- Segment register manipulation instructions
- The e300 core implements the following instructions which are defined as optional by the PowerPC architecture:
 - Floating Select (**fsel**)
 - Floating Reciprocal Estimate Single-Precision (**fres**)
 - Floating Reciprocal Square Root Estimate (**frsqrte**)
 - Store Floating-Point as Integer Word (**stfiwx**)

Note that this grouping of instructions does not indicate the execution unit that executes a particular instruction or group of instructions.

Integer instructions operate on byte, half-word, and word operands. Floating-point instructions operate on single-precision (one word) and double-precision (one double word) floating-point operands. The PowerPC architecture uses instructions that are 4 bytes long and word-aligned. It provides for byte, half-word, and word operand loads and stores between memory and a set of 32 GPRs. It also provides for word and double-word operand loads and stores between memory and a set of 32 FPRs.

Computational instructions do not modify memory. To use a memory operand in a computation and then modify the same or another memory location, the memory contents must be loaded into a register, modified, and then written back to the target location with distinct instructions.

The core follows the program flow when it is in the normal execution state. However, the flow of instructions can be interrupted directly by the execution of an instruction or by an asynchronous event. Either kind of interrupt may cause one of several components of the system software to be invoked.

7.4.2.2 Implementation-Specific Instruction Set

The e300 core instruction set is defined as follows:

- The core provides hardware support for all 32-bit PowerPC instructions.
- The core provides two implementation-specific instructions used for software table search operations following TLB misses:
 - Load Data TLB Entry (**tlbld**)
 - Load Instruction TLB Entry (**tlbli**)
- The core implements the following instruction which is added to support critical interrupts (also supported on the G2_LE). This is a supervisor-level, context synchronizing instruction.
 - Return from Critical Interrupt (**rftci**)
- The core implements the following instruction which is added to support easy start-up initialization or reloading of the instruction cache.
 - Instruction Cache Block Touch (**icbt**)
- The core provides the following performance monitor instructions:
 - Move to Performance Monitor Register (**mtpmr**)
 - Move from Performance Monitor Register (**mfpmr**)

7.4.3 Cache Implementation

The following sections describe the general cache characteristics as implemented in the PowerPC architecture and the core implementation.

7.4.3.1 PowerPC Cache Characteristics

The PowerPC architecture does not define hardware aspects of cache implementations. The e300 core controls the following memory access modes on a page or block basis:

- Write-back/write-through mode
- Caching-inhibited mode
- Memory coherency

Note that in the core, a cache block is defined as eight words. The VEA defines cache management instructions that provide a means by which the application programmer can affect the cache contents.

7.4.3.2 Implementation-Specific Cache Organization

The e300c3 provides 16-Kbyte, four-way set-associative instruction and data caches. The caches are physically addressed, and the data cache can operate in either write-back or write-through mode as specified by the PowerPC architecture.

The data cache is configured as 128 sets of 4 blocks each on the e300c3. Each block consists of 32 bytes, 2 state bits, and an address tag. The two state bits implement the three-state MEI (modified/exclusive/invalid) protocol. Each block contains eight 32-bit words. Note that the PowerPC architecture defines the term ‘block’ as the cacheable unit. For the core, the block size is equivalent to a cache line. A block diagram of the data cache organization is shown in [Figure 7-4](#).

The instruction cache is configured as 128 sets of 4 blocks each on the e300c3. Each block consists of 32 bytes, an address tag, and a valid bit. The instruction cache may not be written to, except through a block fill operation. In the e300 core, the instruction cache is blocked only until the critical load completes. The e300 core supports instruction fetching from other instruction cache lines following the forwarding of the critical-first-double-word of a cache line load operation. Successive instruction fetches from the cache line being loaded are forwarded, and accesses to other instruction cache lines can proceed during the cache line load operation. The instruction cache is not snooped, and cache coherency must be maintained by software. A fast hardware invalidation capability is provided to support cache maintenance.

The e300c3 data cache is configured as 128 sets of four blocks per set. The organization of the data cache is shown in Figure 7-4.

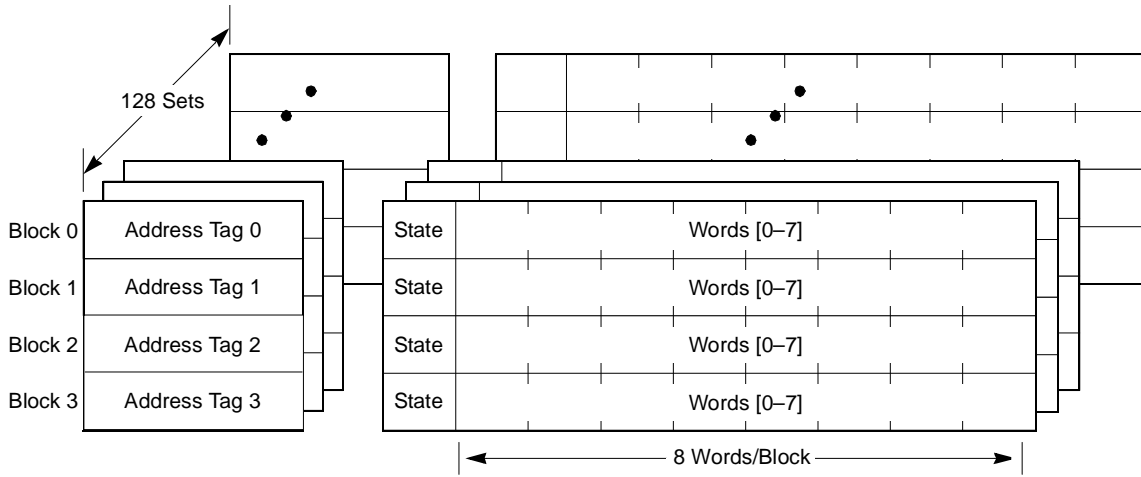


Figure 7-4. e300c3 Data Cache Organization

Each cache block contains eight contiguous words from memory that are loaded from an 8-word boundary (that is, bits A[27–31] of the effective addresses are zero); thus, a cache block never crosses a page boundary. Misaligned accesses across a page boundary can incur a performance penalty.

The e300 core cache blocks are loaded in four beats of 64 bits each on the 64-bit data bus. The burst load is performed as critical-double-word-first. The data cache is blocked to internal accesses until the load completes; the instruction cache allows sequential fetching during a cache block load. In the core, the critical-double-word is simultaneously written to the cache and forwarded to the requesting unit, thus minimizing stalls due to load delays.

To ensure coherency among caches in a multiprocessor (or multiple caching-device) implementation, the core implements the MEI protocol during normal operation of the data cache. The new data cache MESI extension supports the additional fourth cache coherency shared state for the data cache. To support this feature, the shared signal, *shd*, has been added to the bus interface. Although the MESI protocol is supported by the e300 core, it is not implemented on MPC8315E. The following four states indicate the state of the cache block:

- **Modified**—The cache block is modified with respect to system memory; that is, data for this address is valid only in the cache and not in system memory.
- **Exclusive**—This cache block holds valid data that is identical to the data at this address in system memory. No other cache has this data.
- **Shared**—Only available if HID2[MESISTATE] register bit is set. The address block is valid in the cache and in at least one other cache. This block is always consistent with system memory. That is, the shared state is shared-unmodified; there is no shared-modified state. Although the MESI protocol is supported by the e300 core, it is not implemented on MPC8315E.
- **Invalid**—This cache block does not hold valid data.

Cache coherency is enforced by on-chip bus snooping logic. Because the e300 core data cache tags are single-ported, a simultaneous load/store and snoop access represents a resource contention. The snoop access is given first access to the tags. The load or store then occurs on the clock following the snoop.

Parity is now integrated into both instruction and data cache memory. A machine check interrupt is now taken upon the detection of an instruction or data cache parity error. Parity is checked whenever valid data is returned from the instruction or data cache for a cache hit or whenever valid data is read out of the cache for a castout or snoop-push operation.

7.4.3.3 Instruction and Data Cache Way-Locking

The e300 core implements instruction and data cache way-locking, which guarantees that certain memory accesses will hit in the cache. This provides deterministic access times for those accesses.

7.4.4 Interrupt Model

This section describes the PowerPC interrupt model and the e300 core implementation specifically.

7.4.4.1 PowerPC Interrupt Model

The PowerPC interrupt mechanism allows the core to change to supervisor state as a result of external signals, errors, or unusual conditions arising in the execution of instructions. The conditions that can cause interrupts are called exceptions. When interrupts occur, information about the state of the core is saved to certain registers and the core begins execution at an address (interrupt vector) predetermined for each interrupt type. Interrupts are processed in supervisor mode.

Some interrupts, such as program interrupts, can be triggered by a broad range of exception conditions. Other interrupts, such as the decremter interrupt, have only a single exception condition. Although multiple exception conditions can map to a single interrupt vector, a more specific condition may be determined by examining a register associated with the interrupt—for example, the DSISR and the FPSCR. Additionally, some exception conditions can be explicitly enabled or disabled by software.

The PowerPC architecture requires that interrupts be handled in program order; therefore, although a particular implementation may recognize exception conditions out of order, they are presented strictly in order. When an instruction-caused interrupt is recognized, any unexecuted instructions that appear earlier in the instruction stream, including any that have not yet entered the execute stage, are required to complete before the interrupt is taken. Any interrupts caused by those instructions are handled first. Likewise, asynchronous, precise interrupts are recognized when they occur, but are not handled until the instruction currently in the completion stage successfully completes execution or generates an interrupt, and the completed store queue is emptied.

Unless a catastrophic condition causes a system reset or machine check interrupt, only one interrupt is handled at a time. If, for example, a single instruction encounters multiple interrupt conditions, those conditions are handled sequentially. After the interrupt handler completes, the instruction execution continues until the next interrupt condition is encountered. However, in many cases there is no attempt to re-execute the instruction. This method of recognizing and handling interrupts sequentially guarantees that interrupts are recoverable.

To prevent the program state from being lost due to a system reset, a machine check interrupt, or an instruction-caused interrupt in the interrupt handler, interrupt handlers should save the information stored in SRR0 and SRR1 early and before enabling external interrupts.

The PowerPC architecture supports four types of interrupts:

- Synchronous, precise—These are caused by instructions. All instruction-caused interrupts are handled precisely; that is, the machine state at the time the interrupt occurs is known and can be completely restored. This means that (excluding the trap and system call interrupts) the address of the faulting instruction is provided to the interrupt handler and neither the faulting instruction nor subsequent instructions in the code stream will complete execution before the interrupt is taken. Once the interrupt is processed, execution resumes at the address of the faulting instruction (or at an alternate address provided by the interrupt handler). When an interrupt is taken due to a trap or system call instruction, execution resumes at an address provided by the handler.
- Synchronous, imprecise—The PowerPC architecture defines two imprecise floating-point exception modes: recoverable and nonrecoverable. Even though the core provides a means to enable the imprecise modes, it implements these modes identically to the precise mode (that is, all enabled floating-point exceptions are always precise on the core).
- Asynchronous, maskable—The external system management interrupt (SMI) and decremter interrupts are maskable, asynchronous interrupts. When these interrupts occur, their handling is postponed until the next instruction and any of its associated interrupts complete execution. If there are no instructions in the execution units, the interrupt is taken immediately upon determination of the correct restart address (for loading SRR0).
- Asynchronous, nonmaskable—The system reset and the machine check interrupt are nonmaskable, asynchronous interrupts. They may not be recoverable, or they may provide a limited degree of recoverability. All interrupts report recoverability through MSR[RI].

7.4.4.2 Implementation-Specific Interrupt Model

As specified by the Power Architecture, all interrupts can be described as either precise or imprecise and either synchronous or asynchronous. Asynchronous interrupts (some of which are maskable) are caused by events external to the processor’s execution; synchronous interrupts, which are all handled precisely by the e300 core, are caused by instructions. A system management interrupt is an implementation-specific interrupt. The interrupt classes are shown in [Table 7-7](#).

Table 7-7. Interrupt Classifications

| Synchronous/Asynchronous | Precise/Imprecise | Interrupt Type |
|---------------------------|-------------------|--|
| Asynchronous, nonmaskable | Imprecise | Machine check System reset |
| Asynchronous, maskable | Precise | External interrupt Decrementer System management interrupt Critical interrupt |
| Synchronous | Precise | Instruction-caused interrupts |

Although interrupts have other characteristics, such as whether they are maskable, the distinctions shown in Table 7-7 define categories of interrupts that the core handles uniquely. Note that Table 7-7 includes no synchronous, imprecise instructions. While the PowerPC architecture supports imprecise handling of floating-point exceptions, the core implements floating-point exception modes as precise.

The e300 core interrupts and exception conditions that cause them are listed in Table 7-8.

Table 7-8. Exceptions and Interrupts

| Interrupt Type | Vector Offset (hex) | Exception Conditions |
|--------------------|---------------------|--|
| Reserved | 00000 | — |
| System reset | 00100 | Caused by the assertion of either \overline{hreset} . |
| Machine check | 00200 | Caused by the assertion of the \overline{tea} signal during a data bus transaction, assertion of \overline{mcp} , an address or data parity error, or an instruction or data cache parity error. Note that the e300 has SRR1 register values that are different from the G2/G2_LE cores' when a machine check occurs. |
| DSI | 00300 | Determined by the bit settings in the DSISR, listed as follows: 1 Set if the translation of an attempted access is not found in the primary hash table entry group (HTEG), or in the rehashed secondary HTEG, or in the range of a DBAT register; otherwise cleared 4 Set if a memory access is not permitted by the page or DBAT protection mechanism; otherwise cleared 6 Set for a store operation and cleared for a load operation 9 Set if a data address breakpoint interrupt occurs when the data [0–28] in the DABR or DABR2 matches the next data access (load or store instruction) to complete in the completion unit. The different breakpoints are enabled as follows: <ul style="list-style-type: none"> • Write breakpoints enabled when DABR[30] is set • Read breakpoints enabled when DABR[31] is set |
| ISI | 00400 | Caused when an instruction fetch cannot be performed for any of the following reasons: <ul style="list-style-type: none"> • The effective (logical) address cannot be translated. That is, there is a page fault for this portion of the translation, so an ISI interrupt must be taken to load the PTE (and possibly the page) into memory. • The fetch access violates memory protection (indicated by SRR1[4] set). If the key bits (Ks and Kp) in the segment register and the PP bits in the PTE are set to prohibit read access, instructions cannot be fetched from this location. |
| External interrupt | 00500 | Caused when MSR[EE] = 1 and the \overline{int} signal is asserted. |
| Alignment | 00600 | Caused when the core cannot perform a memory access for any of the reasons described below: <ul style="list-style-type: none"> • The operand of a floating-point load or store instruction is not word-aligned. • The operands of lmw, stmw, lwarx, and stwcx instructions are not aligned. • The instruction is lswi, lswx, stswi, stswx, and the core is in little-endian mode. Note that PowerPC little-endian mode is not supported on the e300 core. • The operand of dcbz is in memory that is write-through-required or caching-inhibited. |

Table 7-8. Exceptions and Interrupts (continued)

| Interrupt Type | Vector Offset (hex) | Exception Conditions |
|--------------------------------|---------------------|---|
| Program | 00700 | <p>Caused by one of the following exception conditions, which correspond to bit settings in SRR1 and arise during execution of an instruction.</p> <p>Floating-point enabled exception—A floating-point enabled exception condition is generated when the following condition is met: (MSR[FE0] MSR[FE1]) and FPSCR[FEX] is 1.</p> <ul style="list-style-type: none"> • FPSCR[FEX] is set by the execution of a floating-point instruction that causes an enabled exception or by the execution of one of the Move to FPSCR instructions that results in both an exception condition bit and its corresponding enable bit being set in the FPSCR. • Illegal instruction—An illegal instruction program interrupt is generated when execution of an instruction is attempted with an illegal opcode or illegal combination of opcode and extended opcode fields (including PowerPC instructions not implemented in the core), or when execution of an optional instruction not provided in the core is attempted (these do not include those optional instructions that are treated as no-ops). • Privileged instruction—A privileged instruction program interrupt is generated when the execution of a privileged instruction is attempted and the MSR register user privilege bit, MSR[PR], is set. In the e300 core, this interrupt is generated for mtspr or mfspr with an invalid SPR field if SPR[0] = 1 and MSR[PR] = 1. This may not be true for all cores that implement the PowerPC architecture. • Trap—A trap type program interrupt is generated when any of the conditions specified in a trap instruction are met. |
| Floating-point unavailable | 00800 | Caused by an attempt to execute a floating-point instruction (including floating-point load, store, and move instructions) when the floating-point available bit (MSR[FP]) is cleared. |
| Decrementer | 00900 | Occurs when DEC[0] changes from 0 to 1. This interrupt is enabled with MSR[EE]. |
| Critical interrupt | 00A00 | Taken when \overline{cint} is asserted and MSR[CE] = 1. |
| Reserved | 00B00–00BFF | — |
| System call | 00C00 | Occurs when a System Call (sc) instruction is executed. |
| Trace | 00D00 | Taken when MSR[SE] = 1 or when the currently completing instruction is a branch and MSR[BE] = 1. |
| Reserved | 00E00 | The e300 core does not generate an interrupt to this vector. Other devices may use this vector for floating-point assist interrupts. |
| Performance monitor | 00F00 | Caused when a configured PM counter using the pm_event_in to transition overflows. |
| Instruction translation miss | 01000 | Caused when the effective address for an instruction fetch cannot be translated by the ITLB. |
| Data load translation miss | 01100 | Caused when the effective address for a data load operation cannot be translated by the DTLB. |
| Data store translation miss | 01200 | Caused when the effective address for a data store operation cannot be translated by the DTLB, or when a DTLB hit occurs and the change bit in the PTE must be set due to a data store operation. |
| Instruction address breakpoint | 01300 | Occurs when the address (bits 0–29) in the IABR matches the next instruction to complete in the completion unit, and IABR[30] is set. Note that the e300 core also implements IABR2, which functions identically to IABR. |

Table 7-8. Exceptions and Interrupts (continued)

| Interrupt Type | Vector Offset (hex) | Exception Conditions |
|-----------------------------|---------------------|--|
| System management interrupt | 01400 | Caused when MSR[EE] = 1 and the \overline{smi} input signal is asserted. |
| Reserved | 01500–02FFF | — |

7.4.5 Memory Management

The following sections describe the memory management features of the PowerPC architecture and the e300 core implementation, respectively.

7.4.5.1 PowerPC Memory Management

The primary functions of the MMU are to translate logical (effective) addresses to physical addresses for memory accesses and to provide access protection on blocks and pages of memory.

The core generates two types of accesses that require address translation: instruction accesses and data accesses to memory generated by load and store instructions.

The PowerPC MMU and interrupt model support demand-paged virtual memory. Virtual memory management permits execution of programs larger than the size of physical memory; demand-paged implies that individual pages are loaded into physical memory from system memory only when they are first accessed by an executing program.

The hashed page table is a variable-sized data structure that defines the mapping between virtual page numbers and physical page numbers. The page table size is a power of two, and its starting address is a multiple of its size.

The page table contains a number of page-table entry groups (PTEGs). A PTEG contains eight page-table entries (PTEs) of 8 bytes each; therefore, each PTEG is 64 bytes long. PTEG addresses are entry points for table search operations.

Address translations are enabled by setting bits in the MSR—MSR[IR] enables instruction address translations, and MSR[DR] enables data address translations.

7.4.5.2 Implementation-Specific Memory Management

The instruction and data memory management units in the e300 core provide 4 Gbytes of logical address space accessible to supervisor and user programs with a 4-Kbyte page size and 256-Mbyte segment size. Block sizes range from 128 Kbytes to 256 Mbytes and are software selectable. In addition, the core uses an interim 52-bit virtual address and hashed page tables for generating 32-bit physical addresses. The MMUs in the e300 core rely on the interrupt processing mechanism for the implementation of the paged virtual memory environment and for enforcing protection of designated memory areas.

Instruction and data TLBs provide address translation in parallel with the on-chip cache access, incurring no additional time penalty in the event of a TLB hit. A TLB is a cache of the most recently used page table

entries. Software is responsible for maintaining the consistency of the TLB with memory. The core TLBs are 64-entry, two-way, set-associative caches that contain instruction and data address translations. The core provides hardware assist for software table search operations through the hashed page table on TLB misses. Supervisor software can invalidate TLB entries selectively.

For instructions and data that correspond to block address translation, the e300 core provides independent eight-entry BAT arrays. These entries define blocks that can vary from 128 Kbytes to 256 Mbytes. The BAT arrays are maintained by system software. HID2[HBE] is added to the e300 for enabling or disabling the four additional pairs of BAT registers. However, regardless of the setting of HID2[HBE], these BATs are accessible by **mfspr** and **mtspr**.

As specified by the PowerPC architecture, the hashed page table is a variable-sized data structure that defines the mapping between virtual page numbers and physical page numbers. The page table size is a power of two, and its starting address is a multiple of its size.

Also as specified by the PowerPC architecture, the page table contains a number of PTEGs. A PTEG contains 8 PTEs of 8 bytes each; therefore, each PTEG is 64 bytes long. PTEG addresses are entry points for table search operations.

7.4.6 Instruction Timing

The e300 core is a pipelined superscalar processor core. Because instruction processing is reduced into a series of stages, an instruction does not require all of the resources of an execution unit at the same time. For example, after an instruction completes the decode stage, it can pass on to the next stage, while the subsequent instruction can advance into the decode stage. This improves the throughput of the instruction flow. For example, it may take three cycles for a single floating-point instruction to execute, but if there are no stalls in the floating-point pipeline, a series of floating-point instructions can have a throughput of one instruction per cycle.

The core instruction pipeline has four major pipeline stages, described as follows:

- The fetch pipeline stage primarily involves retrieving instructions from the memory system and determining the location of the next instruction fetch. Additionally, if possible, the BPU decodes branches during the fetch stage and folds out branch instructions before the dispatch stage.
- The dispatch pipeline stage is responsible for decoding the instructions supplied by the instruction fetch stage and determining which of the instructions are eligible to be dispatched in the current cycle. In addition, the source operands of the instructions are read from the appropriate register file and dispatched with the instruction to the execute pipeline stage. At the end of the dispatch pipeline stage, the dispatched instructions and their operands are latched by the appropriate execution unit.
- In the execute pipeline stage, each execution unit with an instruction executes the selected instruction (perhaps over multiple cycles), writes the instruction's result into the appropriate rename register, and notifies the completion stage when the execution has finished. In the case of an internal interrupt, the execution unit reports the interrupt to the completion/write-back pipeline stage and discontinues instruction execution until the interrupt is handled. The interrupt is not signaled until that instruction is the next to be completed. Execution of most floating-point instructions is pipelined within the FPU, allowing up to three instructions to execute in the FPU concurrently. The FPU pipeline stages are multiply, add, and round-convert. The LSU has two

pipeline stages: the first stage, for effective address calculation and MMU translation, and the second, for accessing data in the cache.

- The complete/write-back pipeline stage maintains the correct architectural machine state and transfers the contents of the rename registers to the GPRs and FPRs as instructions are retired. If the completion logic detects an instruction causing an interrupt, all subsequent instructions are canceled, their execution results in rename registers are discarded, and instructions are fetched from the correct instruction stream.

A superscalar processor core issues multiple, independent instructions into multiple pipelines, allowing instructions to execute in parallel. The e300c1 core has independent execution units for: integer instructions, floating-point instructions, branch instructions, load/store instructions, and system register instructions. The e300c3 provides two IUs, which improves the throughput of integer instructions. The e300c3 provides two integer units for greater integer instruction throughput along with enhanced multipliers in each IU that reduce the multiply instruction latency to a maximum of two cycles. The IU and the FPU each have dedicated register files for maintaining operands (GPRs and FPRs, respectively), allowing integer and floating-point calculations to occur simultaneously without interference.

The core provides support for single-cycle store, and it provides an adder/comparator in the system register unit that allows the dispatch and execution of multiple integer add and compare instructions on each cycle.

Because the PowerPC architecture can be applied to such a wide variety of implementations, instruction timing among processor cores varies accordingly.

7.4.7 Core Interface

The core interface is specific for each processor core implementation.

The MPC8315E contains an internal coherent system bus (CSB) that interfaces the processor core to the peripheral logic. This internal bus is very similar in function to the external 60x bus interface on the MPC603e. In the case of the MPC8315E, the CSB system logic decodes e300-initiated transactions and directs all accesses to the appropriate interface.

The e300 core can operate at a variety of frequencies allowing the designer to trade off performance for power consumption. The processor core is clocked from a separate PLL, which is referenced to the CSB frequency. This allows the processor core and the peripheral logic to operate at different frequencies.

The e300 core provides a versatile core interface that allows for a wide range of implementations. The interface includes a 32-bit address bus, a 64-bit data bus, and 56 control and information signals (see [Figure 7-5](#)). The core interface allows for address-only transactions, as well as address and data transactions. The core control and information signals include the address arbitration, address start, address transfer, transfer attribute, address termination, data arbitration, data transfer, data termination, and core state signals. Test and control signals provide diagnostics for selected internal circuits.




Figure 7-5. Core Interface

The core interface supports bus pipelining, allowing the address tenure of one transaction to overlap the data tenure of another. The extent of the pipelining depends on external arbitration and control circuitry. Similarly, the core supports split-bus transactions for systems with multiple potential bus masters—one device can have mastership of the address bus while another has mastership of the data bus. Allowing multiple bus transactions to occur simultaneously increases the available bus bandwidth for other activity and, as a result, improves performance.

The core clocking structure allows the bus to operate at integer multiples of the core cycle time.

The following sections describe the core bus support for memory operations. Note that some signals perform different functions depending on the addressing protocol used.

7.4.7.1 Memory Accesses

The e300 core CSB is a 64-bit data bus.

With a 64-bit CSB, memory accesses allow transfer sizes of 8, 16, 24, 32, 40, 48, 56, or 64 bits in one bus clock cycle. Data transfers occur in either single-beat transactions or four-beat burst transactions. Single-beat transactions are caused by noncached accesses that access memory directly (that is, reads and writes when caching is disabled, caching-inhibited accesses, and stores in write-through mode). Four-beat burst transactions, which always transfer an entire cache block (32 bytes), are initiated when a line is read from or written to memory.

7.4.7.2 Signals

The e300 core signals are grouped as follows:

- **Interrupts/Resets**—These signals include the external interrupt signal (\overline{int}), critical interrupt signal (\overline{cint}), checkstop signals, performance monitor signal (pm_event_in) via the PM counters, and both soft reset and hard reset signals. They are used to interrupt and, under various conditions, to reset the core.

- JTAG/debug interface signals—The JTAG (based on the IEEE 1149.1 standard) interface and debug unit provides a serial interface to the system for performing monitoring and boundary tests. Two additional signals are added to the e300 core to allow observation of the internal clock state of the core (*stopped*) and to allow the external input to force the core into a halted state (*ext_halt*).
- Core status and control—These signals include the memory reservation signal, machine quiesce control signals, time base/decrementer clock base enable signal, and the *tlbisync* signal.
- Clock control—These signals provide for system clock input and frequency control.
- Test interface signals—Signals like address matching, combinational matching, and watchpoint are used in the core for production testing.
- Transfer attribute signals—These signals provide information about the type of transfer, such as the transfer size and whether the transaction is bursted, write-through, or cache-inhibited.

7.4.8 Debug Features

Some new debug features are specific to the e300 core. Accesses to the debug facilities are available only in supervisor mode by using the **mtspr** and **mfspr** instructions. The e300 provides the following additional feature in the JTAG/debug interface: Inclusion of breakpoint status and control pins: *stopped* and *ext_halt*.

7.4.8.1 Breakpoint Signaling

The breakpoint signaling provided on the e300 core allows observability of breakpoint matches external to the core. The *iabr*, *iabr2*, *dabr*, and *dabr2* breakpoint signals are asserted for at least one bus clock cycle when the respective breakpoint occurs. The status of the run state of the e300 core is indicated by the *stopped* pin. An asynchronous external breakpoint can be asserted to the e300 core using the *ext_halt* pin:

- When DBCR and IBCR are configured for an OR combinational signal type, the breakpoint signals *iabr*, *iabr2* and *dabr*, *dabr2* reflect their respective breakpoints.
- When the DBCR and IBCR are configured for AND combinational signal type, only the *iabr2* and *dabr2* breakpoint signals are asserted after the AND condition is met (that is, both instruction breakpoints occurred or both data breakpoints occurred).
- When the core_stopped pin is asserted, the e300 core has entered a stopped state and all internal clocking has stopped, indicating that a hardware debug event has occurred.
- The *ext_halt* input pin can be used to force the core into halted state. The halted state may be a hardstop, conditional upon the HARDSTOP condition being set through the JTAG/debug interface

7.5 Differences Between Cores

The e300 core has similar functionality to the G2_LE core. [Table 7-9](#) describes the differences between the G2_LE and the e300.

Table 7-9. Differences Between e300 and G2_LE Cores

| e300 Core | G2_LE Core | Impact |
|---|--|---|
| New HID0 bits | — | The e300 core has a new HID0 bit defined to enable cache parity error reporting (ECPE). |
| New HID1 bits | — | The e300 core has new HID1 bits defined to extend the number of PLL configuration signals to seven (PC5, PC6). |
| New HID2 bits | — | The e300 core has new HID2 bits defined to support instruction fetch bursting (IFEB), MESI coherency protocol (MESI), instruction fetch cancels (IFEC), data cache queue sharing (EBQS), pipelining extension (EBPX), additional cache way locking (IWLCK and DWLCK), and instruction cache way protection (ICWP). |
| New PVR register value | — | The processor version register values differ. |
| New IBCR and DBCR bits | — | The e300 core has new IBCR[IABRSTAT, IABR2STAT] and DBCR[DABR1STAT, DABR2STAT] fields to provide instruction and data address breakpoint status. |
| — | 16-Kbyte, four-way, set-associative, instruction and data caches | Some e300 cores may have different cache sizes than the G2_LE. See the <i>e300 PowerPC Core Reference Manual</i> for detailed information. |
| L1 cache parity | — | The e300 core supports parity for both instruction and data caches; the G2_LE does not support cache parity. |
| MEI or MESI coherency protocols | MEI protocol only | The e300 supports two coherency protocols: MEI and MESI; the G2_LE only supports the MEI protocol. |
| Instruction cancel extension | — | The e300 instruction cancel mechanism improves utilization of instruction cache by supporting 'hits-under-cancels' and 'misses-under-cancels'; the G2_LE requires the cancel to complete before new instruction fetches can begin. |
| Instruction fetch bursts to caching-inhibited space | Single-beat instruction fetches to caching-inhibited space | The e300's instruction fetch burst extension allows all caching-inhibited instruction fetches to be performed on the bus as burst transactions, even though the instructions are not cached. This improves performance for instruction space that is caching-inhibited, because up to eight instructions are returned with one bus operation. The G2_LE core must use single-beat instruction fetches for caching-inhibited space, returning only two instructions per bus operation. |
| Instruction cache way protection | — | The e300 core can protect locked ways in the instruction cache from invalidation; the G2_LE does not support instruction cache way protection. |

Table 7-9. Differences Between e300 and G2_LE Cores (continued)

| e300 Core | G2_LE Core | Impact |
|---------------------------------------|---|---|
| Data cache queue sharing | — | The e300 has a new data cache queue sharing extension that allows the two burst-write queues in the bus unit to be used interchangeably for cache replacements and snoop pushes. Thus, the data cache can support two outstanding cache replacements or two outstanding snoop push operations on the bus at any given time. |
| icbt instruction | — | The e300 supports a new instruction cache block touch instruction that facilitates preloading the instruction cache before locking; the G2_LE core requires speculatively fetching instructions before locking the instruction cache. |
| 1-1/2-level bus pipelining | 1-level bus pipelining | For the e300, a new transaction can complete an address tenure when the previous transaction has been granted the data bus; for the G2_LE, a new transaction must wait until the previous data tenure has completed before completing its address tenure. |
| PowerPC little-endian not supported | PowerPC little-endian supported | PowerPC little-endian will not be supported in the e300 core, although true little-endian will be fully supported. |
| Data retry mode removed | Data retry mode available | <i>drtry</i> and <i>drtrymode</i> will no longer be supported on the e300 and future versions. |
| External control instructions removed | External control instructions available | The eciwx and ecowx instruction pair will not be supported on the e300 core. These are optional instructions in the PowerPC architecture. |
| Reduced pin mode removed | Reduced pin mode available | Reduced pinout mode and the signal <i>redpinmode</i> will not be supported in the e300 core. |

Chapter 8

Integrated Programmable Interrupt Controller (IPIC)

This chapter describes the integrated programmable interrupt controller (IPIC), including a definition of the external signals and their functions. Also, the configuration, control, and status registers are described in this chapter. Note that individual chapters in this reference manual describe specific initialization aspects for each individual block.

8.1 IPIC Introduction

This chapter describes the IPIC interrupt protocol, various types of interrupt sources controlled by the IPIC unit, and the IPIC registers with some programming guidelines. The programming model is similar to the interrupt controller of the MPC8260. The interrupt controller provides interrupt management that is responsible for receiving hardware-generated interrupts from different sources (both internal and external). It also prioritizes and delivers the interrupts to the CPU for servicing. The IPIC prioritizes and manages interrupts from the following controller units:

- DDR memory controller (DDR)
- Enhanced local bus memory controller (eLBC)
- PCI
- Four-channel DMA controller (DMA)
- Message unit (MU)
- Dual three-speed Ethernet controllers (eTSEC1 and eTSEC2)
- DUART communication module (DUART)
- USB 2.0 dual role controller (USB DR)
- Security engine (SEC)
- System bus arbiter (SBA)
- Periodic interval timer (PIT)
- Real time clock timer (RTC ALR and RTC SEC)
- Eight global timers (GTM1–GTM8)
- Software watchdog timer (WDT)
- I²C controller (I²C)
- SPI controller (SPI)
- Power management controller (PMC)
- General-purpose I/O controller (GPIO)
- External pins ($\overline{\text{IRQ}}[0:7]$)
- Serial ATA (SATA) controller (SATA1 and SATA2)

- PCI Express controllers (PCI Express1 and PCI Express2)
- Message Shared Interrupt (MSI)
- Time Division Multiplexing (TDM) interface

The interrupt sources controlled by the IPIC unit cause exceptions in the processor core. The internal interrupt (\overline{int}) signal is the main interrupt output from the IPIC to the core and it causes the regular interrupt exception. The \overline{cint} signal is the critical interrupt output from the IPIC to the processor core and causes the critical interrupt exception. The \overline{smi} signal is the system management interrupt output from the IPIC to the processor core and causes the system management interrupt exception. The machine check exception is caused by the internal \overline{mcp} signal generated by the IPIC, informing the host processor of error conditions, assertion of the external $\overline{IRQ0}$ machine-check request (enabled when $SEMSR[SIRQ0] = 1$), and other conditions.

Table 8-1 shows the relationship of the various functional blocks and external signals of the device to the IPIC unit.

The IPIC receives interrupt request signals from the following two sources:

- External to the integrated device
- Internal to the integrated device

The unit selects the highest priority interrupt from all current interrupts and forwards it to the internal processor core, or off-chip for external servicing.

The IPIC also manages an internal non-maskable machine-check processor (\overline{mcp}) signal and the interrupt generated by the off-chip interrupt sources ($\overline{IRQ}[0:7]$).

The interrupt router of the IPIC monitors the outputs of the internal configuration registers. When the priority is highest in one of the received interrupt signals, the IPIC sets the corresponding bit in one of the interrupt pending registers—system internal interrupt pending register (SIPNR)/system external interrupt pending register (SEPNR). If the interrupt is not masked, the IPIC asserts the \overline{int} signal to indicate an interrupt request to the processor. When the processor is running the specific \overline{int} , \overline{cint} , or \overline{smi} interrupt handler code, the processor must vectorize the external interrupt handler by explicitly (in software) reading the corresponding interrupt vector register (SIVCR, SCVCR or SMVCR). In response to this read, the IPIC unit returns the vector (associated with the interrupt source) to the interrupt handler routine. In addition, the handler can vectorize different branches of interrupt handling.




Figure 8-1. Interrupt Sources Block Diagram

The IPIC receives the following types of interrupts:

- External interrupt—triggered by the off-chip signals (\overline{IRQn}) listed in [Table 8-1](#)
- Internal interrupts—on-chip interrupts, triggered by the sources listed in [Table 8-7](#) and [Table 8-9](#)
- External and internal non-maskable machine check conditions, signaled by the sources listed in [Table 8-23](#) through \overline{mcp}

The interrupt controller provides the ability to mask each interrupt source. Any source that can be caused by multiple events are also maskable.

When the IPIC receives an internal or external interrupt, its configuration register is checked to determine if it should be routed off-chip (to the external $\overline{PCI_INTA}$) or serviced as a normal external interrupt by the processor core (through the \overline{int} signal). As a third alternative, if the incoming interrupt has been configured as a critical or system management interrupt, the IPIC completes the processing of the interrupt by asserting \overline{cint} or \overline{smi} to the core. The assertion of the \overline{cint} or \overline{smi} signal to the core causes the interrupt to be serviced as a critical or a system management interrupt, respectively.

8.2 IPIC Features

The IPIC unit implements the following features:

- Functional and programming compatibility with the MPC8260 interrupt controller
- Support for external and internal discrete vectorized interrupt sources
- Support for external and internal non-maskable machine check conditions, signaled by \overline{mcp}
- Programmable highest priority request (can be programmed to support a critical (\overline{cint}) or system management interrupt (\overline{smi}) type)
- Two programmable priority mixed groups of four on-chip and four external interrupt signals with two priority schemes for each group: grouped and spread
- Four programmable priority internal groups of eight on-chip interrupt signals with two priority schemes for each group: grouped and spread
- Two highest priority interrupts from each group can be programmed to support a critical or system management interrupt type
- External and internal interrupts directed to host processor
- Unique vector number for each interrupt source

8.3 IPIP Modes of Operation

The IPIC unit can operate in the core enable or core disable mode.

8.3.1 Core Enable Mode

In core enable mode, all internal interrupts (including those from the PCI block) are routed to and from the IPIC; the interrupts are sent to the PowerPC core. The DMA controller can optionally (depending on the programming of the DMA registers) steer its interrupt to the PCI host through the $\overline{PCI_INTA}$ signal.

In this mode all machine check interrupts are gathered by the IPIC unit and sent to the PowerPC core. If the device performs as a PCI host, the interrupts of the other PCI agents should be connected to the implementation's $\overline{\text{IRQx}}$ signals and treated like normal external interrupts (sent to the core).

8.3.2 Core Disable Mode

In core disable mode, all internal interrupts (including those from the PCI block) are routed to and from the IPIC, the interrupts are then sent through the $\overline{\text{PCI_INTA}}$ signal to the PCI host CPU. Note that the core interrupt signal is masked. The user should use in this mode only the $\overline{\text{int}}$ output interrupt type (should not use $\overline{\text{cint}}$ or $\overline{\text{smi}}$ output interrupt types) to read an updated SIVCR. (See Section 8.5.9, “System Internal Interrupt Control Register (SICNR),” and Section 8.5.14, “System External Interrupt Control Register (SECNR).”)

In this mode, machine check interrupts are driven either on $\overline{\text{PCI_INTA}}$ or on $\overline{\text{MCP_OUT}}$ as level-sensitive interrupts. $\text{SERCRCR}[\text{MCPR}]$ (see Section 8.5.17, “System Error Control Register (SERCRCR)”) controls which external signal is used.

8.4 IPIC External Signal Description

The following sections provide an overview and detailed descriptions of the IPIC signals.

8.4.1 IPIC External Signals Overview

The device has 8 distinct external interrupt request input signals ($\overline{\text{IRQ}}[0:7]$) and one interrupt request output signal ($\overline{\text{PCI_INTA}}$). The IPIC interface signals are defined in Table 8-1.

Table 8-1. IPIC Signal Properties

| Name | Port | Function | I/O | Reset | Requires Pull Up |
|-------------------------------|-------------------------------|--------------------------|-----|-------|------------------|
| $\overline{\text{IRQ}}[0:7]$ | $\overline{\text{IRQ}}[0:7]$ | External interrupts | I | — | Yes |
| $\overline{\text{PCI_INTA}}$ | $\overline{\text{PCI_INTA}}$ | Interrupt request output | O | Z | Yes |
| $\overline{\text{MCP_OUT}}$ | $\overline{\text{MCP_OUT}}$ | Interrupt request output | O | Z | Yes |

8.4.2 IPIC Detailed Signal Descriptions

Table 8-2 provides detailed descriptions of the external IPIC signals.

Table 8-2. IPIC External Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|------------------------------|-----|--|
| $\overline{\text{IRQ}}[0:7]$ | I | Interrupt request 0–7. The sense (level or edge) of each of these signals is programmable. All of these inputs can be driven completely asynchronously. |
| | | State Meaning Asserted—When an external interrupt request signal is asserted the priority is checked by the IPIC unit, and the interrupt is conditionally passed to the processor. Negated—There is no incoming interrupt from that source. |
| | | Timing Assertion—All of these inputs can be asserted completely asynchronously. Negation—Interrupts programmed as level-sensitive must remain asserted until serviced. |

Table 8-2. IPIC External Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|----------|-----|--|
| PCI_INTA | OD | Interrupt request out. Active-low, open drain. When the IPIC is programmed in core disable mode, this output reflects the raw interrupts generated by on-chip sources. See Section 8.3, “IPIPIC Modes of Operation,” for details. |
| | | State Meaning Asserted—At least one interrupt is currently being signalled to the external system. Negated—Indicates no interrupt source currently routed to $\overline{\text{IRQ_OUT}}$. |
| | | Timing Because external interrupts are asynchronous with respect to the system clock, both assertion and negation of $\overline{\text{IRQ_OUT}}$ occur asynchronously with respect to the interrupt source. All timing given here is approximate. Assertion—Internal interrupt source: 3 system bus clock cycles after the interrupt occurs. External interrupt source: 4 cycles after the interrupt occurs. Negation—Follows interrupt source negation with the following delay: Internal interrupt: 3 system bus clock cycles. External interrupt: 4 cycles. |
| MCP_OUT | OD | Non-maskable Interrupt (machine check) request out. Active-low, open drain. When the IPIC is programmed in core disable mode, this output reflects the <i>mcp</i> interrupts generated by on-chip sources. See Section 8.3, “IPIPIC Modes of Operation.” |
| | | State Meaning Asserted—At least one machine check interrupt is currently being signalled to the external system. Negated—Indicates no interrupt source currently routed to MCP_OUT. |
| | | Timing Because external interrupts are asynchronous with respect to the system clock, both assertion and negation of MCP_OUT occurs asynchronously with respect to the interrupt source. All timing given here is approximate. Assertion—Internal interrupt source: 2 system bus clock cycles after interrupt occurs. External interrupt source: 4 cycles after interrupt occurs. Negation—Follows interrupt source negation with the following delay: Internal interrupt: 2 system bus clock cycles. External interrupt: 4 cycles. |

8.5 IPIC Memory Map/Register Definition

The IPIC programmable register map occupies 256 bytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect.

All IPIC registers are 32-bits wide and they are located on 32-bit address boundaries. Software can perform byte, half-word or word accesses to any IPIC registers. All addresses used in this chapter are offsets from the IPIC base, as defined in [Chapter 2, “Memory Map.”](#)

[Table 8-3](#) shows the memory map of the IPIC unit.

Table 8-3. IPIC Register Address Map

| Offset | Register | Access | Reset Value | Section/ Page |
|---|--|--------|-------------|----------------------------|
| Integrated Programmable Interrupt Controller—Block Base Address 0x0_0700 | | | | |
| 0x00 | System global interrupt configuration register (SICFR) | R/W | All zeros | 8.5.1/8-8 |
| 0x04 | System regular interrupt vector register (SIVCR) | R | All zeros | 8.5.2/8-9 |
| 0x08 | System internal interrupt pending register (SIPNR_H) | R | All zeros | 8.5.3/8-12 |
| 0x0C | System internal interrupt pending register (SIPNR_L) | R | All zeros | 8.5.3/8-12 |

Table 8-3. IPIC Register Address Map (continued)

| Offset | Register | Access | Reset Value | Section/ Page |
|-----------|---|--------|-------------|-----------------------------|
| 0x10 | System internal interrupt group A priority register (SIPRR_A) | R/W | 0x0530_9770 | 8.5.4/8-15 |
| 0x14 | System internal interrupt group B priority register (SIPRR_B) | R/W | 0x0530_9770 | 8.5.5/8-16 |
| 0x18 | System internal interrupt group C priority register (SIPRR_C) | R/W | 0x0530_9770 | 8.5.6/8-16 |
| 0x1C | System internal interrupt group D priority register (SIPRR_D) | R/W | 0x0530_9770 | 8.5.7/8-17 |
| 0x20 | System internal interrupt mask register (SIMSR_H) | R/W | All zeros | 8.5.8/8-18 |
| 0x24 | System internal interrupt mask register (SIMSR_L) | R/W | All zeros | 8.5.8/8-18 |
| 0x28 | System internal interrupt control register (SICNR) | R/W | All zeros | 8.5.9/8-19 |
| 0x2C | System external interrupt pending register (SEP NR) | R/W | Special | 8.5.10/8-21 |
| 0x30 | System mixed interrupt group A priority register (SMPRR_A) | R/W | 0x0530_9770 | 8.5.11/8-22 |
| 0x34 | System mixed interrupt group B priority register (SMPRR_B) | R/W | 0x0530_9770 | 8.5.12/8-23 |
| 0x38 | System external interrupt mask register (SEMSR) | R/W | All zeros | 8.5.13/8-23 |
| 0x3C | System external interrupt control register (SECNR) | R/W | All zeros | 8.5.14/8-24 |
| 0x40 | System error status register (SERSR) | R/W | All zeros | 8.5.15/8-26 |
| 0x44 | System error mask register (SERMR) | R/W | | 8.5.16/8-27 |
| 0x48 | System error control register (SERCR) | R/W | All zeros | 8.5.17/8-27 |
| 0x4C | System external interrupt polarity control register (SEPCR) | R/W | 0x0000_0000 | 8.5.18/8-28 |
| 0x4F | Reserved | — | — | — |
| 0x50 | System internal interrupt force register (SIFCR_H) | R/W | All zeros | 8.5.19/8-29 |
| 0x54 | System internal interrupt force register (SIFCR_L) | R/W | All zeros | 8.5.19/8-29 |
| 0x58 | System external interrupt force register (SEFCR) | R/W | All zeros | 8.5.20/8-30 |
| 0x5C | System error force register (SERFR) | R/W | All zeros | 8.5.21/8-30 |
| 0x60 | System critical interrupt vector register (SCVCR) | R | All zeros | 8.5.22/8-31 |
| 0x64 | System management interrupt vector register (SMVCR) | R | All zeros | 8.5.23/8-31 |
| 0x68–0xBF | Reserved | — | — | — |

8.5.1 System Global Interrupt Configuration Register (SICFR)

SICFR, shown in [Figure 8-2](#), defines the highest priority interrupt and whether interrupts are grouped or spread in the priority table. See [Table 8-4](#) for more information.




Figure 8-2. System Global Interrupt Configuration Register (SICFR)

[Table 8-4](#) defines the bit fields of SICFR.

Table 8-4. SICFR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0 | — | Write ignored, read = 0 |
| 1–7 | HPI | Highest priority interrupt. Specifies the 7-bit unique interrupt number/vector (see Table 8-6) of the single interrupt controller interrupt source that is advanced to the highest priority in the IPIC priority table (see Table 8-34). HPI can be modified dynamically. |
| 8 | — | Write ignored, read = 0 |
| 9 | MPSB | Mixed interrupts priority scheme for group B. Selects the relative MIXB priority scheme. It cannot be changed dynamically. 0 Grouped. The MIXBs are grouped by priority at the top of the table. 1 Spread. The MIXBs are spread by priority in the table. |
| 10 | MPSA | Mixed interrupts priority scheme for group A. Selects the relative MIXA priority scheme. It cannot be changed dynamically. 0 Grouped. The MIXAs are grouped by priority at the top of the table. 1 Spread. The MIXAs are spread by priority in the table. |
| 11 | — | Write ignored, read = 0 |
| 12 | IPSD | Internal interrupts priority scheme for group D. Selects the relative SYSD priority scheme. It cannot be changed dynamically. 0 Grouped. The SYSDs are grouped by priority at the top of the table. 1 Spread. The SYSDs are spread by priority in the table. |
| 13 | IPSC | Internal interrupts priority scheme for group C. Selects the relative SYSC priority scheme. It cannot be changed dynamically. 0 Grouped. The SYSCs are grouped by priority at the top of the table. 1 Spread. The SYSCs are spread by priority in the table. |
| 14 | IPSB | Internal interrupts priority scheme for group B. Selects the relative SYSB priority scheme. It cannot be changed dynamically. 0 Grouped. The SYSBs are grouped by priority at the top of the table. 1 Spread. The SYSBs are spread by priority in the table. |
| 15 | IPSA | Internal interrupts priority scheme for group A. Selects the relative SYSA priority scheme. It cannot be changed dynamically. 0 Grouped. The SYSAs are grouped by priority at the top of the table. 1 Spread. The SYSAs are spread by priority in the table. |

Table 8-5. SIVCR Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 6–24 | — | Write ignored, read = 0 |
| 25–31 | IVEC | Regular interrupt vector. Specifies a 7-bit unique number of the IPIC's highest priority regular interrupt source, pending to the core. Note that the when a regular interrupt request occurs, SIVCR can be read. If there are multiple regular interrupt sources, SIVCR latches the highest priority regular interrupt. Note that the IVEC field correctly reflects all interrupt vectors (see Table 8-6 for details). The value of SIVCR cannot change while it is being read. |

[Table 8-6](#) shows the definition of IVEC.

Table 8-6. IVEC/CVEC/MVEC Field Definition

| Interrupt ID Number | Interrupt Meaning | Interrupt Vector |
|---------------------|----------------------|-----------------------|
| 0 | Error (no interrupt) | 0b000_0000 |
| 1 | PEX1 CNT | 0b000_0001 |
| 2 | PEX2 CNT | 0b000_0010 |
| 3 | DMAC | 0b000_0011 |
| 4 | MSIR1 | 0b0000_0100 |
| 5–8 | Reserved | 0b000_0101–0b000_1000 |
| 9 | UART1 | 0b000_1001 |
| 10 | UART2 | 0b000_1010 |
| 11 | SEC | 0b000_1011 |
| 12 | eTSEC1 1588 timer | 0b000_1100 |
| 13 | eTSEC2 1588 timer | 0b000_1101 |
| 14 | I2C | 0b000_1110 |
| 15 | Reserved | 0b000_1111 |
| 16 | SPI | 0b001_0000 |
| 17 | IRQ1 | 0b001_0001 |
| 18 | IRQ2 | 0b001_0010 |
| 19 | IRQ3 | 0b001_0011 |
| 20 | IRQ4 | 0b001_0100 |
| 21 | IRQ5 | 0b001_0101 |
| 22 | IRQ6 | 0b001_0110 |
| 23 | IRQ7 | 0b001_0111 |
| 24–31 | Reserved | 0b001_1000–0b001_1111 |
| 32 | TSEC1 Tx | 0b010_0000 |
| 33 | TSEC1 Rx | 0b010_0001 |
| 34 | TSEC1 Err | 0b010_0010 |
| 35 | TSEC2 TX | 0b010_0011 |

Table 8-6. IVEC/CVEC/MVEC Field Definition (continued)

| Interrupt ID Number | Interrupt Meaning | Interrupt Vector |
|---------------------|-------------------|-----------------------|
| 36 | TSEC2 Rx | 0b010_0100 |
| 37 | TSEC2 Err | 0b010_0101 |
| 38 | USB DR | 0b010_0110 |
| 39 | Reserved | 0b010_0111 |
| 40 | TDM Tx | 0b010_1000 |
| 41 | TDM Rx | 0b010_1001 |
| 42–43 | Reserved | 0b010_1010–0x010_1011 |
| 44 | SATA1 | 0b010_1100 |
| 45 | SATA2 | 0b010_1101 |
| 46–47 | Reserved | 0x010_1110–0x010_1111 |
| 48 | IRQ0 | 0b011_0000 |
| 49–63 | Reserved | 0b011_0001–0b011_1111 |
| 64 | RTC SEC | 0b100_0000 |
| 65 | PIT | 0b100_0001 |
| 66 | PCI | 0b100_0010 |
| 67 | MSIR0 | 0b100_0011 |
| 68 | RTC ALR | 0b100_0100 |
| 69 | MU | 0b100_0101 |
| 70 | SBA | 0b100_0110 |
| 71 | DMA | 0b100_0111 |
| 72 | GTM4 | 0b100_1000 |
| 73 | GTM8 | 0b100_1001 |
| 74 | GPIO | 0b100_1010 |
| 75 | Reserved | 0b100_1011 |
| 76 | DDR | 0b100_1100 |
| 77 | LBC | 0b100_1101 |
| 78 | GTM2 | 0b100_1110 |
| 79 | GTM6 | 0b100_1111 |
| 80 | PMC | 0b101_0000 |
| 81 | MSIR2 | 0b101_0001 |
| 82 | MSIR3 | 0b101_0010 |
| 83 | TDM Err | 0b101_0011 |
| 84 | GTM3 | 0b101_0100 |
| 85 | GTM7 | 0b101_0101 |
| 86 | MSIR4 | 0b101_0110 |

Table 8-6. IVEC/CVEC/MVEC Field Definition (continued)

| Interrupt ID Number | Interrupt Meaning | Interrupt Vector |
|---------------------|-------------------|-----------------------|
| 87 | MSIR5 | 0b101_0111 |
| 88 | MSIR6 | 0b101_1000 |
| 89 | MSIR7 | 0b101_1001 |
| 90 | GTM1 | 0b101_1010 |
| 91 | GTM5 | 0b101_1011 |
| 92–93 | Reserved | 0b101_1100–0b101_1101 |
| 94 | DMAC Err | 0b101_1110 |
| 95–127 | Reserved | 0b101_1111–0b111_1111 |

8.5.3 System Internal Interrupt Pending Registers (SIPNR_H and SIPNR_L)

Each bit in SIPNR_H and SIPNR_L, shown in [Figure 8-4](#) and [Figure 8-5](#), may be assigned an internal interrupt source. (Implemented bits are listed in [Table 8-7](#).) When an interrupt request is received, the interrupt controller sets the corresponding SIPNR bit. When a pending interrupt is handled, the user clears the SIPNR bit by clearing the corresponding event register bit.

Note that SIPNR bit positions are not changed according to relative priority.




Figure 8-4. System Internal Interrupt Pending Register (SIPNR_H)

[Table 8-7](#) lists implemented SIPNR_H fields. Note that these field descriptions are also valid for SIFCR_H and SIMSR_H.

Table 8-7. SIPNR_H/SIFCR_H/SIMSR_H Bit Assignments

| Bits | Field |
|------|-----------|
| 0 | TSEC1 Tx |
| 1 | TSEC1 Rx |
| 2 | TSEC1 Err |
| 3 | TSEC2 Tx |
| 4 | TSEC2 Rx |
| 5 | TSEC2 Err |
| 6 | USB DR |
| 7 | — |
| 8 | TDM Tx |

Table 8-7. SIPNR_H/SIFCR_H/SIMSR_H Bit Assignments (continued)

| Bits | Field |
|-------|-------------------|
| 9 | TDM Rx |
| 10–11 | — |
| 12 | SATA1 |
| 13 | SATA2 |
| 14–15 | — |
| 16 | PEX1 CNT |
| 17 | PEX2 CNT |
| 18 | DMAC |
| 19 | MSIR1 |
| 20-23 | — |
| 24 | UART1 |
| 25 | UART2 |
| 26 | SEC |
| 27 | eTSEC1 1588 timer |
| 28 | eTSEC2 1588 timer |
| 29 | I2C |
| 30 | — |
| 31 | SPI |

Table 8-8 defines the bit fields of SIPNR_H.

Table 8-8. SIPNR_H Field Descriptions

| Bits | Name | Description |
|------|------------------|--|
| 0–31 | INT _n | Each implemented bit (listed in Table 8-7) corresponds to an internal interrupt source. When an interrupt is received, the interrupt controller sets the corresponding SIPNR bit. When a pending interrupt is handled, the user clears the SIPNR bit by clearing the corresponding event register bit. SIPNR bits are read only. Writing to this register has no effect. Note that the SIPNR bit positions are not changed according to their relative priority. For unimplemented bits, writes are ignored, read = 0. |

SIPNR_L is shown in Figure 8-4.




Figure 8-5. System Internal Interrupt Pending Register (SIPNR_L)

Table 8-9 lists implemented SIPNR_L fields. Note that these field assignments are also valid for SIFCR_L and SIMSR_L.

Table 8-9. SIPNR_L/SIFCR_L/SIMSR_L Bit Assignments

| Bits | Field |
|-------|----------|
| 0 | RTC SEC |
| 1 | PIT |
| 2 | PCI |
| 3 | MSIR0 |
| 4 | RTC ALR |
| 5 | MU |
| 6 | SBA |
| 7 | DMA |
| 8 | GTM4 |
| 9 | GTM8 |
| 10 | GPIO |
| 11 | — |
| 12 | DDR |
| 13 | LBC |
| 14 | GTM2 |
| 15 | GTM6 |
| 16 | PMC |
| 17 | MSIR2 |
| 18 | MSIR3 |
| 19 | TDM Err |
| 20 | GTM3 |
| 21 | GTM7 |
| 22 | MSIR4 |
| 23 | MSIR5 |
| 24 | MSIR6 |
| 25 | MSIR7 |
| 26 | GTM1 |
| 27 | GTM5 |
| 28-29 | — |
| 30 | DMAC Err |
| 31 | — |

Table 8-10 defines the bit fields of SIPNR_L.

Table 8-10. SIPNR_L Field Descriptions

| Bits | Name | Description |
|------|---------|--|
| 0–31 | INT n | Each implemented bit (listed in Table 8-9) corresponds to an internal interrupt source. When an interrupt is received, the interrupt controller sets the corresponding SIPNR bit. When a pending interrupt is handled, the user clears the SIPNR bit by clearing the corresponding event register bit. SIPNR bits are read only. Writing to this register has no effect. Note that the SIPNR bit positions are not changed according to their relative priority. For unimplemented bits, writes are ignored, read = 0. |

8.5.4 System Internal Interrupt Group A Priority Register (SIPRR_A)

SIPRR_A, shown in Figure 8-6, defines the priority between TSEC1 transmit request (TSEC1 Tx), TSEC1 receive request (TSEC1 Rx), TSEC1 transmit/receive error (TSEC1 Err), TSEC2 transmit request (TSEC2 Tx), TSEC2 receive request (TSEC2 Rx), TSEC2 transmit/receive error (TSEC2 Err), USB DR, internal interrupt signals.

For more information, see Section 8.6.3, “Internal Interrupts Group Relative Priority.”




Figure 8-6. System Internal Interrupt Group A Priority Register (SIPRR_A)

Table 8-11 defines the bit fields of SIPRR_A.

Table 8-11. SIPRR_A Field Descriptions

| Bits | Name | Description |
|--------------|---------------|--|
| 0–2 | SYSA0P | SYSA0 priority order. Defines which interrupt source asserts its request in the SYSA0 priority position. The user should not program the same code to multiple priority positions (0–7). These bits can be changed dynamically. The definition of SYSA0P is as follows: 000 TSEC1 Tx asserts its request in the SYSA0 position. 001 TSEC1 Rx asserts its request in the SYSA0 position. 010 TSEC1 Err asserts its request in the SYSA0 position. 011 TSEC2 Tx asserts its request in the SYSA0 position. 100 TSEC2 Rx asserts its request in the SYSA0 position. 101 TSEC2 Err asserts its request in the SYSA0 position. 110 USB DR asserts its request in the SYSA0 position. 111 Reserved |
| 3–11, 16–27 | SYSA1P–SYSA7P | Same as SYSA0P, but for SYSA1P–SYSA7P. |
| 12–15, 28–31 | — | Write ignored, read = 0 |

8.5.5 System Internal Interrupt Group B Priority Register (SIPRR_B)

The system internal interrupt group B priority register (SIPRR_B), shown in Figure 8-7, defines the priority between internal interrupt signals.

For more information, see Section 8.6.3, “Internal Interrupts Group Relative Priority.”




Figure 8-7. System Internal Interrupt Group B Priority Register (SIPRR_B)

Table 8-12 defines the bit fields of SIPRR_B.

Table 8-12. SIPRR_B Field Descriptions

| Bits | Name | Description |
|-----------------|-------------------|---|
| 0–2 | SYSB0P | SYSB0 Priority order. Defines which interrupt source asserts its request in the SYSB0 priority position. The user should not program the same code to more than one priority position (0–7). These bits can be changed dynamically. The definition of SYSB0P is shown as follows: 000 TDM Tx asserts its request in the SYSB0 position. 001 TDM Rx asserts its request in the SYSB0 position. 010 Reserved 011 Reserved 100 SATA1 asserts its request in the SYSB0 position. 101 SATA2 asserts its request in the SYSB0 position. 110 Reserved 111 Reserved |
| 3–11, 16–27 | SYSB1P– SYSB7P | Same as SYSB0P, but for SYSB1P–SYSB7P. |
| 12–15, 28–31 | — | Write ignored, read = 0 |

8.5.6 System Internal Interrupt Group C Priority Register (SIPRR_C)

The system internal interrupt group C priority register (SIPRR_C), shown in Figure 8-8, defines the priority between internal interrupt signals.

For more information about interrupt priorities, see Section 8.6.3, “Internal Interrupts Group Relative Priority.”

Offset 0x18

Access: Read/write



Figure 8-8. System Internal Interrupt Group C Priority Register (SIPRR_C)

Table 8-13 defines the bit fields of SIPRR_C.

Table 8-13. SIPRR_C Field Descriptions

| Bits | Name | Description |
|-----------------|-------------------|---|
| 0–2 | SYSC0P | SYSC0 priority order. Defines which interrupt source asserts its request in the SYSC0 priority position. The user should not program the same code to more than one priority position (0–7). These bits can be changed dynamically. The definition of SYSC0P is shown as follows: 000 PEX1 CNT asserts its request in the SYSC0 position. 001 PEX2 CNT asserts its request in the SYSC0 position. 010 DMAC asserts its request in the SYSC0 position. 011 MSIR1 asserts its request in the SYSC0 position 100–111 Reserved |
| 3–11, 16–27 | SYSC1P– SYSC7P | Same as SYSC0P, but for SYSC1P–SYSC7P. |
| 12–15, 28–31 | — | Write ignored, read = 0 |

8.5.7 System Internal Interrupt Group D Priority Register (SIPRR_D)

SIPRR_D, shown in Figure 8-9, defines the priority among the interrupt sources listed in Table 8-14.

Offset 0x1C

Access: Read/write



Figure 8-9. System Internal Interrupt Group D Priority Register (SIPRR_D)

Table 8-14 defines the bit fields of SIPRR_D.

Table 8-14. SIPRR_D Field Descriptions

| Bits | Name | Description |
|--------------|---------------|---|
| 0–2 | SYSD0P | SYSD0 priority order. Defines which interrupt source asserts its request in the SYSD0 priority position. The user should not program the same code to more than one priority position (0–7). These bits can be changed dynamically. SYSD0P is defined as follows: 000 UART1 asserts its request in the SYSD0 position. 001 UART2 asserts its request in the SYSD0 position. 010 SEC asserts its request in the SYSD0 position. 011 eTSEC1 1588 timer asserts its request in the SYSD0 position. 100 eTSEC2 1588 timer asserts its request in the SYSD0 position. 101 I2C asserts its request in the SYSD0 position. 110 Reserved 111 SPI asserts its request in the SYSD0 position. |
| 3–11, 16–27 | SYSD1P–SYSD7P | Same as SYSD0P, but for SYSD1P–SYSD7P. |
| 12–15, 28–31 | — | Write ignored, read = 0 |

8.5.8 System Internal Interrupt Mask Register (SIMSR_H and SIMSR_L)

Each implemented bit in SIMSR_H and SIMSR_L, shown in Figure 8-10 and Figure 8-11, corresponds to an internal interrupt source. The user masks an interrupt by clearing the corresponding SIMSR bit. When an interrupt request occurs, the corresponding SIPNR bit is set, regardless of the SIMSR bit. However, if the corresponding SIMSR bit is cleared, no interrupt request is passed to the core.

When an SIMSR bit is cleared by the user at the same time corresponding interrupt source requests an interrupt service, the request stops. If the user sets the SIMSR bit later, the core processes any pending corresponding interrupt requests according to its priority.




Figure 8-10. System Internal Interrupt Mask Register (SIMSR_H)

Table 8-15 defines the bit fields of SIMSR_H.

Table 8-15. SIMSR_H Field Descriptions

| Bits | Name | Description |
|------|------------------|---|
| 0–31 | INT _n | <p>Each implemented bit (listed in Table 8-7) corresponds to an external interrupt source. The user masks an interrupt by clearing the corresponding SIMSR bit. An interrupt is unmasked (enable) by setting the corresponding SIMSR bit. The SIMSR can be read by the user at any time.</p> <p>Note:</p> <ul style="list-style-type: none"> SIMSR bit positions do not change according to their relative priority. The user can clear pending register bits that were set by multiple interrupt events only by clearing all unmasked events in the corresponding event register. If an SIMSR bit is masked at the same time that the corresponding SIPNR bit causes an interrupt request to the core, the error vector is issued (if no other interrupts are pending). Thus, the user should always include an error vector routine, even if it contains only an <code>rfi</code> instruction. The error vector cannot be masked. <p>Unimplemented bits, shown as reserved in Table 8-7, are ignored on writes; read = 0.</p> |

Figure 8-11 shows SIMSR_L.



Figure 8-11. System Internal Interrupt Mask Register (SIMSR_L)

Table 8-16 defines the bit fields of SIMSR_L.

Table 8-16. SIMSR_L Field Descriptions

| Bits | Name | Description |
|------|------------------|---|
| 0–31 | INT _n | <p>Each implemented bit (listed in Table 8-9) corresponds to an external interrupt source. The user masks an interrupt by clearing the corresponding SIMSR bit. An interrupt is unmasked (enabled) by setting the corresponding SIMSR bit. The SIMSR can be read by the user at any time.</p> <p>Note:</p> <ul style="list-style-type: none"> SIMSR bit positions are not changed according to their relative priority. The user can clear pending register bits that were set by multiple interrupt events only by clearing all unmasked events in the corresponding event register. If an SIMSR bit is masked at the same time that the corresponding SIPNR bit causes an interrupt request to the core, the error vector is issued (if no other interrupts are pending). Thus, the user should always include an error <p>Unimplemented bits, shown as reserved in Figure 8-11, are ignored on writes; read = 0.</p> |

8.5.9 System Internal Interrupt Control Register (SICNR)

SICNR, shown in Figure 8-12, defines the IPIC output interrupt type ($\overline{\text{int}}$, $\overline{\text{cint}}$, or $\overline{\text{smi}}$) in the SYSA0–SYSA1, SYSB0–SYSB1, SYSC0–SYSC1, and SYSD0–SYSD1 priority positions. All other priority positions assert $\overline{\text{int}}$ to the core.

Note that in core disabled mode the user should use the \overline{int} output interrupt type (should not use \overline{cint} or \overline{smi} output interrupt types) to read an updated SIVCR.

Offset 0x28

Access: Read/write




Figure 8-12. System Internal Interrupt Control Register (SICNR)

Table 8-17 defines the bit fields of SICNR.

Table 8-17. SICNR Field Descriptions

| Bits | Name | Description |
|-------|--------|--|
| 0–1 | SYSD0T | SYSD0 priority position IPIC output interrupt type. Defines which type of the IPIC output interrupt signal (\overline{int} , \overline{cint} , or \overline{smi}) asserts its request to the core in the SYSD0 priority position. These bits cannot be changed dynamically. (to change it, software must make sure the corresponding interrupt source is masked or it cannot happen during the change). The definition of SYSD0T is as follows: 00 \overline{int} request is asserted to the core for SYSD0. 01 \overline{smi} request is asserted to the core for SYSD0. 10 \overline{cint} request is asserted to the core for SYSD0. 11 Reserved |
| 2–3 | SYSD1T | Same as SYSD0T, but for SYSD1T. |
| 4–7 | — | Write ignored, read = 0 |
| 8–9 | SYSC0T | SYSC0 priority position IPIC output interrupt Type. Defines which type of the IPIC output interrupt signal (\overline{int} , \overline{cint} , or \overline{smi}) asserts its request to the core in the SYSC0 priority position. These bits can not be changed dynamically. (If s/w really wants to change it, it has to make sure the corresponding interrupt source is masked or it won't happen during the change). The definition of SYSC0T is as follows: 00 \overline{int} request is asserted to the core for SYSC0. 01 \overline{smi} request is asserted to the core for SYSC0. 10 \overline{cint} request is asserted to the core for SYSC0. 11 Reserved |
| 10–11 | SYSC1T | Same as SYSC0T, but for SYSC1T. |
| 12–15 | — | Write ignored, read = 0 |
| 16–17 | SYSB0T | SYSB0 priority position IPIC output interrupt Type. Defines which type of the IPIC output interrupt signal (\overline{int} , \overline{cint} , or \overline{smi}) asserts its request to the core in the SYSB0 priority position. These bits can not be changed dynamically. (If s/w really wants to change it, it has to make sure the corresponding interrupt source is masked or it won't happen during the change). The definition of SYSB0T is as follows: 00 \overline{int} request is asserted to the core for SYSB0. 01 \overline{smi} request is asserted to the core for SYSB0. 10 \overline{cint} request is asserted to the core for SYSB0. 11 Reserved |
| 18–19 | SYSB1T | Same as SYSB0T, but for SYSB1T. |
| 20–23 | — | Write ignored, read = 0 |

Table 8-17. SICNR Field Descriptions (continued)


| Bits | Name | Description |
|-------|--------|--|
| 24–25 | SYSA0T | <p>SYSA0 priority position IPIC output interrupt type. Defines which type of the IPIC output interrupt signal (\overline{int}, \overline{smi}, or \overline{cint}) asserts its request to the core in the SYSA0 priority position. These bits can not be changed dynamically. (If s/w really wants to change it, it must ensure the corresponding interrupt source is masked or it does not happen during the change).</p> <p>The definition of SYSA0T is as follows:</p> <p>00 \overline{int} request is asserted to the core for SYSA0.</p> <p>01 \overline{smi} request is asserted to the core for SYSA0.</p> <p>10 \overline{cint} request is asserted to the core for SYSA0.</p> <p>11 Reserved.</p> |
| 26–27 | SYSA1T | Same as SYSA0T, but for SYSA1T |
| 28–31 | — | Write ignored, read = 0 |

8.5.10 System External Interrupt Pending Register (SEPNR)


Each bit in the SEPNR, shown in Figure 8-13, corresponds to an external interrupt source. When an interrupt is received, the interrupt controller sets the corresponding SEPNR bit.

Offset 0x2C

Access: Read/write



Reset The reset values of implemented bits reflect the values of the external IRQ signals. Reserved bits are zeros. ²



Reset All zeros

¹ This bit is valid only if the IRQ0 signal is configured as an external maskable interrupt (SEMSR[SIRQ0] = 0)

² The user should drive all IRQ inputs to an inactive state prior to reset negation

Figure 8-13. System External Interrupt Pending Register (SEPNR)

Table 8-18 defines the bit fields of SEPNR.

Table 8-18. SEPNR Field Descriptions

| Bits | Name | Description |
|------|------------------|---|
| 0–7 | IRQ _n | Each bit corresponds to an external interrupt source. When an external interrupt is received, the interrupt controller sets the corresponding SEPNR bit. When a pending interrupt is handled, the user must clear the corresponding SEPNR bit. For level triggered cases, the software needs to cause the IRQ _n to negate which automatically clears the bit in SEPNR. For edge-triggered cases, the software needs to clear the corresponding bit in SEPNR. SEPNR bits are cleared by writing ones to them. Because the user can only clear bits in this register, writing zeros to this register has no effect. Note that the SEPNR bit positions are not changed according to their relative priority. |
| 8–31 | — | Write ignored, read = 0 |

8.5.11 System Mixed Interrupt Group A Priority Register (SMPRR_A)

The SMPRR_A, shown in Figure 8-14, defines the priority among the sources listed in Table 8-19.




Figure 8-14. System Mixed Interrupt Group A Priority Register (SMPRR_A)

Table 8-19 defines the bit fields of SMPRR_A.

Table 8-19. SMPRR_A Field Descriptions

| Bits | Name | Description |
|--------------|---------------|--|
| 0–2 | MIXA0P | MIXA0 priority order. Defines which interrupt source asserts its request in the MIXA0 priority position. The user must not program the same code to more than one priority position (0–7). These bits can be changed dynamically. The definition of MIXA0P is as follows: 000 RTC SEC asserts its request to the MIXA0 position. 001 PIT asserts its request to the MIXA0 position. 010 PCI asserts its request to the MIXA0 position. 011 MSIR0 asserts its request to the MIXA0 position. 100 IRQ0 asserts its request to the MIXA0 position. This field for MIXA0 position is valid (must not be ignored) if IRQ0 signal configured as an external maskable interrupt (SEMSR[SIRQ0] = 0). 101 IRQ1 asserts its request to the MIXA0 position. 110 IRQ2 asserts its request to the MIXA0 position. 111 IRQ3 asserts its request to the MIXA0 position. |
| 3–11, 16–27 | MIXA1P–MIXA7P | Same as MIXA0P, but for MIXA1P–MIXA7P. |
| 12–15, 28–31 | — | Write ignored, read = 0 |

8.5.12 System Mixed Interrupt Group B Priority Register (SMPRR_B)

SMPRR_B, shown in [Figure 8-15](#), defines the priority among the sources listed in [Table 8-20](#).




Figure 8-15. System Mixed Interrupt Group B Priority Register (SMPRR_B)

[Table 8-20](#) defines the bit fields of SMPRR_B.

Table 8-20. SMPRR_B Field Descriptions

| Bits | Name | Description |
|-----------------------|------------|--|
| 0–2 3–11, 16–27 | MIXB n P | MIXB n priority order. Defines which interrupt source asserts its request in the MIXB n priority position. The user must not program the same code to more than one priority position (0–7). These bits can be changed dynamically. The definition of MIXB n P is as follows: 000 RTC ALR asserts its request to the MIXB n position. 001 MU asserts its request to the MIXB n position. 010 SBA asserts its request to the MIXB n position. 011 DMA asserts its request to the MIXB n position. 100 IRQ4 asserts its request to the MIXB n position. 101 IRQ5 asserts its request to the MIXB n position. 110 IRQ6 asserts its request to the MIXB n position. 111 IRQ7 asserts its request to the MIXB n position. |
| 12–15, 28–31 | — | Write ignored, read = 0 |

8.5.13 System External Interrupt Mask Register (SEMSR)

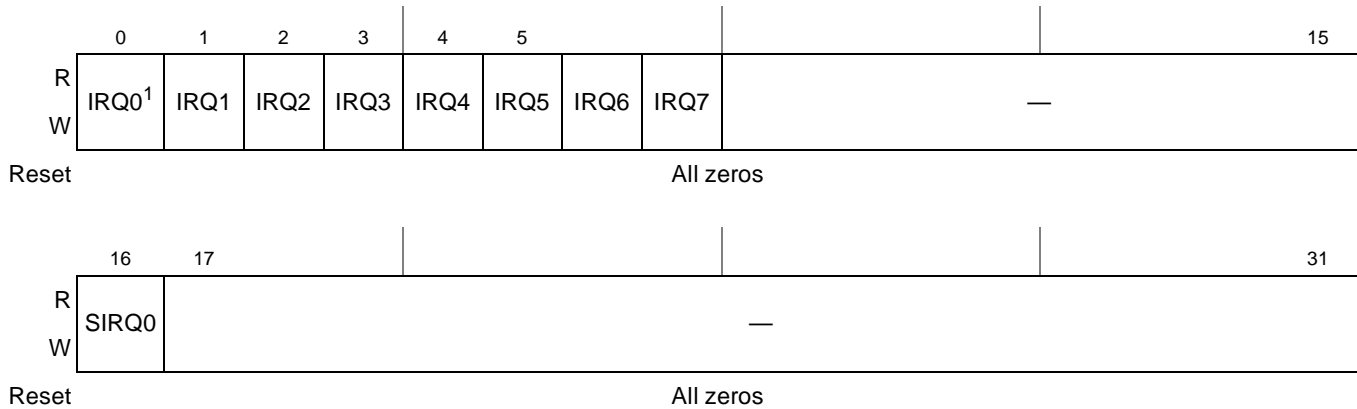
Each bit in the system external interrupt mask register (SEMSR), shown in [Figure 8-13](#), corresponds to an external interrupt source. The user masks an interrupt by clearing the corresponding SEMSR bit. An interrupt is unmasked (enabled) by setting the corresponding SEMSR bit.

When an external interrupt request occurs, the corresponding SEP n R bit is set regardless of the setting of the corresponding SEMSR bit. However, if the corresponding SEMSR bit is cleared, no interrupt request is passed to the core.

When an SEMSR bit is cleared by the user at the same time that an interrupt source requests an interrupt service, the request stops. If the user sets the SEMSR bit later, a previously pending interrupt request is processed by the core according to its assigned priority. SEMSR can be read by the user at any time.

Offset 0x38

Access: Read/write



¹ This bit is valid only if the IRQ0 signal is configured as an external maskable interrupt (SEMSR[SIRQ0] = 0)
²

Figure 8-16. System External Interrupt Mask Register (SEMSR)

Table 8-21 defines the bit fields of SEMSR.

Table 8-21. SEMSR Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0–7 | — | Each bit corresponds to an external interrupt source. The user masks an interrupt by clearing the SEMSR bit. An interrupt can be enabled by setting the corresponding SEMSR bit. SEMSR can be read by the user at any time. Note: <ul style="list-style-type: none"> SEMSR bit positions are not affected by their relative priority. The user can clear pending register bits that were set by multiple interrupt events only by clearing all unmasked events in the corresponding event register. If an SEMSR bit is masked at the same time that the corresponding SEPNR bit causes an interrupt request to the core, the error vector is issued (if no other interrupts pending). Thus, the user must always include an error vector routine, even if it contains only an <i>rfi</i> instruction. The error vector cannot be masked. |
| 8–15 | — | Write ignored, read = 0 |
| 16 | SIRQ0 | Steer IRQ0. 0 IRQ0 is used as external interrupt request 1 IRQ0 is used as external MCP request |
| 17–31 | — | Write ignored, read = 0 |

8.5.14 System External Interrupt Control Register (SECNR)

SECNR, shown in Figure 8-17, defines the edge detect mode for external \overline{IRQn} interrupt signals and determines whether the corresponding \overline{IRQn} signal asserts an interrupt request upon either a high-to-low change or assertion on the pin. It also defines the IPIC output interrupt type (\overline{int} , \overline{cint} , or \overline{smi}) in the MIXA0–MIXA1 and MIXB0–MIXB1 priority positions.

Note that in core disabled mode of operation the user should use the \overline{int} output interrupt type (should not use \overline{cint} or \overline{smi} output interrupt types) in order to read an updated SIVCR.

Offset 0x3C

Access: Read/write




Figure 8-17. System External Interrupt Control Register (SECNR)

Table 8-22 defines the bit fields of SECNR.

Table 8-22. SECNR Field Descriptions

| Bits | Name | Description |
|-------|--------|--|
| 0–1 | MIXB0T | MIXB0 priority position IPIC output interrupt type. Defines which type of the IPIC output interrupt signal (\overline{int} , \overline{cint} , or \overline{smi}) asserts its request to the core in the MIXB0 priority position. These bits can be changed dynamically. The definition of MIXB0T is as follows: 00 \overline{int} request is asserted to the core for MIXB0. 01 \overline{smi} request is asserted to the core for MIXB0. 10 \overline{cint} request is asserted to the core for MIXB0. 11 Reserved |
| 2–3 | MIXB1T | Same as MIXB0T, but for MIXB1T. |
| 4–7 | — | Write ignored, read = 0 |
| 8–9 | MIXA0T | MIXA0 priority position IPIC output interrupt Type. Defines which type of the IPIC output interrupt signal (\overline{int} , \overline{cint} , or \overline{smi}) asserts its request to the core in the MIXA0 priority position. These bits can be changed dynamically. The definition of MIXA0T is as follows: 00 \overline{int} request is asserted to the core for MIXA0. 01 \overline{smi} request is asserted to the core for MIXA0. 10 \overline{cint} request is asserted to the core for MIXA0. 11 Reserved |
| 10–11 | MIXA1T | Same as MIXA0T, but for MIXA1T. |
| 12–15 | — | Write ignored, read = 0 |
| 16–23 | EDIx | Each bit defines the edge detect mode for the external \overline{IRQn} interrupt signals, determines whether the corresponding \overline{IRQn} signal asserts an interrupt request upon either a high-to-low (high assertion for active high polarity) change or low assertion (high assertion for active high polarity) on the pin. The corresponding \overline{IRQn} signal asserts an interrupt request as follows: 0 Low assertion (high assertion for active high polarity) on \overline{IRQn} generates an interrupt request (level sensitive). 1 High-to-low (high assertion for active high polarity) change on \overline{IRQn} generates an interrupt request (edge sensitive). |
| 24–31 | — | Write ignored, read = 0 |

8.5.15 System Error Status Register (SERSR)

The bits in the SERSR, shown in [Figure 8-18](#), correspond to the external and internal non-maskable error source machine check (mcp) conditions listed in [Table 8-23](#). When an error interrupt signal is received, the interrupt controller sets the corresponding SERSR bit.




Figure 8-18. System Error Status Register (SERSR)

[Table 8-23](#) lists the implemented SERSR bits. Note that these field assignments are valid for SERMR and SERFR.

Table 8-23. SERSR/SERMR/SERFR Bit Assignments

| Bits | Field |
|-------|-------------------|
| 0 | IRQ0 ¹ |
| 1 | WDT |
| 2 | SBA |
| 3 | — |
| 4 | — |
| 5 | PCI |
| 6 | — |
| 7 | MU |
| 8–14 | — |
| 15 | — |
| 16–31 | — |

¹ This bit is valid only if the IRQ0 signal is configured as an external MCP interrupt (SEMSR[SIRQ0] = 1)

[Table 8-24](#) defines the bit fields of SERSR.

Table 8-24. SERSR Field Descriptions

| Bits | Name | Description |
|------|---------|--|
| 0–31 | INT n | Each implemented bit in the SERSR, listed in Table 8-23 , corresponds to an external and an internal error source (mcp). When an error interrupt signal is received, the interrupt controller sets the corresponding SERSR bit. SERSR bits are cleared by writing ones to them. Unmasked event register bits should be cleared before clearing SERSR bits. Because the user can only clear bits in this register, writing zeros to this register has no effect. SERSR bits are cleared by power-on reset. Subsequent soft and hard resets do not affect SERSR bit states. For unimplemented bits (listed as reserved in Table 8-23), writes are ignored, read = 0 |

8.5.16 System Error Mask Register (SERMR)

Each implemented bit in SERMR, shown in [Figure 8-19](#), corresponds to an external and an internal \overline{mcp} source (MCP). The user masks an MCP by clearing and enables an interrupt by setting the corresponding SERMR bit. When a masked MCP occurs, the corresponding SERSR bit is set, regardless of the setting of the corresponding SERMR bit although no MCP request is passed to the core in this case. The SERMR can be read by the user at any time.




Figure 8-19. System Error Mask Register (SERMR)

[Table 8-25](#) defines the bit fields of SERMR.

Table 8-25. SERMR Field Descriptions

| Bits | Name | Description |
|------|---------|---|
| 0–31 | INT n | Each implemented SERMR bit, listed in Table 8-23 , corresponds to an external and an internal MCP source. The user masks an MCP by clearing and enables an interrupt by setting the corresponding SERMR bit. When a masked MCP occurs, the corresponding SERSR bit is set, regardless of the setting of the SERMR bit although no MCP request is passed to the core. The SERMR can be read by the user at any time. Writes to unimplemented (reserved) bits are ignored; read = 0 |

8.5.17 System Error Control Register (SERCR)

SERCR, shown in [Figure 8-20](#), defines the control bits that route MCP requests in core disable mode to either $\overline{MCP_OUT}$ or $\overline{PCI_INTA}$ in core-disable mode.




Figure 8-20. System Error Control Register (SERCR)

Table 8-26 defines the bit fields of SERCR.

Table 8-26. SERCR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–30 | — | Write ignored, read = 0 |
| 31 | MCPR | MCP route. Route MCP request to either $\overline{\text{MCP_OUT}}$ or $\overline{\text{PCI_INTA}}$ (in core disable mode). 0 MCP routed to $\overline{\text{PCI_INTA}}$ (in core disable mode). 1 MCP routed to $\overline{\text{MCP_OUT}}$ (in core disable mode). |

8.5.18 System External interrupt Polarity Control Register (SEPCR)

SEPCR, shown in Figure 8-21, defines the polarity for each one of the external $\overline{\text{IRQ}}_n$ interrupt signals and determines whether the corresponding $\overline{\text{IRQ}}_n$ signal is treated as active low or active high signal. The active low signals will assert an interrupt request upon either a high-to-low change or assertion (low state) on the pin. The active high signals will asserts an interrupt request upon either a low-to-high change or assertion (high state) on the pin. See Section 8.5.14, “System External Interrupt Control Register (SECNR),” for more details.

NOTE

Note that the $\overline{\text{IRQ}}_n$ signals are overbarred although the SEPCR could be programmed to accept active high signals. The overbar should be ignored in this case.




Figure 8-21. System External Interrupt Polarity Control Register (SEPCR)

Table 8-27 defines the bit fields of SEPCR.

Table 8-27. SEPCR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–7 | EIPx | Each bit defines the active state for the external $\overline{\text{IRQ}}_n$ interrupt signals. 0 Active Low. 1 Active High. |
| 8–31 | — | Write ignored, read = 0 |

8.5.19 System Internal Interrupt Force Registers (SIFCR_H and SIFCR_L)

Each implemented bit SIFCR_H and SIFCR_L, shown in [Figure 8-22](#) and [Figure 8-23](#), corresponds to an internal interrupt source. When a bit is set, the interrupt controller generates the corresponding interrupt (sets the corresponding SIPNR bit). The SIFCR can be read by the user at any time.



Figure 8-22. System Internal Interrupt Force Register (SIFCR_H)

[Table 8-28](#) defines the bit fields of SIFCR_H.

Table 8-28. SIFCR_H Field Descriptions

| Bits | Name | Description |
|------|---------|--|
| 0-31 | INT n | Each implemented bit, listed in Table 8-7 , corresponds to an internal interrupt source. The user forces an interrupt by setting the corresponding SIFCR x bit. SIFCR n bit positions are not changed according to their relative priority. Writes to unimplemented (reserved) bits are ignored; read = 0 |

SIFCR_L is shown in [Figure 8-23](#).




Figure 8-23. System Internal Interrupt Force Register (SIFCR_L)

[Table 8-29](#) defines the bit fields of SIFCR_L.

Table 8-29. SIFCR_L Field Descriptions

| Bits | Name | Description |
|------|---------|--|
| 0-31 | INT n | Each implemented bit, listed in Table 8-9 , corresponds to an internal interrupt source. The user forces an interrupt by setting the corresponding SIFCR x bit. SIFCR x bit positions are not changed according to their relative priority. Writes to unimplemented (reserved) bits are ignored; read = 0 |

8.6.2 Interrupt Configuration

Figure 8-28 shows the interrupt configuration.

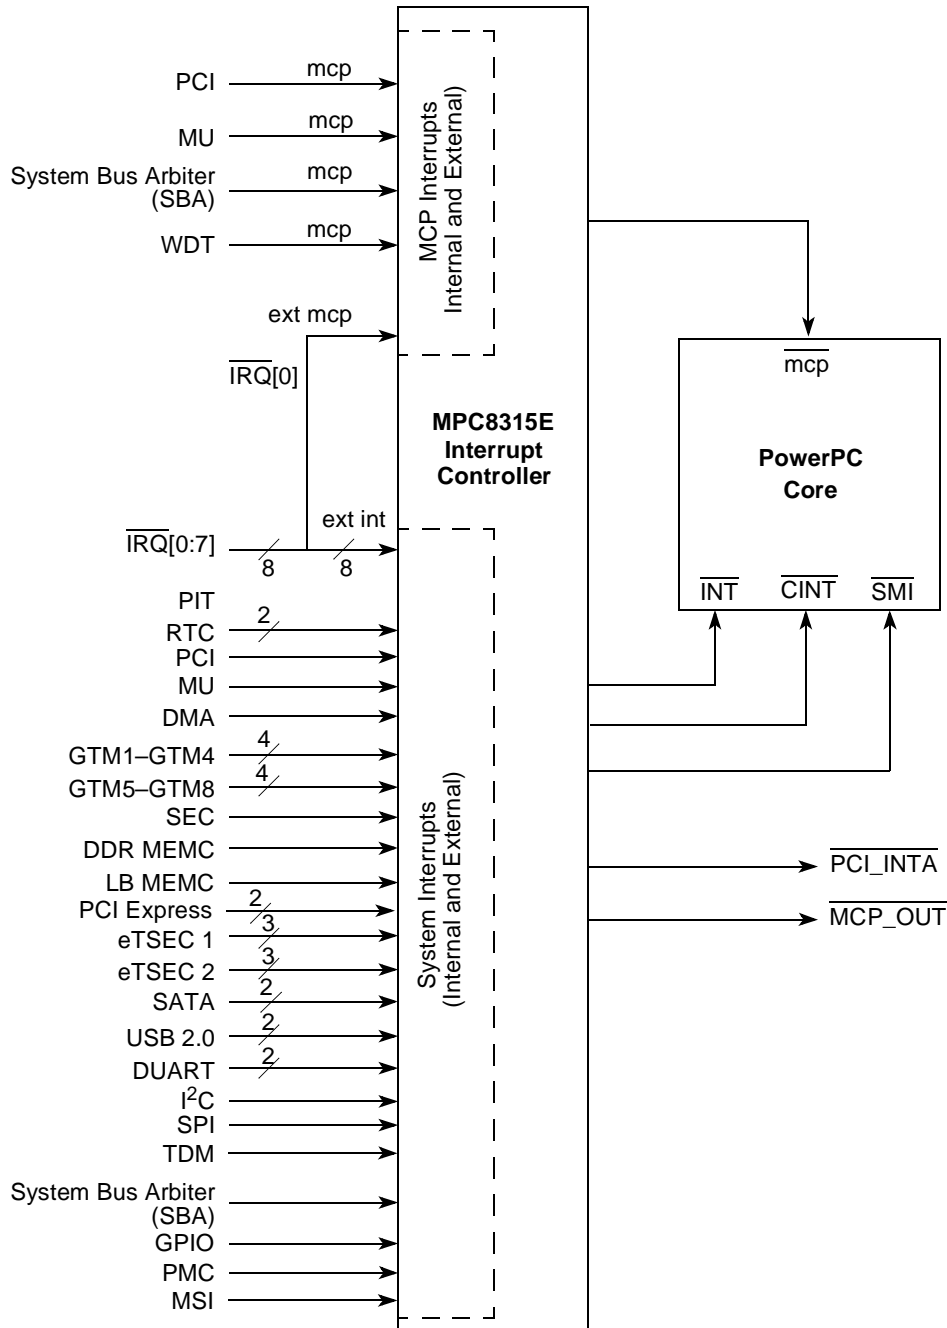


Figure 8-28. Interrupt Structure

The interrupt controller allows masking of each interrupt source. When an unmasked interrupt source is pending in the SIPNR register, the interrupt controller sends an interrupt request to the core. When an interrupt is taken, the interrupt mask bit in the machine state register is cleared to disable further interrupt requests to the PowerPC core until software can handle them.

All interrupt sources are prioritized and bits are set in the system interrupt pending register (SIPNR, SEPNR) as interrupts occur regardless of whether they are masked in the IPIC. The prioritization of the interrupt sources is flexible within the following groups:

- The relative priority of the eTSEC1 Tx, eTSEC1 Rx, eTSEC1 Err, eTSEC2 TX, eTSEC2 Rx, TSEC2 Err, and USB DR internal interrupt signals can be modified.
- The relative priority of the TDM, SATA1, and SATA2 internal interrupt signals can be modified.
- The relative priority of the PCI Express1, PCI Express2, DMAC, and MSIR1 internal interrupts can be modified.
- The relative priority of the UART1, UART2, SPI, I2C, eTSEC1 1588 timer, eTSEC2 1588 timer, and SEC internal interrupt signals can be modified.
- The relative priority of the IRQ0, IRQ1, IRQ2, and IRQ3 external interrupts, and RTC SEC, PCI, and MSIR0 internal interrupts can be modified.
- The relative priority of the IRQ4, IRQ5, IRQ6, and IRQ7 external interrupts, and RTC ALR, MU, SBA, and DMA internal interrupts can be modified.
- One interrupt source can be assigned to be the programmable highest priority.

The SIVCR is updated with a 7-bit vector corresponding to the sub-block with the highest current priority.

8.6.3 Internal Interrupts Group Relative Priority

The relative priority in each internal group is programmable and can be changed dynamically. The group priorities are programmed in the IPIC internal interrupt priority registers (SIPRR_x) and can be changed dynamically to implement a rotating priority.

In addition, the grouping of the locations of the interrupt entries has the following two options:

- Grouped. In the group scheme, all interrupts are grouped together at the top of [Table 8-34](#), ahead of most other interrupt sources. This scheme is ideal for applications where all interrupt sources function at a very high data rate and interrupt latency is very important.
- Spread. In the spread scheme, priorities are spread over [Table 8-34](#) so other sources can have lower interrupt latencies. This scheme is also programmed but cannot be changed dynamically.

8.6.4 Mixed Interrupts Group Relative Priority

The relative priority between up to four internal and four external interrupts in each group is programmable and can be changed dynamically. The group priorities are programmed in the IPIC mixed interrupt priority registers (SMPRR_x) and can be changed dynamically to implement a rotating priority.

In addition, the grouping of the locations of the mixed interrupt entries has the following two options:

- Grouped. In the group scheme, all interrupts are grouped together at the top of the priority table, ahead of most other interrupt sources. See [Table 8-34](#) for more information. This scheme is ideal for applications where all interrupt sources function at a very high data rate and interrupt latency is very important.
- Spread. In the spread scheme, priorities are spread over the table so other sources can have lower interrupt latencies. This scheme is also programmed but cannot be changed dynamically.

8.6.5 Highest Priority Interrupt

In addition to the group relative priority option, SICFR[HPI] can be used to specify one interrupt source as having the highest priority. This interrupt remains within the same interrupt level as the other interrupt controller interrupts, but is serviced before any other interrupt in [Table 8-34](#).

If the highest priority feature is not used, the IPIC selects the interrupt request in MIXA0 to be the highest priority interrupt and the standard interrupt priority order is used from [Table 8-34](#). SICFR[HPI] can be updated dynamically to allow the user to change a normally low priority source into a high priority-source for a period as needed.

8.6.6 Interrupt Source Priorities

Each of the IPIC's internal and external interrupt sources can independently assert one interrupt request to the core. [Table 8-34](#) shows the prioritization of these interrupt sources. As described in previous sections, flexibility exists in the relative ordering of the interrupts, but, in general, relative priorities are as shown. A single interrupt priority number is associated with each table entry.

Table 8-34. Interrupt Source Priority Levels

| Priority | Interrupt Source Description |
|----------|------------------------------|
| 1 | Highest |
| 2 | MIXA0 (Grouped/Spread) |
| 3 | MIXA1 (Grouped) |
| 4 | MIXA2 (Grouped) |
| 5 | MIXA3 (Grouped) |
| 6 | MIXB0 (Spread) |
| 7 | SYSB0 (Grouped) |
| 8 | SYSB1 (Grouped) |
| 9 | SYSB2 (Grouped) |
| 10 | SYSB3 (Grouped) |
| 11 | MIXA1 (Spread) |
| 12 | SYSB4 (Grouped) |
| 13 | SYSB5 (Grouped) |
| 14 | SYSB6 (Grouped) |
| 15 | SYSB7 (Grouped) |
| 16 | MIXB0 (Grouped) |
| 17 | MIXB1 (Grouped) |
| 18 | MIXB2 (Grouped) |
| 19 | MIXB3 (Grouped) |
| 20 | MIXB1 (Spread) |
| 21 | SYSA0 (Grouped) |
| 22 | SYSA1 (Grouped) |

Table 8-34. Interrupt Source Priority Levels (continued)

| Priority | Interrupt Source Description |
|----------|------------------------------|
| 23 | SYSA2 (Grouped) |
| 24 | SYSA3 (Grouped) |
| 25 | MIXA2 (Spread) |
| 26 | SYSA4 (Grouped) |
| 27 | SYSA5 (Grouped) |
| 28 | SYSA6 (Grouped) |
| 29 | SYSA7 (Grouped) |
| 30 | MIXA4 (Grouped) |
| 31 | MIXA5 (Grouped) |
| 32 | MIXA6 (Grouped) |
| 33 | MIXA7 (Grouped) |
| 34 | MIXB2 (Spread) |
| 35 | SYSC0 (Grouped) |
| 36 | SYSC1 (Grouped) |
| 37 | SYSC2 (Grouped) |
| 38 | SYSC3 (Grouped) |
| 39 | MIXA3 (Spread) |
| 40 | SYSC4 (Grouped) |
| 41 | SYSC5 (Grouped) |
| 42 | SYSC6 (Grouped) |
| 43 | SYSC7 (Grouped) |
| 44 | MIXB4 (Grouped) |
| 45 | MIXB5 (Grouped) |
| 46 | MIXB6 (Grouped) |
| 47 | MIXB7 (Grouped) |
| 48 | MIXB3 (Spread) |
| 49 | SYSD0 (Grouped) |
| 50 | SYSD1 (Grouped) |
| 51 | SYSD2 (Grouped) |
| 52 | SYSD3 (Grouped) |
| 53 | MIXA4 (Spread) |
| 54 | SYSD4 (Grouped) |
| 55 | SYSD5 (Grouped) |
| 56 | SYSD6 (Grouped) |
| 57 | SYSD7 (Grouped) |
| 58 | MIXB4 (Spread) |

Table 8-34. Interrupt Source Priority Levels (continued)

| Priority | Interrupt Source Description |
|----------|------------------------------|
| 59 | GTM4 |
| 60 | SYSB0 (Spread) |
| 61 | SYSA0 (Spread) |
| 62 | GTM8 |
| 63 | SYSC0 (Spread) |
| 64 | SYSD0 (Spread) |
| 65 | Reserved |
| 66 | GPIO |
| 67 | MIXA5 (Spread) |
| 68 | Reserved |
| 69 | SYSB1 (Spread) |
| 70 | SYSA1 (Spread) |
| 71 | DDR |
| 72 | SYSC1 (Spread) |
| 73 | SYSD1 (Spread) |
| 74 | Reserved |
| 75 | LBC |
| 76 | MIXB5 (Spread) |
| 77 | GTM2 |
| 78 | SYSB2 (Spread) |
| 79 | SYSA2 (Spread) |
| 80 | GTM6 |
| 81 | SYSC2 (Spread) |
| 82 | SYSD2 (Spread) |
| 83 | Reserved |
| 84 | PMC |
| 85 | MIXA6 (Spread) |
| 86 | MSIR2 |
| 87 | SYSB3 (Spread) |
| 88 | SYSA3 (Spread) |
| 89 | MSIR3 |
| 90 | SYSC3 (Spread) |
| 91 | SYSD3 (Spread) |
| 92 | Reserved |
| 93 | TDM Err |
| 94 | MIXB6 (Spread) |

Table 8-34. Interrupt Source Priority Levels (continued)

| Priority | Interrupt Source Description |
|----------|------------------------------|
| 95 | GTM3 |
| 96 | SYSB4 (Spread) |
| 97 | SYSA4 (Spread) |
| 98 | GTM7 |
| 99 | SYSC4 (Spread) |
| 100 | SYSD4 (Spread) |
| 101 | Reserved |
| 102 | MSIR4 |
| 103 | MIXA7 (Spread) |
| 104 | MSIR5 |
| 105 | SYSB5 (Spread) |
| 106 | SYSA5 (Spread) |
| 107 | MSIR6 |
| 108 | SYSC5 (Spread) |
| 109 | SYSD5 (Spread) |
| 110 | Reserved |
| 111 | MSIR7 |
| 112 | MIXB7 (Spread) |
| 113 | GTM1 |
| 114 | SYSB6 (Spread) |
| 115 | SYSA6 (Spread) |
| 116 | GTM5 |
| 117 | SYSC6 (Spread) |
| 118 | SYSD6 (Spread) |
| 119 | Reserved |
| 120 | Reserved |
| 121 | Reserved |
| 122 | SYSB7 (Spread) |
| 123 | SYSA7 (Spread) |
| 124 | DMAC Err |
| 125 | SYSC7 (Spread) |
| 126 | SYSD7 (Spread) |
| 127 | Reserved |
| 128 | Reserved |

8.6.7 Masking Interrupt Sources

By programming the system interrupt mask registers, SIMSR_x and SEMSR, the user can mask interrupt requests to the core. Each SIMSR_x and SEMSR bit corresponds to an interrupt source. To enable an interrupt, set the corresponding SIMSR or SEMSR bit. When a masked interrupt source has a pending interrupt request, the corresponding SIPNR_x or SEMSR bit is set, even though the interrupt is not generated to the core. The user can mask all interrupt sources to implement a polling interrupt servicing scheme.

When an interrupt source has multiple interrupting events, the user can individually mask these events by programming a mask register within that particular block. Table 8-34 shows which interrupt sources have multiple interrupting events.

Figure 8-29 shows an example of how the masking occurs, using a DDR block as an example.

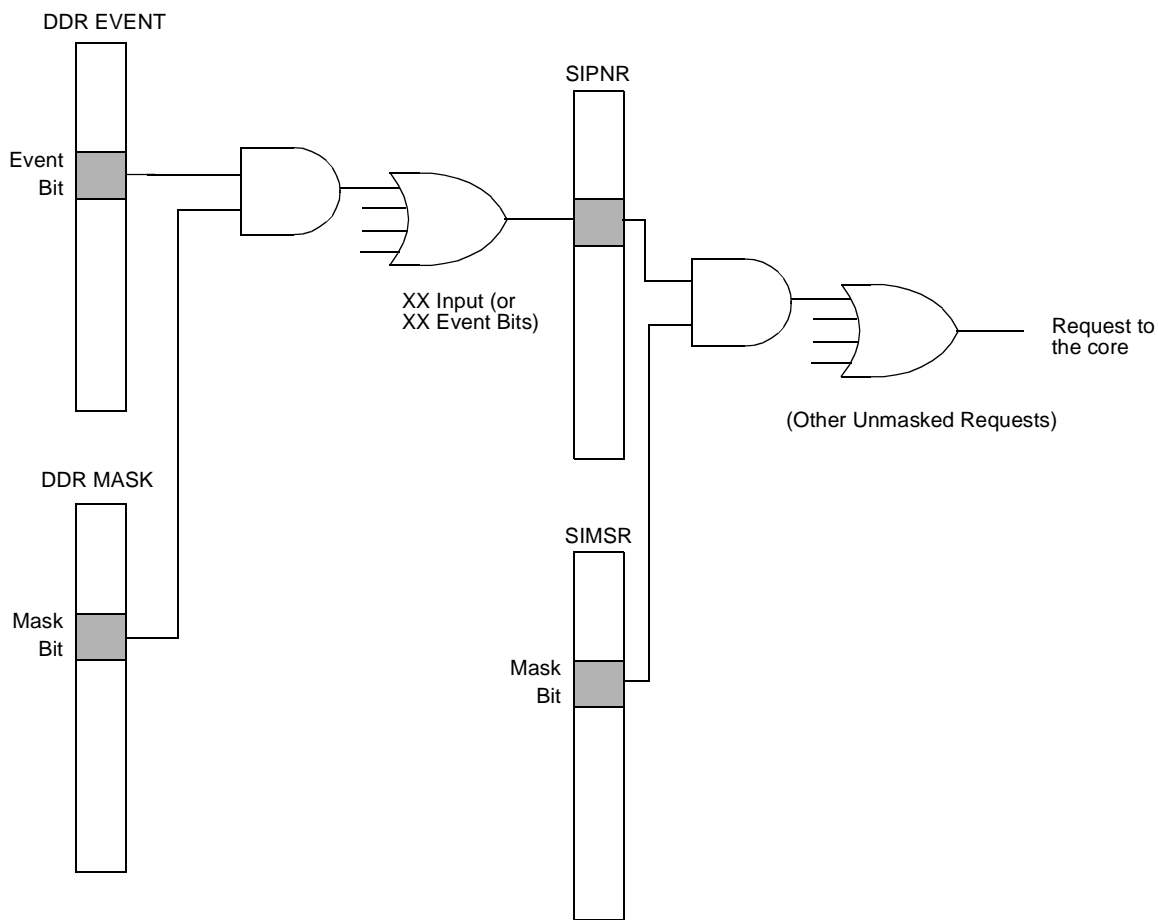


Figure 8-29. DDR Interrupt Request Masking

8.6.8 Interrupt Vector Generation and Calculation

Pending unmasked interrupts are presented to the core in order of priority according to Table 8-34. The interrupt vector that allows the core to locate the interrupt service routine is made available to the core by interrupt handler software reading SIVCR. The interrupt controller passes an interrupt vector

corresponding to the highest-priority, unmasked, pending interrupt in response to a read of SIVCR. [Table 8-5](#) lists the encodings for the seven low-order bits of the interrupt vector.

8.6.9 Machine Check Interrupts

The PIC supports the non-maskable machine check interrupts. When an error interrupt signal is received, the interrupt controller indicates the source by setting the corresponding SERSR bit. These sources are listed in [Table 8-23](#).

8.7 Message Shared Interrupts

The message shared interrupt (MSI) registers enable the system end points (PCI agents or PCI express end points) to generate interrupt requests to the local e300 CPU. Each end point can generate an interrupt and set a unique bit in one of the eight MSIR registers. Clearing the MSIR register happens immediately after the read of its content, and by then a new set operation can begin. MSIR n is considered active if it contains at least one bit set. Each active non masked MSIR n register will generate an interrupt.

8.7.1 Memory Map/Register Definition

The MSI programmable register map occupies 64 bytes of memory-mapped space. The MSI registers are 32-bit wide and must be accessed in a 32-bit read or write operation. The listed addresses are offset from the IPIC base address.

Table 8-35. Message Shared Registers Address Map

| Offset | Register | Access | Reset | Section/page |
|-----------|--|---------|-------------|------------------------------|
| 0xC0 | MSIR0—Message shared interrupt register 0 | Special | 0x0000_0000 | 8.7.2.1/8-41 |
| 0xC4 | MSIR1—Message shared interrupt register 1 | Special | 0x0000_0000 | 8.7.2.1/8-41 |
| 0xC8 | MSIR2—Message shared interrupt register 2 | Special | 0x0000_0000 | 8.7.2.1/8-41 |
| 0xCC | MSIR3—Message shared interrupt register 3 | Special | 0x0000_0000 | 8.7.2.1/8-41 |
| 0xD0 | MSIR4—Message shared interrupt register 4 | Special | 0x0000_0000 | 8.7.2.1/8-41 |
| 0xD4 | MSIR5—Message shared interrupt register 5 | Special | 0x0000_0000 | 8.7.2.1/8-41 |
| 0xD8 | MSIR6—Message shared interrupt register 6 | Special | 0x0000_0000 | 8.7.2.1/8-41 |
| 0xDC | MSIR7—Message shared interrupt register 7 | Special | 0x0000_0000 | 8.7.2.1/8-41 |
| 0xE0–0xEC | Reserved | — | — | — |
| 0xF0 | MSIMR—Message shared interrupt mask register | R/W | 0x0000_0000 | 8.7.2.2/8-41 |
| 0xF4 | MSISR—Message shared interrupt status register | R | 0x0000_0000 | 8.7.2.3/8-42 |
| 0xF8 | MSIIR—Message shared interrupt index register | W | 0x0000_0000 | 8.7.2.4/8-43 |
| 0xFC | Reserved | — | — | — |



Chapter 9

DDR Memory Controller

9.1 Introduction

The fully programmable DDR SDRAM controller supports most JEDEC standard $\times 8$, $\times 16$, or $\times 32$ DDR2 and DDR memories available. In addition, unbuffered and registered DRAM modules are supported. However, mixing different memory types or unbuffered and registered DRAM modules in the same system is not supported. Dynamic power management and auto-precharge modes simplify memory system design. A large set of special features support rapid system debug.

NOTE

In this chapter, the word ‘bank’ refers to a physical bank specified by a chip select; ‘logical bank’ refers to one of the four or eight sub-banks in each SDRAM chip. A sub-bank is specified by the 2 or 3 bits on the bank address (MBA) pins during a memory access.

Figure 9-1 is a high-level block diagram of the DDR memory controller with its associated interfaces. Section 9.5, “Functional Description,” contains detailed figures of the controller.




Figure 9-1. DDR Memory Controller Simplified Block Diagram

9.2 Features

The DDR memory controller includes these distinctive features:

- Support for DDR2 and DDR SDRAM
- 32-bit SDRAM, 16-bit SDRAM for DDR and DDR2 and
- Programmable settings for meeting all SDRAM timing parameters
- Support for the following SDRAM configurations:
 - As many as two physical banks (chip selects), each bank independently addressable
 - 64-Mbit to 4-Gbit devices depending on internal device configuration with $\times 8/\times 16/\times 32$ data ports (no direct $\times 4$ support)
 - Unbuffered and registered DRAM modules
- Open page management (dedicated entry for each logical bank)
- Automatic DRAM initialization sequence or software-controlled initialization sequence
- Automatic DRAM data initialization

- Support for up to eight posted refreshes
- Memory controller clock frequency of two times the SDRAM clock with support for sleep power management

9.2.1 Modes of Operation

The DDR memory controller supports the following modes:

- 32-byte cache line wrap mode
- Dynamic power management mode. The DDR memory controller can reduce power consumption by negating the SDRAM CKE signal when no transactions are pending to the SDRAM.
- Auto-precharge mode. Clearing `DDR_SDRAM_INTERVAL[BSTOPRE]` causes the memory controller to issue an auto-precharge command with every read or write transaction. Auto-precharge mode can be enabled for separate chip selects by setting `CSn_CONFIG[AP_n_EN]`.

9.3 External Signal Descriptions

This section provides descriptions of the DDR memory controller's external signals. It describes each signal's behavior when the signal is asserted or negated and when the signal is an input or an output.

9.3.1 Signals Overview

Memory controller signals are grouped as follows:

- Memory interface signals
- Clock signals
- Debug signals

Table 9-1 shows how DDR memory controller external signals are grouped. The device hardware specification has a pinout diagram showing pin numbers. It also lists all electrical and mechanical specifications.

Table 9-1. DDR Memory Interface Signal Summary

| Name | Function/Description | Reset | Pins | I/O |
|------------------------------|-----------------------|-----------|------|-----|
| MDQ[0:31] | Data bus | All zeros | 32 | I/O |
| MDQS[0:3] | Data strobes | All zeros | 4 | I/O |
| $\overline{\text{MCAS}}$ | Column address strobe | One | 1 | O |
| MA[14:0] | Address bus | All zeros | 15 | O |
| MBA[2:0] | Logical bank address | All zeros | 3 | O |
| $\overline{\text{MCS}}[0:1]$ | Chip selects | All ones | 2 | O |
| $\overline{\text{MWE}}$ | Write enable | One | 1 | O |
| $\overline{\text{MRAS}}$ | Row address strobe | One | 1 | O |

Table 9-1. DDR Memory Interface Signal Summary (continued)

| Name | Function/Description | Reset | Pins | I/O |
|----------------------------|---|-----------|------|-----|
| MDM[0:3] | Data mask | All zeros | 4 | O |
| MCK[0] | DRAM clock outputs | Zero | 1 | O |
| $\overline{\text{MCK}}[0]$ | DRAM clock outputs (complement) | One | 1 | O |
| MCKE | DRAM clock enable | Zero | 1 | O |
| MODT[0:1] | DRAM on-die termination external control. | All zeros | 2 | O |
| MDVAL | Memory debug data valid | Zero | 1 | O |
| MSRCID[0:4] | Memory debug source ID | All zeros | 5 | O |

Table 9-2 shows the memory address signal mappings.

Table 9-2. Memory Address Signal Mappings

| Signal Name (Outputs) | | JEDEC DDR DIMM Signals (Inputs) |
|-----------------------|------|--|
| msb | MA14 | A14 |
| | MA13 | A13 |
| | MA12 | A12 |
| | MA11 | A11 |
| | MA10 | A10 (AP for DDR) ¹ |
| | MA9 | A9 |
| | MA8 | A8 (alternate AP for DDR) ² |
| | MA7 | A7 |
| | MA6 | A6 |
| | MA5 | A5 |
| | MA4 | A4 |
| | MA3 | A3 |
| | MA2 | A2 |
| | MA1 | A1 |
| lsb | MA0 | A0 |
| msb | MBA2 | MBA2 |
| | MBA1 | MBA1 |
| lsb | MBA0 | MBA0 |

¹ Auto-precharge for DDR signaled on A10 when DDR_SDRAM_CFG[PCHB8] = 0.

² Auto-precharge for DDR signaled on A8 when DDR_SDRAM_CFG[PCHB8] = 1.

9.3.2 Detailed Signal Descriptions

The following sections describe the DDR SDRAM controller input and output signals, the meaning of their different states, and relative timing information for assertion and negation.

9.3.2.1 Memory Interface Signals

Table 9-3 describes the DDR controller memory interface signals.

Table 9-3. Memory Interface Signals—Detailed Signal Descriptions

| Signal | I/O | Description | |
|---------------|---------------|---|---|
| MDQ[0:31] | I/O | Data bus. Both input and output signals on the DDR memory controller. | |
| | O | As outputs for the bidirectional data bus, these signals operate as described below. | |
| | | State Meaning | Asserted/Negated—Represent the value of data being driven by the DDR memory controller. |
| | | Timing | Assertion/Negation—Driven coincident with corresponding data strobes (MDQS) signal. High impedance—No READ or WRITE command is in progress; data is not being driven by the memory controller or the DRAM. |
| | I | As inputs for the bidirectional data bus, these signals operate as described below. | |
| | | State Meaning | Asserted/Negated—Represents the state of data being driven by the external DDR SDRAMs. |
| Timing | | Assertion/Negation—The DDR SDRAM drives data during a READ transaction. High impedance—No READ or WRITE command in progress; data is not being driven by the memory controller or the DRAM. | |
| MDQS[0:3] | I/O | Data strobes. Inputs with read data, outputs with write data. | |
| | | O | As outputs, the data strobes are driven by the DDR memory controller during a write transaction. The memory controller always drives these signals low unless a read has been issued and incoming data strobes are expected. This keeps the data strobes from floating high when there are no transactions on the DRAM interface. |
| | | | State Meaning |
| | Timing | Assertion/Negation—If a WRITE command is registered at clock edge n , data strobes at the DRAM assert centered in the data eye on clock edge $n + 1$. See the JEDEC DDR SDRAM specification for more information. | |
| | I | As inputs, the data strobes are driven by the external DDR SDRAMs during a read transaction. The data strobes are used by the memory controller to synchronize data latching. | |
| | | State Meaning | Asserted/Negated—Driven high when positive capture data is received and driven low when negative capture data is received. Centered in the data eye for writes; coincident with the data eye for reads. Treated as a clock. Data is valid when signals toggle. See Table 9-27 for byte lane assignments. |
| Timing | | Assertion/Negation—If a READ command is registered at clock edge n , and the latency is programmed in TIMING_CFG_1[CASLAT] to be m clocks, data strobes at the DRAM assert coincident with the data on clock edge $n + m$. See the JEDEC DDR SDRAM specification for more information. | |

Table 9-3. Memory Interface Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|-------------------|-----|---|--|
| MA[14:0] | O | Address bus. Memory controller outputs for the address to the DRAM. MA[14:0] carry 15 of the address bits for the DDR memory interface corresponding to the row and column address bits. MA0 is the lsb of the address output from the memory controller. | |
| | | State Meaning | Asserted/Negated—Represents the address driven by the DDR memory controller. Contains different portions of the address depending on the memory size and the DRAM command being issued by the memory controller. See Table 9-30 and Table 9-33 for a complete description of the mapping of these signals. |
| | | Timing | Assertion/Negation—The address lines are only driven when the controller has a command scheduled to issue on the address/CMD bus; otherwise they will be at high-Z. It is valid when a transaction is driven to DRAM (when \overline{MCSn} is active). High impedance—When the memory controller is disabled |
| MBA[2:0] | O | Logical bank address. Outputs that drive the logical (or internal) bank address pins of the SDRAM. Each SDRAM supports four or eight addressable logical sub-banks. Bit zero of the memory controller's output bank address must be connected to bit zero of the SDRAM's input bank address. MBA0, the least-significant bit of the three bank address signals, is asserted during the mode register set command to specify the extended mode register. | |
| | | State Meaning | Asserted/Negated—Selects the DDR SDRAM logical (or internal) bank to be activated during the row address phase and selects the SDRAM internal bank for the read or write operation during the column address phase of the memory access. Table 9-30 and Table 9-33 describes the mapping of these signals in all cases. |
| | | Timing | Assertion/Negation—Same timing as MA_n High impedance—Same timing as MA_n |
| \overline{MCAS} | O | Column address strobe. Active-low SDRAM address multiplexing signal. \overline{MCAS} is asserted for read or write transactions and for mode register set, refresh, and precharge commands. | |
| | | State Meaning | Asserted—Indicates that a valid SDRAM column address is on the address bus for read and write transactions. See Table 9-35 for more information on the states required on \overline{MCAS} for various other SDRAM commands. Negated—The column address is not guaranteed to be valid. |
| | | Timing | Assertion/Negation—Assertion and negation timing is directed by the values described in Section 9.4.1.4, "DDR SDRAM Timing Configuration 0 (TIMING_CFG_0)," Section 9.4.1.5, "DDR SDRAM Timing Configuration 1 (TIMING_CFG_1)," Section 9.4.1.6, "DDR SDRAM Timing Configuration 2 (TIMING_CFG_2)," and Section 9.4.1.3, "DDR SDRAM Timing Configuration 3 (TIMING_CFG_3)." High impedance— \overline{MCAS} is always driven unless the memory controller is disabled. |
| \overline{MRAS} | O | Row address strobe. Active-low SDRAM address multiplexing signal. Asserted for activate commands. In addition; used for mode register set commands and refresh commands. | |
| | | State Meaning | Asserted—Indicates that a valid SDRAM row address is on the address bus for read and write transactions. See Table 9-35 for more information on the states required on \overline{MRAS} for various other SDRAM commands. Negated—The row address is not guaranteed to be valid. |
| | | Timing | Assertion/Negation—Assertion and negation timing is directed by the values described in Section 9.4.1.4, "DDR SDRAM Timing Configuration 0 (TIMING_CFG_0)," Section 9.4.1.5, "DDR SDRAM Timing Configuration 1 (TIMING_CFG_1)," Section 9.4.1.6, "DDR SDRAM Timing Configuration 2 (TIMING_CFG_2)," and Section 9.4.1.3, "DDR SDRAM Timing Configuration 3 (TIMING_CFG_3)." High impedance— \overline{MRAS} is always driven unless the memory controller is disabled. |

Table 9-3. Memory Interface Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|-------------------------------|-----|---|
| $\overline{\text{MCS}}[0:1]$ | O | Chip selects. Two chip selects supported by the memory controller. |
| | | State Meaning Asserted—Selects a physical SDRAM bank to perform a memory operation as described in Section 9.4.1.1, “Chip Select Memory Bounds (CS_n_BND_S),” and Section 9.4.1.2, “Chip Select Configuration (CS_n_CONFIG).” The DDR controller asserts one of the $\overline{\text{MCS}}[0:1]$ signals to begin a memory cycle. Negated—Indicates no SDRAM action during the current cycle. |
| | | Timing Assertion/Negation—Asserted to signal any new transaction to the SDRAM. The transaction must adhere to the timing constraints set in TIMING_CFG_0–TIMING_CFG_3. High impedance—Always driven unless the memory controller is disabled. |
| $\overline{\text{MWE}}$ | O | Write enable. Asserted when a write transaction is issued to the SDRAM. This is also used for mode registers set commands and precharge commands. |
| | | State Meaning Asserted—Indicates a memory write operation. See Table 9-35 for more information on the states required on $\overline{\text{MWE}}$ for various other SDRAM commands. Negated—Indicates a memory read operation. |
| | | Timing Assertion/Negation—Similar timing as $\overline{\text{MRAS}}$ and $\overline{\text{MCAS}}$. Used for write commands. High impedance— $\overline{\text{MWE}}$ is always driven unless the memory controller is disabled. |
| $\overline{\text{MDM}}[0:3]$ | O | DDR SDRAM data output mask. Masks unwanted bytes of data transferred during a write. They are needed to support sub-burst-size transactions (such as single-byte writes) on SDRAM where all I/O occurs in multi-byte bursts. $\overline{\text{MDM}}_0$ corresponds to the most significant byte (MSB). Table 9-27 shows byte lane encodings. |
| | | State Meaning Asserted—Prevents writing to DDR SDRAM. Asserted when data is written to DRAM if the corresponding byte(s) should be masked for the write. Note that the $\overline{\text{MDM}}_n$ signals are active-high for the DDR controller. $\overline{\text{MDM}}_n$ is part of the DDR command encoding. Negated—Allows the corresponding byte to be read from or written to the SDRAM. |
| | | Timing Assertion/Negation—Same timing as $\overline{\text{MDQx}}$ as outputs. High-impedance—Always driven unless the memory controller is disabled. |
| $\overline{\text{MODT}}[0:1]$ | O | On-Die termination. Memory controller outputs for the ODT to the DRAM. $\overline{\text{MODT}}[0:1]$ represents the on-die termination for the associated data, data masks, and data strobes. |
| | | State Meaning Asserted/Negated—Represents the ODT driven by the DDR memory controller. |
| | | Timing Assertion/Negation—Driven in accordance with JEDEC DRAM specifications for on-die termination timings. It is configured through the CS _n _CONFIG[ODT_RD_CFG] and CS _n _CONFIG[ODT_WR_CFG] fields. High impedance—Always driven. |

9.3.2.2 Clock Interface Signals

[Table 9-4](#) contains the detailed descriptions of the clock signals of the DDR controller.

Table 9-4. Clock Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|--|-----|--|
| $\overline{\text{MCK}}[0]$, $\overline{\text{MCK}}[0]$ | O | DRAM clock output and its complement. See Section 9.5.4.1, “Clock Distribution.” |
| | | State Meaning Asserted/Negated—The JEDEC DDR SDRAM specifications require true and complement clocks. A clock edge is seen by the SDRAM when the true and complement cross. |
| | | Timing Assertion/Negation—Timing is controlled by the DDR_CLK_CNTL register at offset 0x130. |

Table 9-4. Clock Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description | | |
|--|--|--|--|--|
| MCKE | O | Clock enable. Output signals used as the clock enables to the SDRAM. MCKE can be negated to stop clocking the DDR SDRAM. The MCKE signals should be connected to the same rank of memory as the corresponding MCS and MODT signals. For example, MCKE[0] should be connected to the same rank of memory as MCS[0] and MODT[0]. | | |
| | | <table border="1"> <tr> <td>State Meaning</td> <td>Asserted—Clocking to the SDRAM is enabled. Negated—Clocking to the SDRAM is disabled and the SDRAM should ignore signal transitions on MCK or MCK. MCK/MCK are don't cares while MCKE is negated.</td> </tr> </table> | State Meaning | Asserted—Clocking to the SDRAM is enabled. Negated—Clocking to the SDRAM is disabled and the SDRAM should ignore signal transitions on MCK or MCK. MCK/MCK are don't cares while MCKE is negated. |
| | | State Meaning | Asserted—Clocking to the SDRAM is enabled. Negated—Clocking to the SDRAM is disabled and the SDRAM should ignore signal transitions on MCK or MCK. MCK/MCK are don't cares while MCKE is negated. | |
| <table border="1"> <tr> <td>Timing</td> <td>Assertion/Negation—Asserted when DDR_SDRAM_CFG[MEM_EN] is set. Can be negated when entering dynamic power management or self refresh. Are asserted again when exiting dynamic power management or self refresh. High impedance—Always driven.</td> </tr> </table> | Timing | Assertion/Negation—Asserted when DDR_SDRAM_CFG[MEM_EN] is set. Can be negated when entering dynamic power management or self refresh. Are asserted again when exiting dynamic power management or self refresh. High impedance—Always driven. | | |
| Timing | Assertion/Negation—Asserted when DDR_SDRAM_CFG[MEM_EN] is set. Can be negated when entering dynamic power management or self refresh. Are asserted again when exiting dynamic power management or self refresh. High impedance—Always driven. | | | |

9.3.2.3 Debug Signals

The debug signals MSRCID[0:4] and MDVAL have no function in normal DDR controller operation. A detailed description of these signals can be found in [Section 5.3.2.7, “Debug Configuration.”](#)

9.4 Memory Map/Register Definition

[Table 9-5](#) shows the register memory map for the DDR memory controller.

In this table and in the register figures and field descriptions, the following access definitions apply:

- Reserved fields are always ignored for the purposes of determining access type.
- R/W, R, and W (read/write, read only, and write only) indicate that all the non-reserved fields in a register have the same access type.
- w1c indicates that all of the non-reserved fields in a register are cleared by writing ones to them.
- Mixed indicates a combination of access types.
- Special is used when no other category applies. In this case the register figure and field description table should be read carefully.

Table 9-5. DDR Memory Controller Memory Map

| Offset | Register | Access | Reset | Section/Page |
|--|---|--------|-------------|------------------------------|
| DDR Memory Controller—Block Base Address 0x0_2000 | | | | |
| 0x000 | CS0_BNDS—Chip select 0 memory bounds | R/W | All zeros | 9.4.1.1/9-9 |
| 0x008 | CS1_BNDS—Chip select 1 memory bounds | R/W | All zeros | 9.4.1.1/9-9 |
| 0x080 | CS0_CONFIG—Chip select 0 configuration | R/W | All zeros | 9.4.1.2/9-10 |
| 0x084 | CS1_CONFIG—Chip select 1 configuration | R/W | All zeros | 9.4.1.2/9-10 |
| 0x100 | TIMING_CFG_3—DDR SDRAM timing configuration 3 | R/W | All zeros | 9.4.1.3/9-12 |
| 0x104 | TIMING_CFG_0—DDR SDRAM timing configuration 0 | R/W | 0x0011_0105 | 9.4.1.4/9-12 |
| 0x108 | TIMING_CFG_1—DDR SDRAM timing configuration 1 | R/W | All zeros | 9.4.1.5/9-14 |
| 0x10C | TIMING_CFG_2—DDR SDRAM timing configuration 2 | R/W | All zeros | 9.4.1.6/9-16 |

Table 9-5. DDR Memory Controller Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|-----------------|---|--------|---|---------------|
| 0x110 | DDR_SDRAM_CFG—DDR SDRAM control configuration | R/W | 0x0200_0000 | 9.4.1.7/9-18 |
| 0x114 | DDR_SDRAM_CFG_2—DDR SDRAM control configuration 2 | R/W | All zeros | 9.4.1.8/9-21 |
| 0x118 | DDR_SDRAM_MODE—DDR SDRAM mode configuration | R/W | All zeros | 9.4.1.9/9-22 |
| 0x11C | DDR_SDRAM_MODE_2—DDR SDRAM mode configuration 2 | R/W | All zeros | 9.4.1.10/9-23 |
| 0x120 | DDR_SDRAM_MD_CNTL—DDR SDRAM mode control | R/W | All zeros | 9.4.1.11/9-24 |
| 0x124 | DDR_SDRAM_INTERVAL—DDR SDRAM interval configuration | R/W | All zeros | 9.4.1.12/9-27 |
| 0x128 | DDR_DATA_INIT—DDR SDRAM data initialization | R/W | All zeros | 9.4.1.13/9-27 |
| 0x130 | DDR_SDRAM_CLK_CNTL—DDR SDRAM clock control | R/W | 0x0200_0000 | 9.4.1.14/9-28 |
| 0x140– 0x144 | Reserved | — | — | — |
| 0x148 | DDR_INIT_ADDR—DDR training initialization address | R/W | All zeros | 9.4.1.15/9-28 |
| 0x150– 0xBF4 | Reserved | — | — | — |
| 0xBF8 | DDR_IP_REV1—DDR IP block revision 1 | R | 0xn _{nnn} _n _{nnn} ¹ | 9.4.1.16/9-29 |
| 0xBFC | DDR_IP_REV2—DDR IP block revision 2 | R | 0x00n _n _00n _n ¹ | 9.4.1.17/9-29 |
| 0xE40 | ERR_DETECT—Memory error detect | w1c | All zeros | 9.4.1.18/9-30 |
| 0xE44 | ERR_DISABLE—Memory error disable | R/W | All zeros | 9.4.1.19/9-30 |
| 0xE48 | ERR_INT_EN—Memory error interrupt enable | R/W | All zeros | 9.4.1.20/9-31 |

¹ Implementation-dependent reset values are listed in specified section/page.

9.4.1 Register Descriptions

This section describes the DDR memory controller registers. Shading indicates reserved fields that should not be written.

9.4.1.1 Chip Select Memory Bounds (CS_n_BNDS)

The chip select bounds registers (CS_n_BNDS) define the starting and ending address of the memory space that corresponds to the individual chip selects. Note that the size specified in CS_n_BNDS should equal the size of physical DRAM. Also, note that EA_n must be greater than or equal to SA_n.

If chip select interleaving is enabled, all fields in the lower interleaved chip select are used, and the other chip selects' bounds registers are unused. For example, if chip selects 0 and 1 are interleaved, all fields in CS₀_BNDS are used, and all fields in CS₁_BNDS are unused.

CS_n_BNDS are shown in [Figure 9-2](#).



Figure 9-2. Chip Select Bounds Registers (CS_n_BNDS)

[Table 9-6](#) describes the CS_n_BNDS register fields.

Table 9-6. CS_n_BNDS Field Descriptions

| Bits | Name | Description |
|-------|-----------------|---|
| 0–7 | — | Reserved |
| 8–15 | SA _n | Starting address for chip select (bank) <i>n</i> . This value is compared against the 8 msbs of the 32-bit address. |
| 16–23 | — | Reserved |
| 24–31 | EA _n | Ending address for chip select (bank) <i>n</i> . This value is compared against the 8 msbs of the 32-bit address. |

9.4.1.2 Chip Select Configuration (CS_n_CONFIG)

The chip select configuration (CS_n_CONFIG) registers shown in [Figure 9-3](#) enable the DDR chip selects and set the number of row and column bits used for each chip select. These registers should be loaded with the correct number of row and column bits for each SDRAM. Because CS_n_CONFIG[ROW_BITS_CS_n, COL_BITS_CS_n] establish address multiplexing, the user should take great care to set these values correctly.

If chip select interleaving is enabled, then all fields in the lower interleaved chip select are used, and the other registers' fields are unused, with the exception of the ODT_RD_CFG and ODT_WR_CFG fields. For example, if chip selects 0 and 1 are interleaved, all fields in CS₀_CONFIG are used, but only the ODT_RD_CFG and ODT_WR_CFG fields in CS₁_CONFIG are used.




Figure 9-3. Chip Select Configuration Register (CS_n_CONFIG)

Table 9-7 describes the CS_n_CONFIG register fields.

Table 9-7. CS_n_CONFIG Field Descriptions

| Bits | Name | Description |
|-------|--------------------------|---|
| 0 | CS _n _EN | Chip select <i>n</i> enable 0 Chip select <i>n</i> is not active 1 Chip select <i>n</i> is active and assumes the state set in CS _n _BNDS. |
| 1–7 | — | Reserved |
| 8 | AP _n _EN | Chip select <i>n</i> auto-precharge enable 0 Chip select <i>n</i> is only auto-precharged if global auto-precharge mode is enabled (DDR_SDRAM_INTERVAL[BSTOPRE] = 0). 1 Chip select <i>n</i> always issues an auto-precharge for read and write transactions. |
| 9–11 | ODT_RD_CFG | ODT for reads configuration. Note that CAS latency plus additive latency must be at least 3 cycles for ODT_RD_CFG to be enabled. ODT should only be used with DDR2 or memories. 000 Never assert ODT for reads 001 Assert ODT only during reads to CS _n 010 Assert ODT only during reads to other chip selects 011 Reserved 100 Assert ODT for all reads 101–111 Reserved |
| 12 | — | Reserved |
| 13–15 | ODT_WR_CFG | ODT for writes configuration. Note that write latency plus additive latency must be at least 3 cycles for ODT_WR_CFG to be enabled. ODT should only be used with DDR2 or memories. 000 Never assert ODT for writes 001 Assert ODT only during writes to CS _n 010 Assert ODT only during writes to other chip selects 011 Reserved 100 Assert ODT for all writes 101–111 Reserved |
| 16–17 | BA_BITS_CS _n | Number of bank bits for SDRAM on chip select <i>n</i> . These bits correspond to the sub-bank bits driven on MBA _n in Table 9-33 . 00 2 logical bank bits 01 3 logical bank bits 10–11 Reserved |
| 18–20 | — | Reserved |
| 21–23 | ROW_BITS_CS _n | Number of row bits for SDRAM on chip select <i>n</i> . See Table 9-33 for details. 000 12 row bits 001 13 row bits 010 14 row bits 011 15 row bits 000–111 Reserved |
| 24–28 | — | Reserved |
| 29–31 | COL_BITS_CS _n | Number of column bits for SDRAM on chip select <i>n</i> . For DDR, the decoding is as follows: 000 8 column bits 001 9 column bits 010 10 column bits 011 11 column bits 100–111 Reserved |

9.4.1.3 DDR SDRAM Timing Configuration 3 (TIMING_CFG_3)

DDR SDRAM timing configuration register 3, shown in Figure 9-4, sets the extended refresh recovery time, which is combined with TIMING_CFG_1[REFREC] to determine the full refresh recovery time.




Figure 9-4. DDR SDRAM Timing Configuration 3 (TIMING_CFG_3)

Table 9-8 describes TIMING_CFG_3 fields.

Table 9-8. TIMING_CFG_3 Field Descriptions

| Bits | Name | Description |
|-------|------------|--|
| 0–12 | — | Reserved |
| 13–15 | EXT_REFREC | Extended refresh recovery time (t_{RFC}). Controls the number of clock cycles from a refresh command until an activate command is allowed. This field is concatenated with TIMING_CFG_1[REFREC] to obtain a 7-bit value for the total refresh recovery. Note that hardware adds an additional 8 clock cycles to the final, 7-bit value of the refresh recovery. $t_{RFC} = \{EXT_REFREC \parallel REFREC\} + 8$, such that t_{RFC} is calculated as follows: 000 0 clocks 001 16 clocks 010 32 clocks 011 48 clocks 100 64 clocks 101 80 clocks 110 96 clocks 111 112 clocks |
| 16–31 | — | Reserved |

9.4.1.4 DDR SDRAM Timing Configuration 0 (TIMING_CFG_0)

DDR SDRAM timing configuration register 0, shown in Figure 9-5, sets the number of clock cycles between various SDRAM control commands.




Figure 9-5. DDR SDRAM Timing Configuration 0 (TIMING_CFG_0)

Table 9-10. TIMING_CFG_1 Field Descriptions (continued)

| Bits | Name | Description |
|-------|----------|---|
| 1–3 | PRETOACT | Precharge-to-activate interval (t_{RP}). Determines the number of clock cycles from a precharge command until an activate or refresh command is allowed. 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks |
| 4–7 | ACTTOPRE | Activate to precharge interval (t_{RAS}). Determines the number of clock cycles from an activate command until a precharge command is allowed. 0000 16 clocks 0101 5 clocks 0001 17 clocks 0110 6 clocks 0010 18 clocks 0111 7 clocks 0011 19 clocks ... 0100 4 clocks 1111 15 clocks |
| 8 | — | Reserved |
| 9–11 | ACTTORW | Activate to read/write interval for SDRAM (t_{RCD}). Controls the number of clock cycles from an activate command until a read or write command is allowed. 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks |
| 12–15 | CASLAT | \overline{MCAS} latency from READ command. Number of clock cycles between registration of a READ command by the SDRAM and the availability of the first output data. If a READ command is registered at clock edge n and the latency is m clocks, data is available nominally coincident with clock edge $n + m$. This value must be programmed at initialization as described in Section 9.4.1.8, “DDR SDRAM Control Configuration 2 (DDR_SDRAM_CFG_2).” 0000 Reserved 1000 4.5 clocks 0001 1 clock 1001 5 clocks 0010 1.5 clocks 1010 5.5 clocks 0011 2 clocks 1011 6 clocks 0100 2.5 clocks 1100 6.5 clocks 0101 3 clocks 1101 7 clocks 0110 3.5 clocks 1110 7.5 clocks 0111 4 clocks 1111 8 clocks |

Table 9-10. TIMING_CFG_1 Field Descriptions (continued)

| Bits | Name | Description |
|-------|----------|---|
| 16–19 | REFREC | Refresh recovery time (t_{RFC}). Controls the number of clock cycles from a refresh command until an activate command is allowed. This field is concatenated with TIMING_CFG_3[EXTREFREC] to obtain a 7-bit value for the total refresh recovery. Note that hardware adds an additional 8 clock cycles to the final, 7-bit value of the refresh recovery, such that t_{RFC} is calculated as follows: $t_{RFC} = \{EXT_REFREC REFREC\} + 8$. 0000 8 clocks 0011 11 clocks 0001 9 clocks ... 0010 10 clocks 1111 23 clocks |
| 20 | — | Reserved |
| 21–23 | WRREC | Last data to precharge minimum interval (t_{WR}). Determines the number of clock cycles from the last data associated with a write command until a precharge command is allowed. 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks |
| 24 | — | Reserved |
| 25–27 | ACTTOACT | Activate-to-activate interval (t_{RRD}). Number of clock cycles from an activate command until another activate command is allowed for a different logical bank in the same physical bank (chip select). 000 Reserved 100 4 clocks 001 1 clock 101 5 clocks 010 2 clocks 110 6 clocks 011 3 clocks 111 7 clocks |
| 28 | — | Reserved |
| 29–31 | WRTORD | Last write data pair to read command issue interval (t_{WTR}). Number of clock cycles between the last write data pair and the subsequent read command to the same physical bank. 000 Reserved 100 4 clocks 001 1 clock 101 5 clocks 010 2 clocks 110 6 clocks 011 3 clocks 111 7 clocks |

9.4.1.6 DDR SDRAM Timing Configuration 2 (TIMING_CFG_2)

DDR SDRAM timing configuration 2, shown in [Figure 9-7](#), sets the clock delay to data for writes.




Figure 9-7. DDR SDRAM Timing Configuration 2 Register (TIMING_CFG_2)

Table 9-12 describes the DDR_SDRAM_CFG fields.

Table 9-12. DDR_SDRAM_CFG Field Descriptions

| Bits | Name | Description |
|------|------------|---|
| 0 | MEM_EN | DDR SDRAM interface logic enable. 0 SDRAM interface logic is disabled. 1 SDRAM interface logic is enabled. Must not be set until all other memory configuration parameters have been appropriately configured by initialization code. |
| 1 | SREN | Self refresh enable (during sleep). 0 SDRAM self refresh is disabled during sleep. Whenever self-refresh is disabled, the system is responsible for preserving the integrity of SDRAM during sleep. 1 SDRAM self refresh is enabled during sleep. |
| 2 | — | Reserved. Must be cleared. |
| 3 | RD_EN | Registered DRAM module enable. Specifies the type of DRAM module used in the system. 0 Indicates unbuffered DRAM modules. 1 Indicates registered DRAM modules. Note: RD_EN and 2T_EN must not both be set at the same time. |
| 4 | — | Reserved |
| 5–7 | SDRAM_TYPE | Type of SDRAM device to be used. This field is used when issuing the automatic hardware initialization sequence to DRAM through Mode Register Set and Extended Mode Register Set commands. Default value is 010 designating DDR1 SDRAM. 000–001 Reserved 010 DDR1 SDRAM 011 DDR2 SDRAM 100 Reserved 101 Reserved 110 Reserved 111 Reserved |
| 8–9 | — | Reserved |
| 10 | DYN_PWR | Dynamic power management mode 0 Dynamic power management mode is disabled. 1 Dynamic power management mode is enabled. If there is no ongoing memory activity, the SDRAM CKE signal is negated. |
| 11 | — | Reserved |
| 12 | 32_BE | 32-bit bus enable. 0 Reserved 1 32-bit bus is used. Software MUST set this bit. |
| 13 | 8_BE | 8-beat burst enable. 0 4-beat bursts are used on the DRAM interface. 1 8-beat bursts are used on the DRAM interface. Note: DDR1 (SDRAM_TYPE = 010) must use 8-beat bursts when using 32-bit bus mode (32_BE = 1) and 4-beat bursts when using 64-bit bus mode; DDR2 (SDRAM_TYPE = 011) must use 4-beat bursts, even when using 32-bit bus mode; |
| 14 | NCAP | Non-concurrent auto-precharge. Some older DDR DRAMs do not support concurrent auto precharge. If one of these devices is used, then this bit needs to be set if auto precharge is used. 0 DRAMs in system support concurrent auto-precharge. 1 DRAMs in system do not support concurrent auto-precharge. |
| 15 | — | Reserved |

Table 9-12. DDR_SDRAM_CFG Field Descriptions (continued)

| Bits | Name | Description |
|-------|--------------|---|
| 16 | 2T_EN | Enable 2T timing. 0 1T timing is enabled. The DRAM command/address is held for only 1 cycle on the DRAM bus. 1 2T timing is enabled. The DRAM command/address are held for 2 full cycles on the DRAM bus for every DRAM transaction. However, the chip select is only held for the second cycle. Note: RD_EN and 2T_EN must not both be set at the same time. |
| 17–23 | BA_INTLV_CTL | Bank (chip select) interleaving control. Set this field only if you wish to use bank interleaving. (All unlisted field values are reserved.) 0000000No external memory banks are interleaved 1000000External memory banks 0 and 1 are interleaved |
| 24–26 | — | Reserved |
| 27 | PCHB8 | Precharge bit 8 enable. 0 MA[10] is used to indicate the auto-precharge and precharge all commands. 1 MA[8] is used to indicate the auto-precharge and precharge all commands. If x32_EN is cleared, then PCHB8 should be cleared as well. |
| 28 | HSE | Global half-strength override Sets I/O driver impedance to half strength. This impedance is used by the address/command, data, and clock impedance values, but only if automatic hardware calibration is disabled and the corresponding group's software override is disabled in the DDR control driver register(s) described in Section 5.3.2.8, “DDR Control Driver Register (DDRCDR).” This bit should be cleared if using automatic hardware calibration. 0 I/O driver impedance is configured to full strength. 1 I/O driver impedance is configured to half strength. |
| 29 | — | Reserved |
| 30 | MEM_HALT | DDR memory controller halt. When this bit is set, the memory controller does not accept any new data read/write transactions to DDR SDRAM until the bit is cleared again. This can be used when bypassing initialization and forcing MODE REGISTER SET commands through software. 0 DDR controller accepts new transactions. 1 DDR controller finishes any remaining transactions, and then it remains halted until this bit is cleared by software. |
| 31 | BI | Bypass initialization 0 DDR controller cycles through initialization routine based on SDRAM_TYPE 1 Initialization routine is bypassed. Software is responsible for initializing memory through DDR_SDRAM_MODE2 register. If software is initializing memory, then the MEM_HALT bit can be set to prevent the DDR controller from issuing transactions during the initialization sequence. Note that the DDR controller does not issue a DLL reset to the DRAMs when bypassing the initialization routine, regardless of the value of DDR_SDRAM_CFG[DLL_RST_DIS]. If a DLL reset is required, then the controller should be forced to enter and exit self-refresh after the controller is enabled. |

9.4.1.8 DDR SDRAM Control Configuration 2 (DDR_SDRAM_CFG_2)

The DDR SDRAM control configuration register 2, shown in [Figure 9-9](#), provides more control configuration for the DDR controller.

Offset 0x114

Access: Read/Write




Figure 9-9. DDR SDRAM Control Configuration Register 2 (DDR_SDRAM_CFG_2)

[Table 9-13](#) describes the DDR_SDRAM_CFG_2 fields.

Table 9-13. DDR_SDRAM_CFG_2 Field Descriptions

| Bits | Name | Description |
|------|-------------|---|
| 0 | FRC_SR | Force self-refresh 0 DDR controller operates in normal mode. 1 DDR controller enters self-refresh mode. |
| 1 | SR_IE | Self-refresh interrupt enable. The DDR controller can be placed into self refresh mode if DDR_SR_REQ (IRQ1) is asserted. This is considered a 'panic interrupt' by the DDR controller, and it will enter self refresh as soon as possible. DDR_SDRAM_CFG[SREN] must also be set if the panic interrupt is used. 0 DDR controller will not enter self-refresh mode if panic interrupt is asserted. 1 DDR controller will enter self-refresh mode if panic interrupt is asserted. |
| 2 | DLL_RST_DIS | DLL reset disable The DDR controller typically issues a DLL reset to the DRAMs when exiting self refresh. However, this function may be disabled by setting this bit during initialization. 0 DDR controller issues a DLL reset to the DRAMs when exiting self refresh. 1 DDR controller does not issue a DLL reset to the DRAMs when exiting self refresh. |
| 3 | — | Reserved |
| 4-5 | DQS_CFG | DQS configuration 00 Only true DQS signals are used. 01 Reserved 10 Reserved 11 Reserved |
| 6-8 | — | Reserved |

Table 9-13. DDR_SDRAM_CFG_2 Field Descriptions (continued)

| Bits | Name | Description |
|-------|---------|---|
| 9–10 | ODT_CFG | ODT configuration This field defines how ODT is driven to the on-chip IOs. See Section 5.3.2.8, “DDR Control Driver Register (DDRCDR),” which defines the termination value that is used. (DDR2-specific, must be cleared for DDR1) 00 Never assert ODT to internal IOs 01 Assert ODT to internal IOs only during writes to DRAM 10 Assert ODT to internal IOs only during reads to DRAM 11 Always keep ODT asserted to internal IOs |
| 11–15 | — | Reserved. |
| 16–19 | NUM_PR | Number of posted refreshes This determines how many posted refreshes, if any, can be issued at one time. Note that if posted refreshes are used, then this field, along with DDR_SDRAM_INTERVAL[REFINT], must be programmed such that the maximum t_{ras} specification cannot be violated. For example, some DDR1 SDRAMs are not able to use more than 3 posted refreshes because the required refresh interval could then exceed the maximum constraint for t_{ras} . 0000 Reserved 0001 1 refresh is issued at a time 0010 2 refreshes is issued at a time 0011 3 refreshes is issued at a time ... 1000 8 refreshes is issued at a time 1001–1111 Reserved |
| 20–24 | — | Reserved |
| 20–26 | — | Reserved |
| 27 | D_INIT | DRAM data initialization This bit is set by software, and it is cleared by hardware. If software sets this bit before the memory controller is enabled, the controller automatically initializes DRAM after it is enabled. This bit is automatically cleared by hardware once the initialization is completed. This data initialization bit should only be set when the controller is idle. 0 There is not data initialization in progress, and no data initialization is scheduled 1 The memory controller initializes memory once it is enabled. This bit remains asserted until the initialization is complete. The value in DDR_DATA_INIT register is used to initialize memory. |
| 28–31 | — | Reserved |

9.4.1.9 DDR SDRAM Mode Configuration (DDR_SDRAM_MODE)

The DDR SDRAM mode configuration register, shown in [Figure 9-10](#), sets the values loaded into the DDR’s mode registers.




Figure 9-10. DDR SDRAM Mode Configuration Register (DDR_SDRAM_MODE)

Table 9-14 describes the DDR_SDRAM_MODE fields.

Table 9-14. DDR_SDRAM_MODE Field Descriptions

| Bits | Name | Description |
|-------|---------|---|
| 0–15 | ESDMODE | Extended SDRAM mode Specifies the initial value loaded into the DDR SDRAM extended mode register. The range and meaning of legal values is specified by the DDR SDRAM manufacturer. When this value is driven onto the address bus (during the DDR SDRAM initialization sequence), MA[0] presents the lsb of ESDMODE, which, in the big-endian convention shown in Figure 9-10, corresponds to ESDMODE[15]. The msb of the SDRAM extended mode register value must be stored at ESDMODE[0]. |
| 16–31 | SDMODE | SDRAM mode Specifies the initial value loaded into the DDR SDRAM mode register. The range of legal values is specified by the DDR SDRAM manufacturer. When this value is driven onto the address bus (during DDR SDRAM initialization), MA[0] presents the lsb of SDMODE, which, in the big-endian convention shown in Figure 9-10, corresponds to SDMODE[15]. The msb of the SDRAM mode register value must be stored at SDMODE[0]. Because the memory controller forces SDMODE[7] to certain values depending on the state of the initialization sequence, the corresponding bits of this field are ignored by the memory controller. Note that SDMODE[7] is mapped to MA[8]. |

9.4.1.10 DDR SDRAM Mode 2 Configuration (DDR_SDRAM_MODE_2)

The DDR SDRAM mode 2 configuration register, shown in Figure 9-11, sets the values loaded into the DDR’s extended mode 2 and 3 registers (for DDR2).




Figure 9-11. DDR SDRAM Mode 2 Configuration Register (DDR_SDRAM_MODE_2)

Table 9-15 describes the DDR_SDRAM_MODE_2 fields.

Table 9-15. DDR_SDRAM_MODE_2 Field Descriptions

| Bits | Name | Description |
|-------|----------|--|
| 0–15 | ESDMODE2 | Extended SDRAM mode 2 Specifies the initial value loaded into the DDR SDRAM extended 2 mode register. The range and meaning of legal values is specified by the DDR SDRAM manufacturer. When this value is driven onto the address bus (during the DDR SDRAM initialization sequence), MA[0] presents the lsb bit of ESDMODE2, which, in the big-endian convention shown in Figure 9-11, corresponds to ESDMODE2[15]. The msb of the SDRAM extended mode 2 register value must be stored at ESDMODE2[0]. |
| 16–31 | ESDMODE3 | Extended SDRAM mode 3 Specifies the initial value loaded into the DDR SDRAM extended 3 mode register. The range of legal values of legal values is specified by the DDR SDRAM manufacturer. When this value is driven onto the address bus (during DDR SDRAM initialization), MA[0] presents the lsb of ESDMODE3, which, in the big-endian convention shown in Figure 9-11, corresponds to ESDMODE3[15]. The msb of the SDRAM extended mode 3 register value must be stored at ESDMODE3[0]. |

9.4.1.11 DDR SDRAM Mode Control Register (DDR_SDRAM_MD_CNTL)

The DDR SDRAM mode control register, shown in Figure 9-12, allows the user to carry out the following tasks:

- Issue a mode register set command to a particular chip select
- Issue an immediate refresh to a particular chip select
- Issue an immediate precharge or precharge all command to a particular chip select
- Force the CKE signals to a specific value




Figure 9-12. DDR SDRAM Mode Control Register (DDR_SDRAM_MD_CNTL)

Table 9-16 describes the fields of this register. Table 9-17 shows the user how to set the fields of this register to accomplish the above tasks.

NOTE

Note that MD_EN, SET_REF, and SET_PRE are mutually exclusive; only one of these fields can be set at a time.

Table 9-16. DDR_SDRAM_MD_CNTL Field Descriptions

| Bits | Name | Description |
|------|---------|---|
| 0 | MD_EN | <p>Mode enable</p> <p>Setting this bit specifies that valid data in MD_VALUE is ready to be written to DRAM as one of the following commands:</p> <ul style="list-style-type: none"> • MODE REGISTER SET • EXTENDED MODE REGISTER SET • EXTENDED MODE REGISTER SET 2 • EXTENDED MODE REGISTER SET 3 <p>The specific command to be executed is selected by setting MD_SEL. In addition, the chip select must be chosen by setting CS_SEL. MD_EN is set by software and cleared by hardware once the command has been issued.</p> <p>0 Indicates that no mode register set command needs to be issued. 1 Indicates that valid data contained in the register is ready to be issued as a mode register set command.</p> |
| 1 | — | Reserved |
| 2–3 | CS_SEL | <p>Select chip select</p> <p>Specifies the chip select that is driven active due to any command forced by software in DDR_SDRAM_MD_CNTL.</p> <p>00 Chip select 0 is active 01 Chip select 1 is active 10 Reserved 11 Reserved</p> |
| 4 | — | Reserved |
| 5–7 | MD_SEL | <p>Mode register select</p> <p>MD_SEL specifies one of the following:</p> <ul style="list-style-type: none"> • During a mode select command, selects the SDRAM mode register to be changed • During a precharge command, selects the SDRAM logical bank to be precharged. A precharge all command ignores this field. • During a refresh command, this field is ignored. <p>Note that MD_SEL contains the value that is presented onto the memory bank address pins (MBA_n) of the DDR controller.</p> <p>000 MR 001 EMR 010 EMR2 011 EMR3</p> |
| 8 | SET_REF | <p>Set refresh</p> <p>Forces an immediate refresh to be issued to the chip select specified by DDR_SDRAM_MD_CNTL[CS_SEL]. This bit is set by software and cleared by hardware once the command has been issued.</p> <p>0 Indicates that no refresh command needs to be issued. 1 Indicates that a refresh command is ready to be issued.</p> |
| 9 | SET_PRE | <p>Set precharge</p> <p>Forces a precharge or precharge all to be issued to the chip select specified by DDR_SDRAM_MD_CNTL[CS_SEL]. This bit is set by software and cleared by hardware once the command has been issued.</p> <p>0 Indicates that no precharge all command needs to be issued. 1 Indicates that a precharge all command is ready to be issued.</p> |

Table 9-16. DDR_SDRAM_MD_CNTL Field Descriptions (continued)

| Bits | Name | Description |
|-------|----------|---|
| 10–11 | CKE_CNTL | <p>Clock enable control</p> <p>Allows software to globally clear or set all CKE signals issued to DRAM. Once software has forced the value driven on CKE, that value continues to be forced until software clears the CKE_CNTL bits. At that time, the DDR controller continues to drive the CKE signals to the same value forced by software until another event causes the CKE signals to change (such as, self refresh entry/exit, power down entry/exit).</p> <p>00 CKE signals are not forced by software. 01 CKE signals are forced to a low value by software. 10 CKE signals are forced to a high value by software. 11 Reserved</p> |
| 12–15 | — | Reserved |
| 16–31 | MD_VALUE | <p>Mode register value</p> <p>This field, which specifies the value that is presented on the memory address pins of the DDR controller during a mode register set command, is significant only when this register is used to issue a mode register set command or a precharge or precharge all command.</p> <p>For a mode register set command, this field contains the data to be written to the selected mode register.</p> <p>For a precharge command, only bit five is significant:</p> <p>0 Issue a precharge command; MD_SEL selects the logical bank to be precharged 1 Issue a precharge all command; all logical banks are precharged</p> |

Table 9-17 shows how DDR_SDRAM_MD_CNTL fields should be set for each of the tasks described above.

Table 9-17. Settings of DDR_SDRAM_MD_CNTL Fields

| Field | Mode Register Set | Refresh | Precharge | Clock Enable Signals Control |
|----------|---------------------------------------|---------|---|------------------------------|
| MD_EN | 1 | 0 | 0 | — |
| SET_REF | 0 | 1 | 0 | — |
| SET_PRE | 0 | 0 | 1 | — |
| CS_SEL | Chooses chip select (CS) | | | — |
| MD_SEL | Select mode register. See Table 9-16. | — | Selects logical bank | — |
| MD_VALUE | Value written to mode register | — | Only bit five is significant. See Table 9-16. | — |
| CKE_CNTL | 0 | 0 | 0 | See Table 9-16. |

9.4.1.16 DDR IP Block Revision 1 (DDR_IP_REV1)

The DDR IP block revision 1 register, shown in Figure 9-17, provides read-only fields with the IP block ID, along with major and minor revision information.




Figure 9-17. DDR IP Block Revision 1 (DDR_IP_REV1)

¹ For reset values, see Table 9-22.

Table 9-22 describes the DDR_IP_REV1 fields.

Table 9-22. DDR_IP_REV1 Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0–15 | IP_ID | IP block ID. For the DDR controller, this value is 0x0002. |
| 16–23 | IP_MJ | Major revision. This is currently set to 0x02. |
| 24–31 | IP_MN | Minor revision. This is currently set to 0x01. |

9.4.1.17 DDR IP Block Revision 2 (DDR_IP_REV2)

The DDR IP block revision 2 register, shown in Figure 9-18, provides read-only fields with the IP block integration and configuration options.




Figure 9-18. DDR IP Block Revision 2 (DDR_IP_REV2)

Table 9-23 describes the DDR_IP_REV2 fields.

Table 9-23. DDR_IP_REV2 Field Descriptions

| Bits | Name | Description |
|-------|--------|--------------------------------|
| 0–7 | — | Reserved |
| 8–15 | IP_INT | IP block integration options |
| 16–23 | — | Reserved |
| 24–31 | IP_CFG | IP block configuration options |

9.4.1.18 Memory Error Detect (ERR_DETECT)

The memory error detect register stores the detection bits for memory select errors. It is a read/write register. A bit can be cleared by writing a one to the bit. System software can determine the type of memory error by examining the contents of this register. If an error is disabled with ERR_DISABLE, the corresponding error is never detected or captured in ERR_DETECT.

ERR_DETECT is shown in [Figure 9-19](#).



Figure 9-19. Memory Error Detect Register (ERR_DETECT)

[Table 9-24](#) describes the ERR_DETECT fields.

Table 9-24. ERR_DETECT Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–30 | — | Reserved |
| 31 | MSE | Memory select error. This bit is cleared by software writing a 1. 0 A memory select error has not been detected. 1 A memory select error has been detected. |

9.4.1.19 Memory Error Disable (ERR_DISABLE)

The memory error disable register, shown in [Figure 9-20](#), allows selective disabling of the DDR controller’s error detection circuitry. Disabled errors are not detected or reported.



Figure 9-20. Memory Error Disable Register (ERR_DISABLE)

[Table 9-25](#) describes the ERR_DISABLE fields.

Table 9-25. ERR_DISABLE Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–30 | — | Reserved |
| 31 | MSED | Memory select error disable 0 Memory select errors are enabled. 1 Memory select errors are disabled. |




Figure 9-22. DDR Memory Controller Block Diagram

Read and write accesses to memory are burst oriented; accesses start at a selected location and continue for a programmed number of higher locations (4 or 8) in a programmed sequence. Accesses to closed pages start with the registration of an ACTIVE command followed by a READ or WRITE. (Accessing open pages does not require an ACTIVE command.) The address bits registered coincident with the activate command specifies the logical bank and row to be accessed. The address coincident with the READ or WRITE command specify the logical bank and starting column for the burst access.

The data interface is source synchronous, meaning whatever sources the data also provides a clocking signal to synchronize data reception. These bidirectional data strobes (MDQS[0:3]) are inputs to the controller during reads and outputs during writes. The DDR SDRAM specification requires the data strobe signals to be centered within the data tenure during writes and to be offset by the controller to the center of the data tenure during reads. This delay is implemented in the controller for both reads and writes.

The address and command interface is also source synchronous, although 1/4 cycle adjustments are provided for adjusting the clock alignment.

Figure 9-23 shows an example DDR SDRAM configuration with four logical banks.

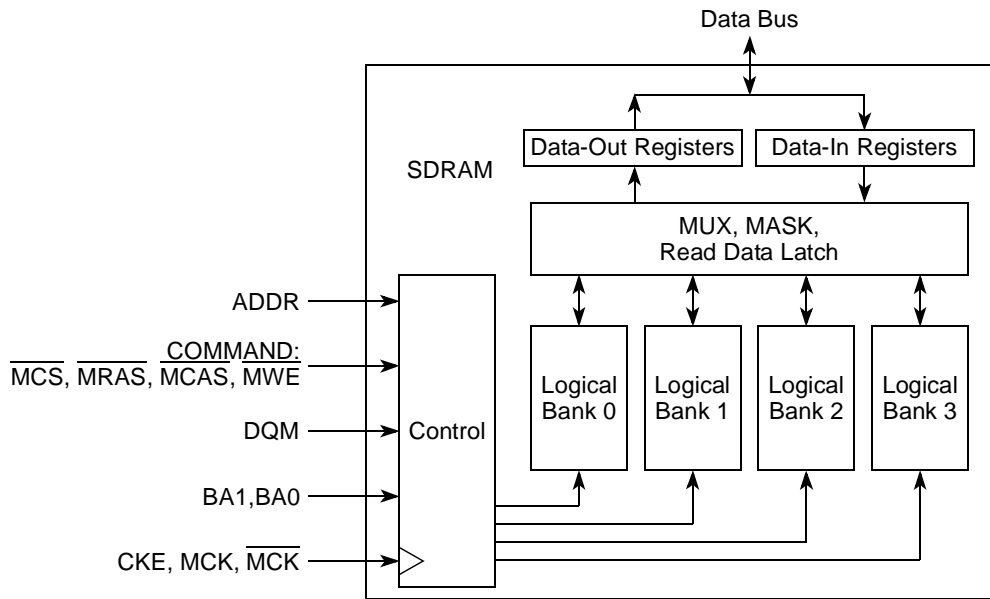


Figure 9-23. Typical Dual Data Rate SDRAM Internal Organization

Figure 9-24 shows some typical signal connections.

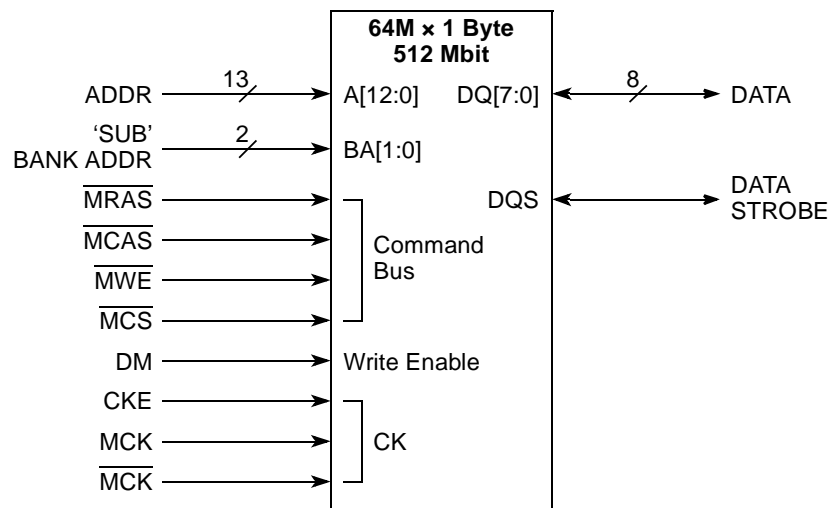
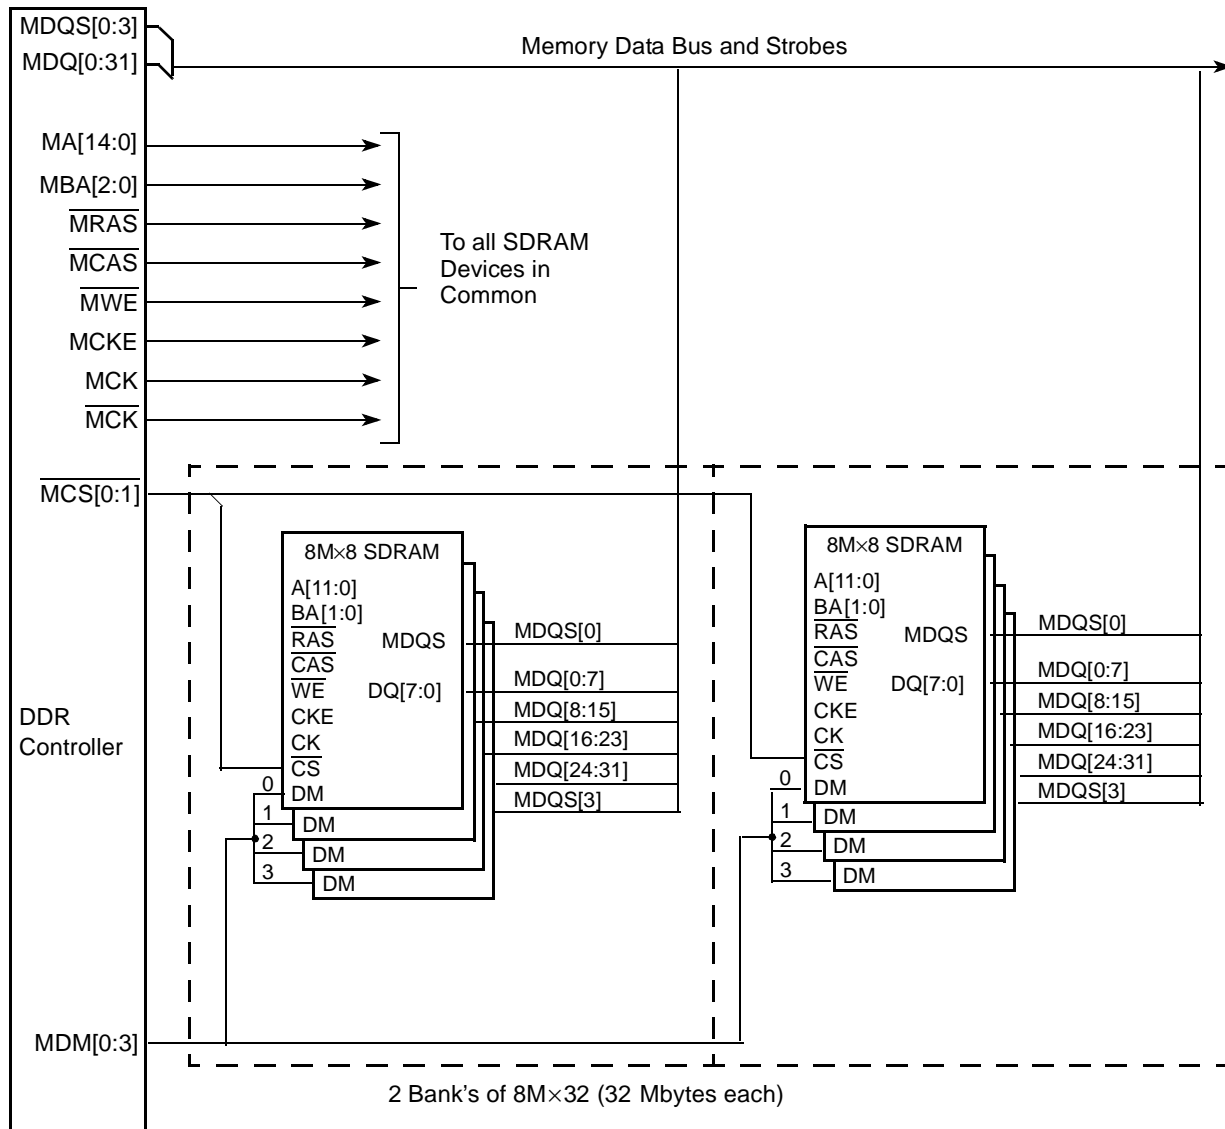


Figure 9-24. Typical DDR SDRAM Interface Signals

Figure 9-25 shows an example DDR SDRAM configuration with two physical banks each comprised of four $8\text{M} \times 8$ DDR modules for a total of 128 Mbytes of system memory. Certain address and control lines may require buffering. Analysis of the device's AC timing specifications, desired memory operating frequency, capacitive loads, and board routing loads can assist the system designer in deciding signal buffering requirements. The DDR memory controller drives 15 address pins, but in this example the DDR SDRAM devices use only 12 bits.



1. All signals are connected in common (in parallel) except for $\overline{MCS}[0:1]$, MCK, MDM[0:3], and the data bus signals.
2. Each of the $\overline{MCS}[0:1]$ signals correspond with a separate physical bank of memory.
3. Buffering may be needed if large memory arrays are used.

Figure 9-25. Example 64-Mbyte DDR SDRAM Configuration

9.5.1 DDR SDRAM Interface Operation

The DDR memory controller supports many different DDR SDRAM configurations. SDRAMs with different sizes can be used in the same system. Fifteen multiplexed address signals and three logical bank select signals support device densities from 64 Mbits to 2 Gbits. Two chip select (\overline{CS}) signals support one bank of memory. The DDR SDRAM physical banks can be built from standard memory modules or directly attached memory devices. The data path to individual physical banks is bits wide. The DDR memory controller supports physical bank sizes from 16 Mbytes to 2 Gbytes. The physical banks can be

constructed using $\times 8$, $\times 16$, or $\times 32$ memory devices. The memory technologies supported are 64 Mbits, 128 Mbits, 256 Mbits, 512 Mbits, 1 Gbit, 2 Gbits, and 4 Gbits. If two ranks of 4-Gbit devices were used, this would consume 4 Gbytes of memory. For a 32-bit address, no space would be left for IMMRBAR or any other address allocations. Four data qualifier (DQM) signals provide byte selection for memory accesses.

NOTE

An 8-bit DDR SDRAM device has a DQM signal and eight data signals (DQ[0:7]). A 16-bit DDR SDRAM device has two DQM signals associated with specific halves of the 16 data signals (DQ[0:7] and DQ[8:15]).

Table 9-27 shows the DDR memory controller's relationships between data byte lane 0–3, MDM[0:3], MDQS[0:3], and MDQ[0:31] when DDR SDRAM memories are used with $\times 8$ or $\times 16$ devices.

Table 9-27. Byte Lane to Data Relationship

| Data Byte Lane | Data Bus Mask | Data Bus Strobe | Data Bus |
|----------------|---------------|-----------------|------------|
| 0 (MSB) | MDM[0] | MDQS[0] | MDQ[0:7] |
| 1 | MDM[1] | MDQS[1] | MDQ[8:15] |
| 2 | MDM[2] | MDQS[2] | MDQ[16:23] |
| 3 | MDM[3] | MDQS[3] | MDQ[24:31] |

9.5.1.1 Supported DDR SDRAM Organizations

Although the DDR memory controller multiplexes row and column address bits onto 15 memory address signals and 3 logical bank select signals, a physical bank may be implemented with memory devices requiring fewer than 30 address bits. The physical bank may be configured to provide from 12 to 15 row address bits, plus 2 or 3 logical bank-select bits and from 8–11 column address bits.

Table 9-28 and Table 9-29 describe DDR SDRAM device configurations supported by the DDR memory controller.

NOTE

DDR SDRAM is limited to 30 total address bits.

Table 9-28. Supported DDR1 SDRAM Device Configurations

| SDRAM Device | Device Configuration | Row \times Column \times Sub-bank Bits | 32-Bit Bank Size | Two Banks of Memory |
|-----------------------|----------------------|--|------------------|---------------------|
| 64 Mbits | 8 Mbits \times 8 | 12 \times 9 \times 2 | 32 Mbytes | 64 Mbytes |
| 64 Mbits ¹ | 4 Mbits \times 16 | 12 \times 8 \times 2 | 16 Mbytes | 32 Mbytes |
| 128 Mbits | 16 Mbits \times 8 | 12 \times 10 \times 2 | 64 Mbytes | 128 Mbytes |
| 128 Mbits | 8 Mbits \times 16 | 12 \times 9 \times 2 | 16 Mbytes | 64 Mbytes |
| 256 Mbits | 32 Mbits \times 8 | 13 \times 10 \times 2 | 128 Mbytes | 256 Mbytes |
| 256 Mbits | 16 Mbits \times 16 | 13 \times 9 \times 2 | 64 Mbytes | 128 Mbytes |
| 512 Mbits | 64 Mbits \times 8 | 13 \times 11 \times 2 | 256 Mbytes | 512 Mbytes |

Table 9-28. Supported DDR1 SDRAM Device Configurations (continued)

| SDRAM Device | Device Configuration | Row × Column × Sub-bank Bits | 32-Bit Bank Size | Two Banks of Memory |
|--------------|----------------------|------------------------------|------------------|---------------------|
| 512 Mbits | 32 Mbits × 16 | 13 × 10 × 2 | 128 Mbytes | 256 Mbytes |
| 1 Gbits | 128 Mbits × 8 | 14 × 11 × 2 | 512 Mbytes | 1 Gbyte |
| 1 Gbits | 64 Mbits × 16 | 14 × 10 × 2 | 256 Mbytes | 512 Mbytes |
| 2 Gbits | 256 Mbits × 8 | 15 × 11 × 2 | 1 Gbyte | 2 Gbytes |
| 2 Gbits | 128 Mbits × 16 | 15 × 10 × 2 | 512 Mbytes | 1 Gbyte |

¹ This configuration is not supported in 16-bit bus mode.

Table 9-29. Supported DDR2 SDRAM Device Configurations

| SDRAM Device | Device Configuration | Row × Column × Sub-bank Bits | 32-Bit Bank Size | Two Banks of Memory |
|--------------|----------------------|------------------------------|------------------|---------------------|
| 256 Mbits | 32 Mbits × 8 | 13 × 10 × 2 | 128 Mbytes | 256 Mbytes |
| 256 Mbits | 16 Mbits × 16 | 13 × 9 × 2 | 64 Mbytes | 128 Mbytes |
| 512 Mbits | 64 Mbits × 8 | 14 × 10 × 2 | 256 Mbytes | 512 Mbytes |
| 512 Mbits | 32 Mbits × 16 | 13 × 10 × 2 | 128 Mbytes | 256 Mbytes |
| 1 Gbits | 128 Mbits × 8 | 14 × 10 × 3 | 512 Mbytes | 1 Gbyte |
| 1 Gbits | 64 Mbits × 16 | 13 × 10 × 3 | 256 Mbytes | 512 Mbytes |
| 2 Gbits | 256 Mbits × 8 | 15 × 10 × 3 | 1 Gbyte | 2 Gbytes |
| 2 Gbits | 128 Mbits × 16 | 14 × 10 × 3 | 512 Mbytes | 1 Gbyte |
| 4 Gbits | 512 Mbits × 8 | 15 × 11 × 3 | 2 Gbytes | 4 Gbytes |
| 4 Gbits | 256 Mbits × 16 | 15 × 10 × 3 | 1 Gbyte | 2 Gbytes |

If a transaction request is issued to the DDR memory controller and the address does not lie within any of the programmed address ranges for an enabled chip select, a memory select error is flagged.

Using a memory-polling algorithm at power-on reset or by querying the JEDEC serial presence detect capability of memory modules, system firmware uses the memory-boundary registers to configure the DDR memory controller to map the size of each bank in memory. The memory controller uses its bank map to assert the appropriate \overline{MCS}_n signal for memory accesses according to the provided bank starting and ending addresses. The memory banks are not required to be mapped to a contiguous address space.

9.5.2 DDR SDRAM Address Multiplexing

, Table 9-30, Table 9-31, Table 9-32, Table 9-33 show the address bit encodings for each DDR SDRAM configuration. The address presented at the memory controller signals MA[14:0] use MA[14] as the msb

and MA[0] as the lsb. Also, MA[10] is used as the auto-precharge bit in DDR1/DDR2 modes for reads and writes, so the column address can never use MA[10].

Table 9-30. DDR1 Address Multiplexing for 32-Bit Data Bus with Interleaving Disabled

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | |
|-----------------|------|--------------------------|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|-------|
| | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | | 29 | 30-31 |
| 15 x 11 x 2 | MRAS | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 15 x 10 x 2 | MRAS | | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 14 x 11 x 2 | MRAS | | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 14 x 10 x 2 | MRAS | | | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 13 x 11 x 2 | MRAS | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 13 x 10 x 2 | MRAS | | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 13 x 9 x 2 | MRAS | | | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 12 x 10 x 2 | MRAS | | | | | | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 12 x 9 x 2 | MRAS | | | | | | | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 12 x 8 x 2 | MRAS | | | | | | | | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | | | | 7 | 6 | 5 | 4 | 3 | 2 | 1 |

Table 9-31. DDR1 Address Multiplexing for 16-Bit Data Bus

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | |
|-------------------|--------------------------|--------------------------|---|---|----|----|----|----|----|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|----|--|
| | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | | 30 | 31 | |
| 15 x 11 x 2 | $\overline{\text{MRAS}}$ | | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | |
| 15 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | |
| 14 x 11 x 2 | $\overline{\text{MRAS}}$ | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 14 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | |
| 13 x 11 x 2 | $\overline{\text{MRAS}}$ | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 13 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | |
| 13 x 9 x 2 | $\overline{\text{MRAS}}$ | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | |
| 12 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | |
| 12 x 9 x 2 | $\overline{\text{MRAS}}$ | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | |
| 12 x 8 x 2 | $\overline{\text{MRAS}}$ | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |

Table 9-32. DDR2 Address Multiplexing for 16-Bit Data Bus

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | | |
|-------------------|--------------------------|--------------------------|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|----|--|--|
| | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | | 30 | 31 | | |
| 15 x 10 x 3 | $\overline{\text{MRAS}}$ | | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 14 x 10 x 3 | $\overline{\text{MRAS}}$ | | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 14 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 13 x 10 x 3 | $\overline{\text{MRAS}}$ | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 13 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 13 x 9 x 2 | $\overline{\text{MRAS}}$ | | | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |

Table 9-33. DDR2 Address Multiplexing

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | | |
|-----------------|--------------------------|--------------------------|----|----|----|----|----|----|----|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|-------|--|--|--|
| | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | | 30-31 | | | |
| 15 x 11 x 3 | $\overline{\text{MRAS}}$ | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 15 x 10 x 3 | $\overline{\text{MRAS}}$ | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 14 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |

Table 9-33. DDR2 Address Multiplexing (continued)

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | |
|-----------------|------|--------------------------|---|---|---|----|----|----|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|----|-------|
| | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | | 28 | 29 | 30-31 |
| 13 x 10 x 3 | MRAS | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 13 x 10 x 2 | MRAS | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 13 x 9 x 2 | MRAS | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |

Chip select interleaving is supported for the memory controller, and is programmed in `DDR_SDRAM_CFG[BA_INTLV_CTL]`. Interleaving is supported between chip selects 0 and 1. When interleaving is enabled, the chip selects being interleaved must use the same size of memory. One extra bit in the address decode is used for the interleaving to determine which chip select to access.

Table 9-34 illustrates examples of address decode when interleaving between two chip selects.

Table 9-34. Example of Address Multiplexing for 32-Bit Data Bus Interleaving between Two Banks with Partial Array Self Refresh Disabled

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | |
|-----------------|------|--------------------------|----|----|----|----|----|---|---|---|----|----|----|----|----|----|----|----|----|--------|--------|----|----|----|----|----|----|----|----|-----|----|-------|---|
| | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | | 29 | 30-31 | |
| 14 x 10 x 3 | MRAS | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | CS SEL | 2 | 1 | 0 | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 14 x 10 x 2 | MRAS | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | CS SEL | 1 | 0 | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 13 x 10 x 3 | MRAS | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | CS SEL | 2 | 1 | 0 | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 13 x 10 x 2 | MRAS | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | CS SEL | 1 | 0 | | | | | | | | | | |
| | MCAS | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |

9.5.3 JEDEC Standard DDR SDRAM Interface Commands

The following section describes the commands and timings the controller uses when operating in DDR2 or DDR modes.

All read or write accesses to DDR SDRAM are performed by the DDR memory controller using JEDEC standard DDR SDRAM interface commands. The SDRAM device samples command and address inputs on rising edges of the memory clock; data is sampled using both the rising and falling edges of DQS. Data read from the DDR SDRAM is also sampled on both edges of DQS.

The following DDR SDRAM interface commands (summarized in [Table 9-35](#)) are provided by the DDR controller. All actions for these commands are described from the perspective of the SDRAM device.

- Row activate—Latches row address and initiates memory read of that row. Row data is latched in SDRAM sense amplifiers and must be restored by a precharge command before another row activate occurs.
- Precharge—Restores data from the sense amplifiers to the appropriate row. Also initializes the sense amplifiers in preparation for reading another row in the memory array (performing another activate command). Precharge must occur after read or write, if the row address changes on the next open page mode access.
- Read—Latches column address and transfers data from the selected sense amplifier to the output buffer as determined by the column address. During each succeeding clock edge, additional data is driven without additional read commands. The amount of data transferred is determined by the burst size which defaults to 4.
- Write—Latches column address and transfers data from the data pins to the selected sense amplifier as determined by the column address. During each succeeding clock edge, additional data is transferred to the sense amplifiers from the data pins without additional write commands. The amount of data transferred is determined by the data masks and the burst size, which is set to four by the DDR memory controller.
- Refresh (similar to $\overline{\text{MCAS}}$ before $\overline{\text{MRAS}}$)—Causes a row to be read in all logical banks (JEDEC SDRAM) as determined by the refresh row address counter. This refresh row address counter is internal to the SDRAM. After being read, the row is automatically rewritten in the memory array. All logical banks must be in a precharged state before executing a refresh. The memory controller also supports posted refreshes, where several refreshes may be executed at once, and the refresh interval may be extended.
- Mode register set (for configuration)—Allows setting of DDR SDRAM options. These options are: $\overline{\text{MCAS}}$ latency, additive latency (for DDR2), write recovery (for DDR2), burst type, and burst length. $\overline{\text{MCAS}}$ latency may be chosen as provided by the preferred SDRAM (some SDRAMs provide $\overline{\text{MCAS}}$ latency {1,2,3}, some provide $\overline{\text{MCAS}}$ latency {1,2,3,4,5}, and so on). Burst type is always sequential. Although some SDRAMs provide burst lengths of 1, 2, 4, 8, and page size, this memory controller supports a burst length of 4. A burst length of 8 is supported for DDR1 memory only. For DDR2 in 32-bit bus mode, all 32-byte burst accesses from the platform are split into two 16-byte (that is, 4-beat) accesses to the SDRAMs in the memory controller. The mode register set command is performed by the DDR memory controller during system initialization. Parameters such as mode register data, $\overline{\text{MCAS}}$ latency, burst length, and burst type, are set by software in `DDR_SDRAM_MODE[SDMODE]` and transferred to the SDRAM array by the DDR

memory controller after DDR_SDRAM_CFG[MEM_EN] is set. If DDR_SDRAM_CFG[BI] is set to bypass the automatic initialization, then the MODE registers can be configured through software through use of the DDR_SDRAM_MD_CNTL register.

- Self refresh (for long periods of standby)—Used when the device is in standby for very long periods of time. Automatically generates internal refresh cycles to keep the data in all memory banks refreshed. Before execution of this command, the DDR controller places all logical banks in a precharged state.

Table 9-35. DDR SDRAM Command Table

| Operation | CKE Prev. | CKE Current | $\overline{\text{MCS}}$ | $\overline{\text{MRAS}}$ | $\overline{\text{MCAS}}$ | $\overline{\text{MWE}}$ | MBA | MA10 | MA |
|-------------------------------|-----------|-------------|-------------------------|--------------------------|--------------------------|-------------------------|---------------------|--------|-----------------|
| Activate | H | H | L | L | H | H | Logical bank select | Row | Row |
| Precharge select logical bank | H | H | L | L | H | L | Logical bank select | L | X |
| Precharge all logical banks | H | H | L | L | H | L | X | H | X |
| Read | H | H | L | H | L | H | Logical bank select | L | Column |
| Read with auto-precharge | H | H | L | H | L | H | Logical bank select | H | Column |
| Write | H | H | L | H | L | L | Logical bank select | L | Column |
| Write with auto-precharge | H | H | L | H | L | L | Logical bank select | H | Column |
| Mode register set | H | H | L | L | L | L | Opcode | Opcode | Opcode and mode |
| Auto refresh | H | H | L | L | L | H | X | X | X |
| Self refresh | H | L | L | L | L | H | X | X | X |

9.5.4 DDR SDRAM Interface Timing

The DDR memory controller supports four-beat bursts to SDRAM. For single-beat reads, the DDR memory controller performs a four- (or eight-) beat burst read, but ignores the last three (or seven) beats. Single-beat writes are performed by masking the last three (or seven) beats of the four- (or eight-) beat burst using the data mask MDM[0:3].

NOTE

If a second read or write is pending, reads shorter than four beats are not terminated early even if some data is irrelevant.

To accommodate available memory technologies across a wide spectrum of operating frequencies, the DDR memory controller allows the setting of the intervals defined in [Table 9-36](#) with granularity of one memory clock cycle, except for CASLAT, which can be programmed with ½ clock granularity.

Table 9-36. DDR SDRAM Interface Timing Intervals

| Timing Intervals | Definition |
|------------------|---|
| ACTTOACT | The number of clock cycles from a bank-activate command until another bank-activate command within a physical bank. This interval is listed in the AC specifications of the SDRAM as t_{RRD} . |
| ACTTOPRE | The number of clock cycles from an activate command until a precharge command is allowed. This interval is listed in the AC specifications of the SDRAM as t_{RAS} . |
| ACTTORW | The number of clock cycles from an activate command until a read or write command is allowed. This interval is listed in the AC specifications of the SDRAM as t_{RCD} . |
| BSTOPRE | The number of clock cycles to maintain a page open after an access. The page open duration counter is reloaded with BSTOPRE each time the page is accessed (including page hits). When the counter expires, the open page is closed with an SDRAM precharge bank command as soon as possible. |
| CASLAT | Used in conjunction with additive latency to obtain the READ latency. The number of clock cycles between the registration of a READ command by the SDRAM and the availability of the first piece of output data. If a READ command is registered at clock edge n , and the read latency is m clocks, the data is available nominally coincident with clock edge $n + m$. |
| PRETOACT | The number of clock cycles from a precharge command until an activate or a refresh command is allowed. This interval is listed in the AC specifications of the SDRAM as t_{RP} . |
| REFINT | Refresh interval. Represents the number of memory bus clock cycles between refresh cycles. Depending on <code>DDR_SDRAM_CFG_2[NUM_PR]</code> , some number of rows are refreshed in each SDRAM bank during each refresh cycle. The value of REFINT depends on the specific SDRAMs used and the frequency of the interface as t_{RP} . |
| REFREC | The number of clock cycles from the refresh command until an activate command is allowed. This can be calculated by referring to the AC specification of the SDRAM device. The AC specification indicates a maximum refresh-to-activate interval in nanoseconds. |
| WR_DATA_DELAY | Provides different options for the timing between a write command and the write data strobe. This allows write data to be sent later than the nominal time to meet the SDRAM timing requirement between the registration of a write command and the reception of a data strobe associated with the write command. The specification dictates that the data strobe may not be received earlier than 75% of a cycle, or later than 125% of a cycle, from the registration of a write command. This parameter is not defined in the SDRAM specification. It is implementation-specific, defined for the DDR memory controller in <code>TIMING_CFG_2</code> . |
| WRREC | The number of clock cycles from the last beat of a write until a precharge command is allowed. This interval, write recovery time, is listed in the AC specifications of the SDRAM as t_{WR} . |
| WRTORD | Last write pair to read command. Controls the number of clock cycles from the last write data pair to the subsequent read command to the same bank as t_{WTR} . |

The value of the above parameters (in whole clock cycles) must be set by boot code at system start-up (in the `TIMING_CFG_0`, `TIMING_CFG_1`, `TIMING_CFG_2`, and `TIMING_CFG_3` registers as described in [Section 9.4.1.4](#), “DDR SDRAM Timing Configuration 0 (`TIMING_CFG_0`),” [Section 9.4.1.5](#), “DDR SDRAM Timing Configuration 1 (`TIMING_CFG_1`),” [Section 9.4.1.6](#), “DDR SDRAM Timing Configuration 2 (`TIMING_CFG_2`),” and [Section 9.4.1.3](#), “DDR SDRAM Timing Configuration 3 (`TIMING_CFG_3`)”) and be kept in the DDR memory controller configuration register space.

The following figures show SDRAM timing for various types of accesses. System software is responsible (at reset) for optimally configuring SDRAM timing parameters. The programmable timing parameters apply to both read and write timing configuration. The configuration process must be completed and the DDR SDRAM initialized before any accesses to SDRAM are attempted.

Figure 9-26 through Figure 9-28 show DDR SDRAM timing for various types of accesses; see Figure 9-26 for a single-beat read operation, Figure 9-27 for a single-beat write operation, and Figure 9-28 for a double word write operation. Note that all signal transitions occur on the rising edge of the memory bus clock and that single-beat read operations are identical to burst-reads. These figures assume the CLK_ADJUST is set to 1/2 DRAM cycle, an additive latency of 0 DRAM cycles is used, and the write latency is 1 DRAM cycle (for DDR1).




Figure 9-26. DDR SDRAM Burst Read Timing—ACTTORW = 3, MCAS Latency = 2




Figure 9-27. DDR SDRAM Single-Beat (Double Word) Write Timing—ACTOR




Figure 9-28. DDR SDRAM Single-Beat (Double Word) Write Timing—ACTORW = 3

9.5.4.1 Clock Distribution

Clock distribution has the following recommendations:

- If running with many devices, zero-delay PLL clock buffers, JEDEC-JESD82 standard, should be used. These buffers were designed for DDR applications.
- PCB traces for DDR clock signals should be short, all on the same layer, and of equal length and loading.

DDR SDRAM manufacturers provide detailed information on PCB layout and termination issues.




Figure 9-29. DDR SDRAM Clock Distribution Example for x8 DDR SDRAMs

9.5.5 DDR SDRAM Mode-Set Command Timing

The DDR memory controller transfers the mode register set commands to the SDRAM array, and it uses the setting of TIMING_CFG_0[MRS_CYC] for the Mode Register Set cycle time.

Figure 9-30 shows the timing of the mode-set command. The first transfer corresponds to the ESDMODE code; the second corresponds to SDMODE. The Mode Register Set cycle time is set to 2 DRAM cycles.




Figure 9-30. DDR SDRAM Mode-Set Command Timing

9.5.6 DDR SDRAM Registered DIMM Mode

To reduce loading, registered DRAM modules latch the DDR SDRAM control signals internally before using them to access the array. Setting `DDR_SDRAM_CFG[RD_EN]` compensates for this delay on the DRAM modules' control bus by delaying the data and data mask writes (on SDRAM buses) by an extra SDRAM clock cycle.

NOTE

Application system board must assert the reset signal on DDR memory devices until software is able to program the DDR memory controller configuration registers, and must deassert the reset signal on DDR memory devices before `DDR_SDRAM_CFG[MEM_EN]` is set. This ensures that the DDR memory devices are held in reset until a stable clock is provided and, further, that a stable clock is provided before memory devices are released from reset.

Figure 9-31 shows the registered DDR SDRAM DIMM single-beat write timing.




Figure 9-31. Registered DDR SDRAM DIMM Burst Write Timing

9.5.7 DDR SDRAM Write Timing Adjustments

The DDR memory controller facilitates system design flexibility by providing a write timing adjustment parameter, write data delay, (`TIMING_CFG_2[WR_DATA_DELAY]`) for data and DQS. The DDR SDRAM specification requires DQS be received no sooner than 75% of an SDRAM clock period—and no later than 125% of a clock period—from the capturing clock edge of the command/address at the SDRAM. `TIMING_CFG_2[WR_DATA_DELAY]` specifies how much to delay the launching of DQS and data from the first clock edge occurring one SDRAM clock cycle after the command is launched. The delay increment step sizes are in 1/4 SDRAM clock periods starting with the default value of 0.

Figure 9-32 shows the use of the WR_DATA_DELAY parameter.




Figure 9-32. Write Timing Adjustments Example for Write Latency = 1

9.5.8 DDR SDRAM Refresh

The DDR memory controller supports auto-refresh and self-refresh. Auto-refresh is used during normal operation and is controlled by the `DDR_SDRAM_INTERVAL[REFINT]` value; self-refresh is used only when the DDR memory controller is set to enter a sleep power management state. The REFINT value, which represents the number of memory bus clock cycles between refresh cycles, must allow for possible outstanding transactions to complete before a refresh request is sent to the memory after the REFINT value is reached. If a memory transaction is in progress when the refresh interval is reached, the refresh cycle waits for the transaction to complete. In the worst case, the refresh cycle must wait the number of bus clock cycles required by the longest programmed access. To ensure that the latency caused by a memory transaction does not violate the device refresh period, it is recommended that the programmed value of REFINT be less than that required by the SDRAM.

When a refresh cycle is required, the DDR memory controller does the following:

1. Completes all current memory requests.
2. Closes all open pages with a PRECHARGE-ALL command to each DDR SDRAM bank with an open page (as indicated by the row open table).
3. Issues one or more auto-refresh commands to each DDR SDRAM bank (as identified by its chip select) to refresh one row in each logical bank of the selected physical bank.

The auto-refresh commands are staggered across the two possible banks to reduce the system’s instantaneous power requirements. Three sets of auto refresh commands are issued on consecutive cycles when the memory is populated with one DIMMs. The initial PRECHARGE-ALL commands are also staggered in three groups for convenience. It is important to note that when entering self-refresh mode, only one refresh command is issued simultaneously to all physical banks. For this entire refresh sequence, no cycle optimization occurs for the usual case where fewer than two banks are installed. After the refresh sequence completes, any pending memory request is initiated after an inactive period specified by TIMING_CFG_1 [REFREC] and TIMING_CFG_3[EXT_REFREC]. In addition, posted refreshes are supported to allow the refresh interval to be set to a larger value.

9.5.8.1 DDR SDRAM Refresh Timing

Refresh timing for the DDR SDRAM is controlled by the programmable timing parameter TIMING_CFG_1 [REFREC], which specifies the number of memory bus clock cycles from the refresh command is allowed. The DDR memory controller implements bank staggering for refreshes, as shown in Figure 9-33 (TIMING_CFG_1 [REFREC] = 10 in this example).




Figure 9-33. DDR SDRAM Bank Staggered Auto Refresh Timing

System software is responsible for optimal configuration of TIMING_CFG_1 [REFREC] and TIMING_CFG_3[EXT_REFREC] at reset. Configuration must be completed before DDR SDRAM accesses are attempted.

9.5.8.2 DDR SDRAM Refresh and Power-Saving Modes

In full-on mode, the DDR memory controller supplies the normal auto refresh to SDRAM. In sleep mode, the DDR memory controller can be configured to take advantage of self-refreshing SDRAMs or to provide no refresh support. Self-refresh support is enabled with the SREN memory control parameter.

Table 9-37 summarizes the refresh types available in each power-saving mode.

Table 9-37. DDR SDRAM Power-Saving Modes Refresh Configuration

| Power Saving Mode | Refresh Type | SREN |
|-------------------|--------------|------|
| Sleep | Self | 1 |
| | None | — |

Note that in the absence of refresh support, system software must preserve DDR SDRAM data (such as by copying the data to disk) before entering the power-saving mode.

The dynamic power-saving mode uses the CKE DDR SDRAM pin to dynamically power down when there is no system memory activity. The CKE pin is negated when both of the following conditions are met:

- No memory refreshes are scheduled
- No memory accesses are scheduled

CKE is reasserted when a new access or refresh is scheduled or the dynamic power mode is disabled. This mode is controlled with DDR_SDRAM_CFG[DYN_PWR_MGMT].

Dynamic power management mode offers tight control of the memory system’s power consumption by trading power for performance through the use of CKE. Powering up the DDR SDRAM when a new memory reference is scheduled causes an access latency penalty, depending on whether active or precharge powerdown is used, along with the settings of TIMING_CFG_0[ACT_PD_EXIT] and TIMING_CFG_0[PRE_PD_EXIT]. A penalty of 1 cycle is shown in Figure 9-34.




Figure 9-34. DDR SDRAM Power-Down Mode

9.5.8.2.1 Self-Refresh in Sleep Mode

The entry and exit timing for self-refreshing SDRAMs is shown in [Figure 9-35](#) and [Figure 9-36](#).




Figure 9-35. DDR SDRAM Self-Refresh Entry Timing




Figure 9-36. DDR SDRAM Self-Refresh Exit Timing

9.5.9 DDR Data Beat Ordering

Transfers to and from memory are always performed in four- or eight-beat bursts. For transfer sizes other than four or eight beats, the data transfers are still operated as four- or eight-beat bursts. The DDR memory controller uses data masks to prevent all unintended full double words from writing to SDRAM. For example, if a write transaction is desired with a size of one word (4 bytes), then the second, third, and fourth beats of data are not written to DRAM, as the width of the data bus is 32 bits.

Table 9-38 lists the data beat sequencing to and from the DDR SDRAM and the data queues for each of the possible transfer sizes with each of the possible starting double-word offsets. All underlined double-word offsets are valid for the transaction.

Table 9-38. Memory Controller—Data Beat Ordering

| Transfer Size | Starting Double-Word Offset | Double-Word Sequence ¹ to/from DRAM and Queues |
|----------------|-----------------------------|---|
| 1 double word | 0 | <u>0</u> - 1 - 2 - 3 |
| | 1 | <u>1</u> - 2 - 3 - 0 |
| | 2 | <u>2</u> - 3 - 0 - 1 |
| | 3 | <u>3</u> - 0 - 1 - 2 |
| 2 double words | 0 | <u>0-1</u> - 2 - 3 |
| | 1 | <u>1-2</u> - 3 - 0 |
| | 2 | <u>2-3</u> - 0 - 1 |
| 3 double words | 0 | <u>0-1-2</u> - 3 |
| | 1 | <u>1-2-3</u> - 0 |

¹ All underlined **Double**-word offsets are valid for the transaction.

9.5.10 Page Mode and Logical Bank Retention

The DDR memory controller supports an open/closed page mode with an allowable open page for each logical bank of DRAM used. In closed page mode for DDR SDRAMs, the DDR memory controller uses the SDRAM auto-precharge feature, which allows the controller to indicate that the page must be automatically closed by the DDR SDRAM after the READ or WRITE access. This is performed using MA[10] of the address during the COMMAND phase of the access to enable auto-precharge. Auto-precharge is non-persistent in that it is either enabled or disabled for each individual READ or WRITE command. It can, however, be enabled or disabled separately for each chip select.

When the DDR memory controller operates in open page mode, it retains the currently active SDRAM page by not issuing a precharge command. The page remains open until one of the following conditions occurs:

- Refresh interval is met.
- The user-programmable DDR_SDRAM_INTERVAL[BSTOPRE] value is exceeded.
- There is a logical bank row collision with another transaction that must be issued.

Page mode can dramatically reduce access latencies for page hits. Depending on the memory system design and timing parameters, using page mode can save two to three clock cycles for subsequent burst accesses that hit in an active page. Also, better performance can be obtained using more banks, especially

in systems which use many different channels. Page mode is disabled by clearing DDR_SDRAM_INTERVAL[BSTOPRE] or setting CS_n_CONFIG[AP_nEN].

9.6 Initialization/Application Information

System software must configure the DDR memory controller, using a memory polling algorithm at system start-up, to correctly map the size of each bank in memory. Then, the DDR memory controller uses its bank map to assert the appropriate \overline{MCS}_n signal for memory accesses according to the provided bank depths. System software must also configure the DDR memory controller at system start-up to appropriately multiplex the row and column address bits for each bank. Refer to row-address configuration in [Section 9.4.1.2, “Chip Select Configuration \(CS_n_CONFIG\).”](#) Address multiplexing occurs according to these configuration bits.

At system reset, initialization software (boot code) must set up the programmable parameters in the memory interface configuration registers. See [Section 9.4.1, “Register Descriptions,”](#) for more detailed descriptions of the configuration registers. These parameters are shown in [Table 9-39.](#)

Table 9-39. Memory Interface Configuration Register Initialization Parameters

| Name | Description | Parameter | Section/page |
|-------------------------|---|---|------------------------------|
| CS _n _BNDS | Chip select memory bounds | SA _n EA _n | 9.4.1.1/9-9 |
| CS _n _CONFIG | Chip select configuration | CS _n _EN AP _n _EN ODT_RD_CFG ODT_WR_CFG BA_BITS_CS _n ROW_BITS_CS _n COL_BITS_CS _n | 9.4.1.2/9-10 |
| TIMING_CFG_3 | Extended timing parameters for fields in TIMING_CFG_1 | EXT_REFREC | 9.4.1.3/9-12 |
| TIMING_CFG_0 | Timing configuration | RWT WRT RRT WWT ACT_PD_EXIT PRE_PD_EXIT ODT_PD_EXIT MRS_CYC | 9.4.1.4/9-12 |
| TIMING_CFG_1 | Timing configuration | PRETOACT ACTTOPRE ACTTORW CASLAT REFREC WRREC ACTTOACT WRTORD | 9.4.1.5/9-14 |
| TIMING_CFG_2 | Timing configuration | ADD_LAT CPO WR_LAT RD_TO_PRE WR_DATA_DELAY CKE_PLS FOUR_ACT | 9.4.1.6/9-16 |
| DDR_SDRAM_CFG | Control configuration | SREN RD_EN SDRAM_TYPE DYN_PWR 32_BE 8_BE DBW NCAP 2T_EN BA_INTLV_CTL x32_EN HSE BI | 9.4.1.7/9-18 |

Table 9-39. Memory Interface Configuration Register Initialization Parameters (continued)

| Name | Description | Parameter | Section/page |
|--------------------|--|-----------------------------|----------------------------------|
| DDR_SDRAM_CFG_2 | Control configuration | SR_IE DQS_CFG ODT_CFG | NUM_PR D_INIT 9.4.1.8/9-21 |
| DDR_SDRAM_MODE | Mode configuration | ESDMODE SDMODE | 9.4.1.9/9-22 |
| DDR_SDRAM_MODE_2 | Mode configuration | ESDMODE2 ESDMODE3 | 9.4.1.10/9-23 |
| DDR_SDRAM_INTERVAL | Interval configuration | REFINT BSTOPRE | 9.4.1.12/9-27 |
| DDR_DATA_INIT | Data initialization configuration register | INIT_VALUE | 9.4.1.13/9-27 |
| DDR_SDRAM_CLK_CNTL | Clock adjust | CLK_ADJUST | 9.4.1.14/9-28 |
| DDR_INIT_ADDR | Initialization address | INIT_ADDR | 9.4.1.15/9-28 |

9.6.1 Programming Differences between Memory Types

Depending on the memory type used, certain fields must be programmed differently. Table 9-40 illustrates the differences in certain fields for different memory types. Note that this table does not list all fields that must be programmed.

Table 9-40. Programming Differences between Memory Types

| Parameter | Description | Differences | | Section/page |
|---------------------|--|-------------|--|--------------|
| AP _n _EN | Chip Select <i>n</i> Auto Precharge Enable | | | 9.4.1.2/9-10 |
| | | DDR1 | Can be used to place chip select <i>n</i> in auto precharge mode | |
| | | DDR2 | Can be used to place chip select <i>n</i> in auto precharge mode | |
| ODT_RD_CFG | Chip Select ODT Read Configuration | DDR1 | Should always be set to 000 | 9.4.1.2/9-10 |
| | | DDR2 | Can be enabled to assert ODT if desired. This could be set differently depending on system topology. However, systems with only 1 chip select typically not uses ODT when issuing reads to the memory. | |
| ODT_WR_CFG | Chip Select ODT Write Configuration | DDR1 | Should always be set to 000 | 9.4.1.2/9-10 |
| | | DDR2 | Can be enabled to assert ODT if desired. This could be set differently depending on system topology. However, ODT typically is set to assert for the chip select that is getting written to (value would be set to 001). | |
| ODT_PD_EXIT | ODT Powerdown Exit | DDR1 | Should be set to 0001 | 9.4.1.4/9-12 |
| | | DDR2 | Should be set according to the DDR2 specifications for the memory used. The JEDEC parameter this applies to is t_{AXPD} . | |

Table 9-40. Programming Differences between Memory Types (continued)

| Parameter | Description | Differences | | Section/page |
|-----------|-------------------------------|-------------|--|--------------|
| PRETOACT | Precharge to Activate Timing | DDR1 | Should be set according to the specifications for the memory used (t_{RP}) | 9.4.1.5/9-14 |
| | | DDR2 | Should be set according to the specifications for the memory used (t_{RP}) | |
| ACTTOPRE | Activate to Precharge Timing | DDR1 | Should be set, along with the Extended Activate to Precharge Timing, according to the specifications for the memory used (t_{RAS}) | 9.4.1.5/9-14 |
| | | DDR2 | Should be set, along with the Extended Activate to Precharge Timing, according to the specifications for the memory used (t_{RAS}) | |
| ACTTORW | Activate to Read/Write Timing | DDR1 | Should be set according to the specifications for the memory used (t_{RCD}) | 9.4.1.5/9-14 |
| | | DDR2 | Should be set according to the specifications for the memory used (t_{RCD}) | |
| CASLAT | CAS Latency | DDR1 | Should be set, along with the Extended CAS Latency, to the desired CAS latency | 9.4.1.5/9-14 |
| | | DDR2 | Should be set, along with the Extended CAS Latency, to the desired CAS latency | |
| REFREC | Refresh Recovery | DDR1 | Should be set, along with the Extended Refresh Recovery, to the specifications for the memory used (t_{RFC}) | 9.4.1.5/9-14 |
| | | DDR2 | Should be set, along with the Extended Refresh Recovery, to the specifications for the memory used (T_{RFC}) | |
| WRREC | Write Recovery | DDR1 | Should be set according to the specifications for the memory used (t_{WR}) | 9.4.1.5/9-14 |
| | | DDR2 | Should be set according to the specifications for the memory used (t_{WR}) | |
| ACTTOACT | Activate A to Activate B | DDR1 | Should be set according to the specifications for the memory used (t_{RRD}) | 9.4.1.5/9-14 |
| | | DDR2 | Should be set according to the specifications for the memory used (t_{RRD}) | |
| WRTORD | Write to Read Timing | DDR1 | Should be set according to the specifications for the memory used (t_{WTR}) | 9.4.1.5/9-14 |
| | | DDR2 | Should be set according to the specifications for the memory used (t_{WTR}) | |
| ADD_LAT | Additive Latency | DDR1 | Should be set to 000 | 9.4.1.6/9-16 |
| | | DDR2 | Should be set to the desired additive latency. This must be set to a value less than <code>TIMING_CFG_1[ACTTORW]</code> | |
| WR_LAT | Write Latency | DDR1 | Should be set to 001 | 9.4.1.6/9-16 |
| | | DDR2 | Should be set to CAS latency – 1 cycle. For example, if the CAS latency is 5 cycles, then this field should be set to 100 (4 cycles). | |

Table 9-40. Programming Differences between Memory Types (continued)

| Parameter | Description | Differences | | Section/page |
|-----------|-----------------------------|-------------|---|---------------|
| RD_TO_PRE | Read to Precharge Timing | DDR1 | Should be set to 010 if burst length is 4 and 100 if burst length is 8 | 9.4.1.6/9-16 |
| | | DDR2 | Should be set according to the specifications for the memory used (t_{RTP}). Time between read and precharge for non-zero value of additive latency (AL) is a minimum of $AL + t_{RTP}$ cycles. | |
| CKE_PLS | Minimum CKE Pulse Width | DDR1 | Can be set to 001 | 9.4.1.6/9-16 |
| | | DDR2 | Should be set according to the specifications for the memory used (t_{CKE}) | |
| FOUR_ACT | Four Activate Window | DDR1 | Should be set to 00001 | 9.4.1.6/9-16 |
| | | DDR2 | Should be set according to the specifications for the memory used (t_{FAW}). Only applies to eight logical banks. | |
| RD_EN | Registered DIMM Enable | DDR1 | If registered DRAM modules are used, then this field should be set to 1 | 9.4.1.7/9-18 |
| | | DDR2 | If registered DRAM modules are used, then this field should be set to 1 | |
| 8_BE | 8-beat burst enable | DDR1 | If 8-beat bursts are desired, then this field should be set to 1 | 9.4.1.7/9-18 |
| | | DDR2 | Should be set to 0 | |
| 2T_EN | 2T Timing Enable | DDR1 | In heavily loaded systems, this can be set to 1 to gain extra timing margin on the interface at the cost of address/command bandwidth. | 9.4.1.7/9-18 |
| | | DDR2 | In heavily loaded systems, this can be set to 1 to gain extra timing margin on the interface at the cost of address/command bandwidth. | |
| ODT_CFG | ODT Configuration | DDR1 | Should be set to 00 | 9.4.1.8/9-21 |
| | | DDR2 | Can be set for termination at the IOs according to system topology. Typically, if ODT is enabled, then the internal IOs should be set up for termination only during reads to DRAM. | |
| BSTOPR | Burst To Precharge Interval | DDR1 | Can be set to any value, depending on the application. Auto precharge can be enabled by setting this field to all 0s. | 9.4.1.12/9-27 |
| | | DDR2 | Can be set to any value, depending on the application. Auto precharge can be enabled by setting this field to all 0s. | |

9.6.2 DDR SDRAM Initialization Sequence

After configuration of all parameters is complete, system software must set `DDR_SDRAM_CFG[MEM_EN]` to enable the memory interface. Note that 200 μ s must elapse after DRAM clocks are stable (`DDR_SDRAM_CLK_CNTL[CLK_ADJUST]` is set and any chip select is

enabled) before MEM_EN can be set, so a delay loop in the initialization code may be necessary if software is enabling the memory controller. If DDR_SDRAM_CFG[BI] is not set, the DDR memory controller conducts an automatic initialization sequence to the memory, which follows the memory specifications. If the bypass initialization mode is used, then software can initialize the memory through the DDR_SDRAM_MD_CNTL register.

9.6.3 Using Forced Self-Refresh Mode to Implement a Battery-Backed RAM System

This section describes the options offered by this device to support battery-backed main memory.

9.6.3.1 Software Based Self-Refresh

The DDR controller also has a software-programmable bit, DDR_SDRAM_CFG_2[FRC_SR], that immediately puts main memory into self-refresh mode. See [Section 9.4.1.8, “DDR SDRAM Control Configuration 2 \(DDR_SDRAM_CFG_2\),”](#) for a description of this register.

It is expected that a critical interrupt routine triggered by an external voltage sensing device has time to set this bit.

9.6.3.2 Bypassing Re-initialization During Battery-Backed Operation

The DDR controller offers an initialization bypass feature (DDR_SDRAM_CFG[BI]), which system designers may use to prevent re-initialization of main memory during system power-on following an abnormal shutdown. See [Section 9.4.1.7, “DDR SDRAM Control Configuration \(DDR_SDRAM_CFG\),”](#) for information on this bit and [Section 9.4.1.15, “DDR Initialization Address \(DDR_INIT_ADDR\),”](#) for a discussion of avoiding possible ECC errors in this mode.

Note that the DDR controller automatically waits 200 DRAM cycles before issuing any command after the assertion of MCKE[0:1] when this mode is used.



Chapter 10

Enhanced Local Bus Controller

This chapter describes the enhanced local bus controller (eLBC) block. It describes the external signals and the memory-mapped registers as well as a functional description of the general-purpose chip-select machine (GPCM), NAND Flash control machine (FCM), and user-programmable machines (UPMs) of the eLBC. Finally, it includes an initialization and applications information section with many specific examples of its use.

10.1 Introduction

Figure 10-1 is a functional block diagram of the eLBC, which supports three interfaces: GPCM, FCM, and UPM controllers.




Figure 10-1. Enhanced Local Bus Controller Block Diagram

10.1.1 Overview

The main component of the eLBC is its memory controller, which provides a seamless interface to many types of memory devices and peripherals. The memory controller is responsible for controlling four memory banks shared by a GPCM, an FCM, and up to three UPMs. As such, it supports a minimal glue logic interface to SRAM, EPROM, NOR Flash EEPROM, NAND Flash EEPROM, burstable RAM, regular DRAM devices, extended data output DRAM devices, and other peripherals. The external address latch signal (LALE) allows multiplexing of addresses with data signals to reduce the device pin count. The eLBC also includes a number of data checking and protection features such as write protection, and a bus monitor to ensure that each bus cycle is terminated within a user-specified period.

10.1.2 Features

The eLBC main features are as follows:

- Memory controller with four memory banks
 - 32-bit address decoding with mask
 - Variable memory block sizes (64 Kbytes to 2 Gbytes in FCM mode, 32 Kbytes to 64 Mbytes in GPCM and UPM modes)
 - Selection of control signal generation on a per-bank basis
 - Data buffer controls activated on a per-bank basis
 - Automatic segmentation of large transactions into memory accesses optimized for bus width and addressing capability
 - Write-protection capability
- General-purpose chip-select machine (GPCM)
 - Compatible with SRAM, EPROM, NOR Flash EEPROM, and peripherals
 - Global (boot) chip-select available at system reset
 - Boot chip-select support for 8- and 16-bit devices
 - Minimum three-clock access to external devices
 - Two byte-write-enable signals ($\overline{\text{LWE}}[0:1]$)
 - Output enable signal ($\overline{\text{LOE}}$)
 - External access termination signal ($\overline{\text{LGTA}}$)
- NAND Flash control machine (FCM)
 - Compatible with small (512 + 16 bytes) and large (2048 + 64 bytes) page parallel NAND Flash EEPROM
 - Global (boot) chip-select available at system reset, with 4-Kbyte boot block buffer for execute-in-place boot loading
 - Boot chip-select support for 8-bit devices
 - Dual 2-Kbyte/eight 512-byte buffers allow simultaneous data transfer during flash reads and programming
 - Interrupt-driven block transfer for reads and writes
 - Programmable command and data transfer sequences of up to eight steps supported

- Generic command and address registers support proprietary flash interfaces
- Block write locking to ensure system security and integrity
- Three user-programmable machines (UPMs)
 - Programmable-array-based machine controls external signal timing with a granularity of up to one quarter of an external bus clock period
 - User-specified control-signal patterns run when an internal master requests a single-beat or burst read or write access.
 - UPM refresh timer runs a user-specified control signal pattern to support refresh
 - User-specified control-signal patterns can be initiated by software
 - Each UPM can be defined to support DRAM devices with depths of 64, 128, 256, and 512 Kbytes, and 1, 2, 4, 8, 16, 32, 64, 128, and 256 Mbytes
 - Support for 8- and 16-bit devices
 - Page mode support for successive transfers within a burst
 - Internal address multiplexing supporting 64-, 128-, 256-, and 512-Kbyte, and 1-, 2-, 4-, 8-, 16-, 32-, 64-, 128-, and 256-Mbyte page banks
- Optional monitoring of transfers between local bus internal masters and local bus slaves (local bus error reporting on interrupt and status registers)
- Different machines (FCM/GPCM/UPM) share the address, data, and control signals. While the eLBC is servicing a transaction, subsequent transactions are queued until the current transaction has completed.

10.1.3 Modes of Operation

The eLBC provides one GPCM, one FCM, and three UPMs for the local bus, with no restriction on how many of the four banks (chip selects) can be programmed to operate with any given machine. The internal transaction address is limited to 32 bits, so all chip selects must fall within the 4-Gbyte window addressed by the internal transaction address. When a memory transaction is dispatched to the eLBC, the internal transaction address is compared with the address information of each bank (chip select). The corresponding machine assigned to that bank (GPCM, FCM, or UPM) then takes ownership of the external signals that control the access and maintains control until the transaction ends. Thus, with the eLBC in GPCM or FCM, or UPM mode, only one of the four chip selects is active at any time for the duration of the transaction except in the case of UPM refresh where all UPM machines that are enabled for refresh have concurrent chip select assertion.

10.1.3.1 eLBC Bus Clock and Clock Ratios

The eLBC supports ratios of 2, 4, and 8 between the faster internal (csb) clock and slower external bus clock (LCLK[0:1]). This ratio is software programmable through the clock ratio register (LCRR[CLKDIV]). This ratio affects the resolution of signal timing shifts in GPCM and FCM modes and the interpretation of UPM array words in UPM mode. The bus clock is driven identically onto pins, LCLK[0:1], to allow the clock load to be shared equally across a set of signal nets, thereby enhancing the edge rates of the bus clock.

10.1.3.2 Source ID Debug Mode

The eLBC provides the ID of a transaction source on external device pins. When those pins are selected, the 5-bit internal ID of the current transaction source appears on LSRCID[0:4] whenever valid address or data is available on the eLBC external pins. The reserved value of 0x1F, which indicates invalid address or data, appears on the source ID pins at all other times. The combination of a valid source ID (any value except 0x1F) and the value of external address latch enable (LALE) and data valid (LDVAL) facilitate capturing useful debug data as follows:

- If a valid source ID is detected on LSRCID[0:4] and LALE is asserted, a valid full 26-bit address may be latched from LAD[0:15] and combined with LA[16:25].
- If a valid source ID is detected on LSRCID[0:4] and LDVAL is asserted, valid data may be latched from LAD.

The LSRCID[0:4] and LDVAL signals are multiplexed with other functions sharing the same external pins. Refer to [Chapter 2, “Signal Descriptions,”](#) and [Chapter 4, “Reset, Clocking, and Initialization,”](#) to learn how to enable the LSRCID/LDVAL pins.

10.2 External Signal Descriptions

[Table 10-1](#) contains a list of external signals related to the eLBC and summarizes their function. The table also shows the reset state of all external signals during assertion of $\overline{\text{HRESET}}$. For more information on the use of some of these signals as reset configuration signals on the device, see “Power on Reset Flow.”

Table 10-1. Signal Properties—Summary

| Name | Alternate Function(s) | Mode | Descriptions | No. of Signals | I/O | Reset State (Outputs) |
|--|--------------------------|------|-------------------------------|----------------|-----|-----------------------|
| LALE | — | — | External address latch enable | 1 | O | Reset_cfg |
| $\overline{\text{LCS}}[0:3]$ | — | — | Chip selects 0–3 | 4 | O | Reset_cfg |
| $\overline{\text{LWE}}0/$ $\overline{\text{LFW}}E/$ $\overline{\text{LBS}}0$ | $\overline{\text{LWE}}0$ | GPCM | Write enable 0 | 1 | O | Reset_cfg |
| | $\overline{\text{LFW}}E$ | FCM | Write enable | 1 | | |
| | $\overline{\text{LBS}}0$ | UPM | Byte (lane) select 0 | 1 | | |
| $\overline{\text{LWE}}1/$ $\overline{\text{LBS}}1$ | $\overline{\text{LWE}}1$ | GPCM | Write enable 1 | 1 | O | Reset_cfg |
| | $\overline{\text{LBS}}1$ | UPM | Byte (lane) select 1 | 1 | | |
| LGPL0/ LFCLE | LGPL0 | UPM | General purpose line 0 | 1 | O | Reset_cfg |
| | LFCLE | FCM | Flash command latch enable | 1 | | — |
| LGPL1/ LFALE | LGPL1 | UPM | General purpose line 1 | 1 | O | Reset_cfg |
| | LFALE | FCM | Flash address latch enable | 1 | | — |
| $\overline{\text{LOE}}/$ LGPL2/ $\overline{\text{LFRE}}$ | $\overline{\text{LOE}}$ | GPCM | Output enable | 1 | O | — |
| | $\overline{\text{LFRE}}$ | FCM | Flash read enable | 1 | | — |
| | LGPL2 | UPM | General purpose line 2 | 1 | | — |

Table 10-1. Signal Properties—Summary (continued)

| Name | Alternate Function(s) | Mode | Descriptions | No. of Signals | I/O | Reset State (Outputs) |
|--|--------------------------|------------|---|----------------|-----|-----------------------|
| $\overline{\text{LGPL3/}}/\overline{\text{LFWP}}$ | LGPL3 | UPM | General purpose line 3 | 1 | O | Reset_cfg |
| | $\overline{\text{LFWP}}$ | FCM | Flash write protect | 1 | | — |
| $\overline{\text{LGTA/}}/\overline{\text{LFRB/}}/\overline{\text{LGPL4/}}/\overline{\text{LUPWAIT}}$ | $\overline{\text{LGTA}}$ | GPCM | Transaction termination | 1 | I | High-Z |
| | $\overline{\text{LFRB}}$ | FCM | Flash ready/ $\overline{\text{busy}}$, open-drain shared pin | 1 | I | — |
| | LGPL4 | UPM | General purpose line 4 | 1 | O | — |
| | LUPWAIT | UPM | External device wait | 1 | I | — |
| LGPL5 | — | UPM | General purpose line 5 | 1 | O | Reset_cfg |
| LBCTL | — | — | Data buffer control | 1 | O | — |
| LA[16:25] | — | — | Non-multiplexed address bus | 10 | O | — |
| LAD[0:15] | — | — | Multiplexed address/data bus | 16 | I/O | — |
| LCLK[0:1] | — | — | Local bus clocks | 2 | O | Driven |
| LDVAL | — | eLBC debug | Local bus data valid | 1 | O | |
| LSRCID[0:4] | — | eLBC debug | Local bus source ID | 5 | O | |

Table 10-2 shows the detailed external signal descriptions for the eLBC.

Table 10-2. Enhanced Local Bus Controller Detailed Signal Descriptions

| Signal | I/O | Description | |
|------------------------------|-----|--|--|
| LALE | O | External address latch enable. The local bus memory controller provides control for an external address latch, which allows address and data to be multiplexed on the device pins. | |
| | | State Meaning | Asserted/Negated—LALE is asserted with the address at the beginning of each memory controller transaction. The number of cycles for which it is asserted is governed by the ORn[EAD] and LCRR[EADC] fields. Note that no other control signals are asserted during the assertion of LALE. |
| $\overline{\text{LCS}}[0:3]$ | O | Chip selects. Four chip selects are provided that are mutually exclusive. | |
| | | State Meaning | Asserted/Negated—Used to enable specific memory devices or peripherals connected to the eLBC. $\overline{\text{LCS}}[0:3]$ are provided on a per-bank basis with $\overline{\text{LCS}}0$ corresponding to the chip select for memory bank 0, which has the memory type and attributes defined by BR0 and OR0. |

Table 10-2. Enhanced Local Bus Controller Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|---|-----|--|--|
| $\overline{\text{LWE0}}$ / $\overline{\text{LFW E}}$ / LBS0 , $\overline{\text{LWE1}}$ / LBS1 | O | GPCM write enable 0/FCM write enable/UPM byte select 0. These signals select or validate each byte lane of the data bus. For an 8-bit port size, bit 0 is the only defined signal. The least-significant address bits of each access also determine which byte lanes are considered valid for a given data transfer. | |
| | | State Meaning | Asserted/Negated—For GPCM operation, $\overline{\text{LWE}}[0:1]$ assert for each byte lane enabled for writing. $\overline{\text{LFW E}}$ enables command, address, and data writes to NAND Flash EEPROMs controlled by FCM. $\text{LBS}[0:1]$ are programmable byte-select signals in UPM mode. See Section 10.4.4.4, “RAM Array,” for programming details about $\text{LBS}[0:1]$. |
| | | Timing | Assertion/Negation—See Section 10.4.2, “General-Purpose Chip-Select Machine (GPCM),” for details regarding the timing of $\overline{\text{LWE}}[0:1]$. |
| LGPL0/ LFCLE | O | General purpose line 0/FCM command latch enable. | |
| | | State Meaning | Asserted/Negated—In UPM mode, LGPL0 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode, LFCLE enables command cycles to NAND Flash EEPROMs. |
| LGPL1/ LFALE | O | General-purpose line 1/FCM address latch enable. | |
| | | State Meaning | Asserted/Negated—In UPM mode, LGPL1 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode, LFALE enables address cycles to NAND Flash EEPROMs. |
| $\overline{\text{LOE}}$ /LGPL2/ LFRE | O | GPCM output enable/General-purpose line 2/FCM read enable. | |
| | | State Meaning | Asserted/Negated—Controls the output buffer of memory when accessing memory/devices in GPCM mode. In UPM mode, LGPL2 is one of six general purpose signals; it is driven with a value programmed into the UPM array. LFRE enables data read cycles from NAND Flash EEPROMs controlled by FCM. |
| LGPL3/ LFWP | O | General-purpose line 3/FCM write protect. | |
| | | State Meaning | Asserted/Negated—In UPM mode, LGPL3 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode LFWP protects NAND Flash EEPROMs from accidental erasure and programming when LFWP is asserted low—see Section 10.3.1.16, “Flash Mode Register (FMR),” for programming of FCM operations to control LFWP. |
| $\overline{\text{LGT A}}$ /LGPL4/ LFRB/ LUPWAIT | I/O | GPCM transfer acknowledge/General-purpose line 4/FCM Flash ready-busy/UPM wait. | |
| | | State Meaning | Asserted/Negated—Input in GPCM or FCM modes used for transaction termination. It may also be configured as one of six general-purpose output signals when in UPM mode or as an input to force the UPM controller to wait for the memory/device. FCM uses LFRB to stall during long-latency read and programming operations, continuing once LFRB returns high. |
| LGPL5 | O | General-purpose line 5 | |
| | | State Meaning | Asserted/Negated—One of six general purpose signals when in UPM mode, and drives a value programmed in the UPM array. |

Table 10-2. Enhanced Local Bus Controller Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|-------------|-----|--|
| LBCTL | O | Data buffer control. The memory controller activates LBCTL for the local bus when a GPCM-, UPM-, or FCM-controlled bank is accessed. Buffer control is disabled by setting $OR\eta[BCTL\overline{D}]$. |
| | | State Meaning Asserted/Negated—The LBCTL pin normally functions as a write/ $\overline{\text{read}}$ control for a bus transceiver connected to the LAD lines. Note that an external data buffer must not drive the LAD lines in conflict with the eLBC when LBCTL is high, because LBCTL remains high after reset and during address phases. |
| LA[16:25] | O | Nonmultiplexed address bus. All bits driven on LA[16:25] are defined for 8-bit port sizes. For 16-bit port sizes LA[25] is a don't care. |
| | | State Meaning Asserted/Negated—LA is the address bus used to transmit addresses to external RAM devices. Refer to Section 10.5, "Initialization/Application Information," for address signal multiplexing. |
| LAD[0:15] | I/O | Multiplexed address/data bus. For a port size of 16 bits, LAD[0:7] connect to the most-significant byte lane (at address offset 0), while LAD[8:15] connect to the least-significant byte lane (at address offset 1). For a port size of 8 bits, only LAD[0:7] are connected to the external RAM. |
| | | State Meaning Asserted/Negated—LAD is the shared 16-bit address/data bus through which external RAM devices transfer data and receive addresses. |
| | | Timing Assertion/Negation—During assertion of LALE, LAD are driven with the RAM address for the access to follow. External logic should propagate the address on LAD while LALE is asserted, and latch the address upon negation of LALE. After LALE is negated, LAD are either driven by write data or are made high-impedance by the eLBC in order to sample read data driven by an external device. Following the last data transfer of a write access, LAD are again taken into a high-impedance state. |
| LCLK[0:1] | O | Local bus clocks |
| | | State Meaning Asserted/Negated—LCLK[0:1] drive an identical bus clock signal for distributed loads. |
| LDVAL | O | Local bus data valid (eLBC debug mode only) |
| | | State Meaning Asserted/Negated—For a read, LDVAL asserts for one bus cycle in the cycle immediately preceding the sampling of read data on LAD. For a write, LDVAL asserts for one bus cycle during the final cycle for which the current write data on LAD is valid. During burst transfers, LDVAL asserts for each data beat. |
| | | Timing Assertion/Negation—Valid only while the eLBC is in system debug mode. In debug mode, LDVAL asserts when the eLBC generates a data transfer acknowledge. |
| LSRCID[0:4] | O | Local bus source ID (eLBC debug mode only). In debug mode, all LSRCID[0:4] pins are driven high unless LSRCID[0:4] is driving a debug source ID for identifying the internal system device controlling the eLBC. |
| | | State Meaning Asserted/Negated—Remain high until the last bus cycle of the assertion of LALE, in which case the source ID of the address is indicated, or until LDVAL is asserted, in which case the source ID relating to the data transfer is indicated. In case of address debug, LSRCID[0:4] is valid only when the address on LAD consists of all physical address bits—with optional padding—for reconstructing the system address presented to the eLBC. |

10.3 Memory Map/Register Definition

Table 10-3 shows the memory-mapped registers of the eLBC. Undefined 4-byte address spaces within offset 0x000–0xFFFF are reserved.

Table 10-3. Enhanced Local Bus Controller Registers

| Enhanced Local Bus Controller—Block Base Address 0x0_5000 | | | | |
|---|---|--------|-------------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | BR0—Base register 0 | R/W | 0x0000_nnnn | 10.3.1.1/10-9 |
| 0x008 | BR1—Base register 1 | R/W | All zeros | 10.3.1.1/10-9 |
| 0x010 | BR2—Base register 2 | R/W | All zeros | 10.3.1.1/10-9 |
| 0x018 | BR3—Base register 3 | R/W | All zeros | 10.3.1.1/10-9 |
| 0x020–0x038 | Reserved | R/W | 0x0000_0000 | 10.3.1.1/10-9 |
| 0x004 | OR0—Options register 0 | R/W | 0x0000_0FF7 | 10.3.1.2/10-11 |
| 0x00C | OR1—Options register 1 | R/W | All zeros | 10.3.1.2/10-11 |
| 0x014 | OR2—Options register 2 | R/W | All zeros | 10.3.1.2/10-11 |
| 0x01C | OR3—Options register 3 | R/W | All zeros | 10.3.1.2/10-11 |
| 0x024–0x064 | Reserved | — | — | — |
| 0x068 | MAR—UPM address register | R/W | All zeros | 10.3.1.3/10-19 |
| 0x06C | Reserved | — | — | — |
| 0x070 | MAMR—UPMA mode register | R/W | All zeros | 10.3.1.4/10-20 |
| 0x074 | MBMR—UPMB mode register | R/W | All zeros | 10.3.1.4/10-20 |
| 0x078 | MCMR—UPMC mode register | R/W | All zeros | 10.3.1.4/10-20 |
| 0x07C–0x080 | Reserved | — | — | — |
| 0x084 | MRTPR—Memory refresh timer prescaler register | R/W | All zeros | 10.3.1.5/10-22 |
| 0x088 | MDR—UPM/FCM data register | R/W | All zeros | 10.3.1.6/10-22 |
| 0x08C | Reserved | — | — | — |
| 0x090 | LSOR—Special operation initiation register | R/W | All zeros | 10.3.1.7/10-23 |
| 0x094–0x09C | Reserved | — | — | — |
| 0x0A0 | LURT—UPM refresh timer | R/W | All zeros | 10.3.1.8/10-24 |
| 0x0A4–0x0AC | Reserved | — | — | — |
| 0x0B0 | LTESR—Transfer error status register | w1c | All zeros | 10.3.1.9/10-25 |
| 0x0B4 | LTEDR—Transfer error disable register | R/W | All zeros | 10.3.1.10/10-27 |
| 0x0B8 | LTEIR—Transfer error interrupt register | R/W | All zeros | 10.3.1.11/10-28 |
| 0x0BC | LTEATR—Transfer error attributes register | R/W | All zeros | 10.3.1.12/10-29 |
| 0x0C0 | LTEAR—Transfer error address register | R/W | All zeros | 10.3.1.13/10-30 |

Table 10-3. Enhanced Local Bus Controller Registers (continued)

| Enhanced Local Bus Controller—Block Base Address 0x0_5000 | | | | |
|---|-----------------------------------|--------|-------------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x0C4–0x0CC | Reserved | — | — | — |
| 0x0D0 | LBCR—Configuration register | R/W | 0x0004_0000 | 10.3.1.14/10-30 |
| 0x0D4 | LCRR—Clock ratio register | R/W | 0x8000_0008 | 10.3.1.15/10-32 |
| 0x0D8–0x0DC | Reserved | — | — | — |
| 0x0E0 | FMR—Flash mode register | R/W | 0x0000_0n00 | 10.3.1.16/10-33 |
| 0x0E4 | FIR—Flash instruction register | R/W | All zeros | 10.3.1.17/10-34 |
| 0x0E8 | FCR—Flash command register | R/W | All zeros | 10.3.1.18/10-35 |
| 0x0EC | FBAR—Flash block address register | R/W | All zeros | 10.3.1.19/10-36 |
| 0x0F0 | FPAR—Flash page address register | R/W | All zeros | 10.3.1.20/10-36 |
| 0x0F4 | FBCR—Flash byte count register | R/W | All zeros | 10.3.1.21/10-38 |
| 0x0F8–0x0FC | Reserved | — | — | — |

10.3.1 Register Descriptions

This section provides a detailed description of the eLBC configuration, status, and control registers with detailed bit and field descriptions.

Address offsets in the eLBC address range that are not defined in [Table 10-3](#) should not be accessed for reading or writing. Similarly, only zero should be written to reserved bits of defined registers, as writing ones can have unpredictable results in some cases.

Bits designated as write-one-to-clear are cleared only by writing ones to them. Writing zeros to them has no effect.


10.3.1.1 Base Registers (BR0–BR3)

The base registers (BR n), shown in [Figure 10-2](#), contain the base address and address types for each memory bank. The memory controller uses this information to compare the address bus value with the current address accessed. Each register (bank) includes a memory attribute and selects the machine for memory operation handling. Note that after system reset, BR0[V] is set, BR1[V]–BR3[V] are cleared, and

the value of BR0[PS] reflects the initial port size configured by the boot ROM location field of the reset configuration word.

Offset BR0: 0x000
 BR1: 0x008
 BR2: 0x010
 BR3: 0x018

Access: Read/Write



¹ BR0 has its valid bit (V) set for RCWH[ROMLOC] = LBC. Thus bank 0 is valid with the port size (PS) configured from RCWH[ROMLOC] as loaded during reset. M = 0 for MSEL of GPCM, 1 for MSEL of FCM at boot. All other base registers have all bits cleared to zero during reset.

Figure 10-2. Base Registers (BR_n)

Table 10-4 describes BR_n fields.

Table 10-4. BR_n Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–16 | BA | Base address. The upper 17 bits of each base register are compared to the address on the address bus to determine if the bus master is accessing a memory bank controlled by the memory controller. Used with the address mask bits OR _n [AM]. |
| 17–18 | — | Reserved |
| 19–20 | PS | Port size. Specifies the port size of this memory region. For BR0, PS is configured from the boot ROM location field in the reset configuration word as loaded during reset. For all other banks the value is reset to 00 (port size not defined). 00 Reserved 01 8-bit (supported for GPCM, UPM, FCM) 10 16-bit (supported for GPCM, UPM) 11 Reserved |
| 21–22 | DECC | Specifies the method for data error checking. 00 Data error checking disabled. No ECC generation for FCM. 01 ECC checking is enabled, but ECC generation is disabled, for FCM on full-page transfers. 10 ECC checking and generation are enabled for FCM on full-page transfers. 11 Reserved |
| 23 | WP | Write protect. 0 Read and write accesses are allowed. 1 Only read accesses are allowed. The memory controller does not assert \overline{LCSn} on write cycles to this memory bank. LTESR[WP] is set (if WP is set) if a write to this memory bank is attempted, and a local bus error interrupt is generated (if enabled), terminating the cycle. |

Table 10-4. BR_n Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 24–26 | MSEL | Machine select. Specifies the machine to use for handling memory operations. 000 GPCM (possible reset value) 001 FCM (possible reset value) 010 Reserved 011 Reserved 100 UPMA 101 UPMB 110 UPMC 111 Reserved |
| 27 | — | Reserved |
| 28–29 | — | Reserved |
| 30 | — | Reserved |
| 31 | V | Valid bit. Indicates that the contents of the BR _n and OR _n pair are valid. \overline{LCS}_n does not assert unless V is set (an access to a region that has no valid bit set may cause a bus time-out). After a system reset, only BR0[V] is set. 0 This bank is invalid. 1 This bank is valid. |

10.3.1.2 Option Registers (OR0–OR3)

The OR_n registers define the sizes of memory banks and access attributes. The OR_n attribute bits support the following three modes of operation as defined by BR_n[MSEL]:

- GPCM mode
- FCM mode
- UPM mode

The OR_n registers are interpreted differently depending on which of the three machine types is selected for that bank. Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options. Table 10-5 shows the reset values for OR0.

Table 10-5. Reset value of OR0 Register

| Boot Source | OR0 Reset Value |
|--------------------------------|-----------------|
| FCM (small page NAND Flash) | 0000_03AE |
| FCM (large page NAND Flash) | 0000_07AE |
| GPCM | 0000_0FF7 |
| eLBC not used as a boot source | 0000_0F07 |

10.3.1.2.1 Address Mask

The address mask field of the option registers (OR_n[AM]) masks up to 17 corresponding BR_n[BA] fields. The 15 LSBs of the 32-bit internal transaction address do not participate in bank address matching in selecting a bank for access. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a

resource to reside in more than one area of the address map. [Table 10-6](#) shows memory bank sizes from 32 Kbytes to 4 Gbytes. Memory block sizes vary from 32 Kbytes to 4 Gbytes in FCM mode, and 32 Kbytes to 64 Mbytes in GPCM and UPM modes.


Table 10-6. Memory Bank Sizes in Relation to Address Mask

| AM | Memory Bank Size |
|-----------------------|------------------|
| 0000_0000_0000_0000_0 | 4 Gbytes |
| 1000_0000_0000_0000_0 | 2 Gbytes |
| 1100_0000_0000_0000_0 | 1 Gbyte |
| 1110_0000_0000_0000_0 | 512 Mbytes |
| 1111_0000_0000_0000_0 | 256 Mbytes |
| 1111_1000_0000_0000_0 | 128 Mbytes |
| 1111_1100_0000_0000_0 | 64 Mbytes |
| 1111_1110_0000_0000_0 | 32 Mbytes |
| 1111_1111_0000_0000_0 | 16 Mbytes |
| 1111_1111_1000_0000_0 | 8 Mbytes |
| 1111_1111_1100_0000_0 | 4 Mbytes |
| 1111_1111_1110_0000_0 | 2 Mbytes |
| 1111_1111_1111_0000_0 | 1 Mbyte |
| 1111_1111_1111_1000_0 | 512 Kbytes |
| 1111_1111_1111_1100_0 | 256 Kbytes |
| 1111_1111_1111_1110_0 | 128 Kbytes |
| 1111_1111_1111_1111_0 | 64 Kbytes |
| 1111_1111_1111_1111_1 | 32 Kbytes |

10.3.1.2.2 Option Registers (OR_n)—GPCM Mode

Figure 10-3 shows the bit fields for OR_n when the corresponding BR_n[MSEL] selects the GPCM machine.

Offset OR0: 0x004 Access: Read/Write
 OR1: 0x00C
 OR2: 0x014
 OR3: 0x01C



¹ Refer to Table 10-5 for the OR0 reset value. All other option registers have all bits cleared.

Figure 10-3. Option Registers (OR_n) in GPCM Mode

Table 10-7 describes OR_n fields for GPCM mode.

Table 10-7. OR_n—GPCM Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | |
|---------------|-------|---|---------------|------|---------|---|---|--|---|---|--|--------|---|---|
| 0–16 | AM | GPCM address mask. Masks corresponding BR _n bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked and therefore don't care for address checking. 1 Corresponding address bits are used in the comparison between base and transaction addresses. | | | | | | | | | | | | |
| 17–18 | — | Reserved | | | | | | | | | | | | |
| 19 | BCTLD | Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank. | | | | | | | | | | | | |
| 20 | CSNT | Chip select negation time. Determines when \overline{LCSn} and \overline{LWE} are negated during an external memory write access handled by the GPCM, provided that ACS ≠ 00 (when ACS = 00, only \overline{LWE} is affected by the setting of CSNT). This helps meet address/data hold times for slow memories and peripherals. 0 \overline{LCSn} and \overline{LWE} are negated normally. 1 \overline{LCSn} and \overline{LWE} are negated earlier depending on the value of LCRR[CLKDIV]. | | | | | | | | | | | | |
| | | <table border="1"> <thead> <tr> <th>LCRR [CLKDIV]</th> <th>CSNT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>x</td> <td>0</td> <td>\overline{LCSn} and \overline{LWE} are negated normally.</td> </tr> <tr> <td>2</td> <td>1</td> <td>\overline{LCSn} and \overline{LWE} are negated normally.</td> </tr> <tr> <td>4 or 8</td> <td>1</td> <td>\overline{LCSn} and \overline{LWE} are negated one quarter bus clock cycle earlier.</td> </tr> </tbody> </table> | LCRR [CLKDIV] | CSNT | Meaning | x | 0 | \overline{LCSn} and \overline{LWE} are negated normally. | 2 | 1 | \overline{LCSn} and \overline{LWE} are negated normally. | 4 or 8 | 1 | \overline{LCSn} and \overline{LWE} are negated one quarter bus clock cycle earlier. |
| LCRR [CLKDIV] | CSNT | Meaning | | | | | | | | | | | | |
| x | 0 | \overline{LCSn} and \overline{LWE} are negated normally. | | | | | | | | | | | | |
| 2 | 1 | \overline{LCSn} and \overline{LWE} are negated normally. | | | | | | | | | | | | |
| 4 or 8 | 1 | \overline{LCSn} and \overline{LWE} are negated one quarter bus clock cycle earlier. | | | | | | | | | | | | |

Table 10-7. OR_n—GPCM Field Descriptions (continued)

| Bits | Name | Description | | | | | | | | | | | | | | | | | | |
|---------------|-------|---|---------------|-------|---------|---|----|---|----|-----------|---|----|---|----|---|--------|----|--|----|---|
| 21–22 | ACS | <p>Address to chip-select setup. Determines the delay of the \overline{LCSn} assertion relative to the address change when the external memory access is handled by the GPCM. At system reset, OR0[ACS] = 11.</p> <table border="1" style="margin-left: 40px;"> <thead> <tr> <th>LCRR [CLKDIV]</th> <th>Value</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td rowspan="2">x</td> <td>00</td> <td>\overline{LCSn} is output at the same time as the address lines. Note that this overrides the value of CSNT such that CSNT = 0.</td> </tr> <tr> <td>01</td> <td>Reserved.</td> </tr> <tr> <td rowspan="2">2</td> <td>10</td> <td>\overline{LCSn} is output one half bus clock cycle after the address lines.</td> </tr> <tr> <td>11</td> <td>\overline{LCSn} is output one half bus clock cycle after the address lines.</td> </tr> <tr> <td rowspan="2">4 or 8</td> <td>10</td> <td>\overline{LCSn} is output one quarter bus clock cycle after the address lines.</td> </tr> <tr> <td>11</td> <td>\overline{LCSn} is output one half bus clock cycle after the address lines.</td> </tr> </tbody> </table> | LCRR [CLKDIV] | Value | Meaning | x | 00 | \overline{LCSn} is output at the same time as the address lines. Note that this overrides the value of CSNT such that CSNT = 0. | 01 | Reserved. | 2 | 10 | \overline{LCSn} is output one half bus clock cycle after the address lines. | 11 | \overline{LCSn} is output one half bus clock cycle after the address lines. | 4 or 8 | 10 | \overline{LCSn} is output one quarter bus clock cycle after the address lines. | 11 | \overline{LCSn} is output one half bus clock cycle after the address lines. |
| LCRR [CLKDIV] | Value | Meaning | | | | | | | | | | | | | | | | | | |
| x | 00 | \overline{LCSn} is output at the same time as the address lines. Note that this overrides the value of CSNT such that CSNT = 0. | | | | | | | | | | | | | | | | | | |
| | 01 | Reserved. | | | | | | | | | | | | | | | | | | |
| 2 | 10 | \overline{LCSn} is output one half bus clock cycle after the address lines. | | | | | | | | | | | | | | | | | | |
| | 11 | \overline{LCSn} is output one half bus clock cycle after the address lines. | | | | | | | | | | | | | | | | | | |
| 4 or 8 | 10 | \overline{LCSn} is output one quarter bus clock cycle after the address lines. | | | | | | | | | | | | | | | | | | |
| | 11 | \overline{LCSn} is output one half bus clock cycle after the address lines. | | | | | | | | | | | | | | | | | | |
| 23 | XACS | <p>Extra address to chip-select setup. Setting this bit increases the delay of the \overline{LCSn} assertion relative to the address change when the external memory access is handled by the GPCM. After a system reset, OR0[XACS] = 1.</p> <p>0 Address to chip-select setup is determined by ORx[ACS] and LCRR[CLKDIV]. 1 Address to chip-select setup is extended (see Table 10-30 and Table 10-31).</p> | | | | | | | | | | | | | | | | | | |
| 24–27 | SCY | <p>Cycle length in bus clocks. Determines the number of wait states inserted in the bus cycle, when the GPCM handles the external memory access. Thus it is the main parameter for determining cycle length. The total cycle length depends on other timing attribute settings. After a system reset, OR0[SCY] = 1111.</p> <p>0000 No wait states 0001 1 bus clock cycle wait state ... 1111 15 bus clock cycle wait states</p> | | | | | | | | | | | | | | | | | | |
| 28 | SETA | <p>External address termination.</p> <p>0 Access is terminated internally by the memory controller unless the external device asserts \overline{LGTA} earlier to terminate the access. 1 Access is terminated externally by asserting the \overline{LGTA} external pin. (Only \overline{LGTA} can terminate the access).</p> | | | | | | | | | | | | | | | | | | |
| 29 | TRLX | <p>Timing relaxed. Modifies the settings of timing parameters for slow memories or peripherals.</p> <p>0 Normal timing is generated by the GPCM. 1 Relaxed timing on the following parameters:</p> <ul style="list-style-type: none"> • Adds an additional cycle between the address and control signals (only if ACS is not equal to 00). • Doubles the number of wait states specified by SCY, providing up to 30 wait states. • Works in conjunction with EHTR to extend hold time on read accesses. • \overline{LCSn} (only if ACS is not equal to 00) and \overline{LWE} signals are negated one cycle earlier during writes. | | | | | | | | | | | | | | | | | | |

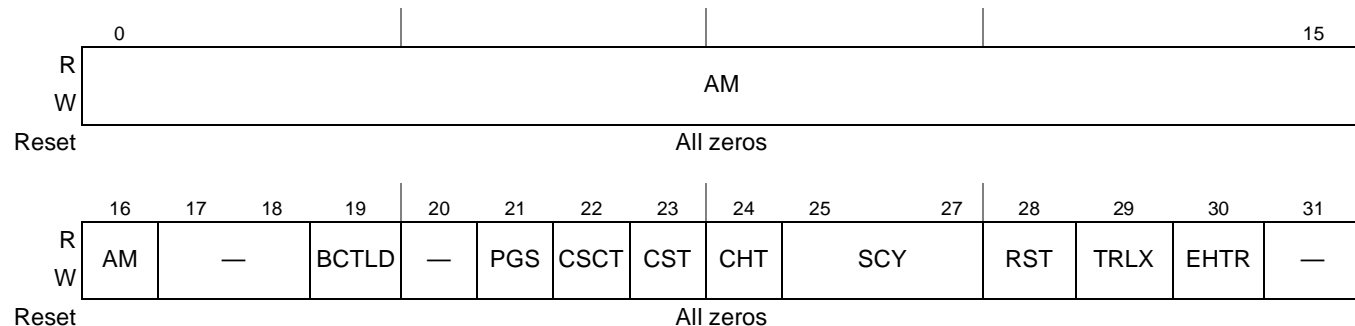
Table 10-7. OR_n—GPCM Field Descriptions (continued)

| Bits | Name | Description | | | | | | | | | | | | | | | |
|------|------|---|------|------|---------|---|---|---|---|---|---------------------------------|---|---|-----------------------------------|---|---|-----------------------------------|
| 30 | EHTR | Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>TRLX</th> <th>EHTR</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The memory controller generates normal timing. No additional cycles are inserted.</td> </tr> <tr> <td>0</td> <td>1</td> <td>1 idle clock cycle is inserted.</td> </tr> <tr> <td>1</td> <td>0</td> <td>4 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>1</td> <td>8 idle clock cycles are inserted.</td> </tr> </tbody> </table> | TRLX | EHTR | Meaning | 0 | 0 | The memory controller generates normal timing. No additional cycles are inserted. | 0 | 1 | 1 idle clock cycle is inserted. | 1 | 0 | 4 idle clock cycles are inserted. | 1 | 1 | 8 idle clock cycles are inserted. |
| TRLX | EHTR | Meaning | | | | | | | | | | | | | | | |
| 0 | 0 | The memory controller generates normal timing. No additional cycles are inserted. | | | | | | | | | | | | | | | |
| 0 | 1 | 1 idle clock cycle is inserted. | | | | | | | | | | | | | | | |
| 1 | 0 | 4 idle clock cycles are inserted. | | | | | | | | | | | | | | | |
| 1 | 1 | 8 idle clock cycles are inserted. | | | | | | | | | | | | | | | |
| 31 | EAD | External address latch delay. Allow extra bus clock cycles when using external address latch (LALE). 0 No additional bus clock cycles (LALE asserted for one bus clock cycle only) 1 Extra bus clock cycles are added (LALE is asserted for the number of bus clock cycles specified by LCRR[EADC]). | | | | | | | | | | | | | | | |

10.3.1.2.3 Option Registers (OR_n)—FCM Mode

Figure 10-4 shows the bit fields for OR_n when the corresponding BR_n[MSEL] selects the FCM machine.

Offset OR0: 0x004 Access: Read/Write
 OR1: 0x00C
 OR2: 0x014
 OR3: 0x01C



¹ Refer to Table 10-5 for the OR0 reset value. All other option registers have all bits cleared.

Figure 10-4. Option Registers (OR_n) in FCM Mode

Table 10-8 describes OR n fields for FCM mode.

Table 10-8. OR n —FCM Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | | | | |
|-------|-------|---|------|------|---------|---|---|---|---|---|---|---|---|--|---|---|---|
| 0–16 | AM | FCM address mask. Masks corresponding BR n bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked. 1 Corresponding address bits are used in the comparison between base and transaction addresses. | | | | | | | | | | | | | | | |
| 17–18 | — | Reserved | | | | | | | | | | | | | | | |
| 19 | BCTLD | Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank. | | | | | | | | | | | | | | | |
| 20 | — | Reserved | | | | | | | | | | | | | | | |
| 21 | PGS | NAND Flash EEPROM page size, buffer size, and block size. 0 Page size of 512 main area bytes plus 16 spare area bytes (small page devices); FCM RAM buffers are 1 Kbyte each; Flash block size of 16 Kbytes. 1 Page size of 2048 main area bytes plus 64 spare area bytes (large page devices); FCM RAM buffers are 4 Kbytes each; Flash block size of 128 Kbytes. | | | | | | | | | | | | | | | |
| 22 | CSCT | Chip select to command time. Determines how far in advance \overline{LCSn} is asserted prior to any bus activity during a NAND Flash access handled by the FCM. This helps meet chip-select setup times for slow memories. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>TRLX</th> <th>CSCT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The chip-select is asserted 1 clock cycle before any command.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The chip-select is asserted 4 clock cycles before any command.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The chip-select is asserted 2 clock cycles before any command.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The chip-select is asserted 8 clock cycles before any command.</td> </tr> </tbody> </table> | TRLX | CSCT | Meaning | 0 | 0 | The chip-select is asserted 1 clock cycle before any command. | 0 | 1 | The chip-select is asserted 4 clock cycles before any command. | 1 | 0 | The chip-select is asserted 2 clock cycles before any command. | 1 | 1 | The chip-select is asserted 8 clock cycles before any command. |
| TRLX | CSCT | Meaning | | | | | | | | | | | | | | | |
| 0 | 0 | The chip-select is asserted 1 clock cycle before any command. | | | | | | | | | | | | | | | |
| 0 | 1 | The chip-select is asserted 4 clock cycles before any command. | | | | | | | | | | | | | | | |
| 1 | 0 | The chip-select is asserted 2 clock cycles before any command. | | | | | | | | | | | | | | | |
| 1 | 1 | The chip-select is asserted 8 clock cycles before any command. | | | | | | | | | | | | | | | |
| 23 | CST | Command setup time. Determines the delay of \overline{LFWEn} assertion relative to the command, address, or data change when the external memory access is handled by the FCM. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>TRLX</th> <th>CST</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The write-enable is asserted coincident with any command.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The write-enable is asserted 0.25 clock cycles after any command, address, or data.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The write-enable is asserted 0.5 clock cycles after any command, address, or data.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The write-enable is asserted 1 clock cycle after any command, address, or data.</td> </tr> </tbody> </table> | TRLX | CST | Meaning | 0 | 0 | The write-enable is asserted coincident with any command. | 0 | 1 | The write-enable is asserted 0.25 clock cycles after any command, address, or data. | 1 | 0 | The write-enable is asserted 0.5 clock cycles after any command, address, or data. | 1 | 1 | The write-enable is asserted 1 clock cycle after any command, address, or data. |
| TRLX | CST | Meaning | | | | | | | | | | | | | | | |
| 0 | 0 | The write-enable is asserted coincident with any command. | | | | | | | | | | | | | | | |
| 0 | 1 | The write-enable is asserted 0.25 clock cycles after any command, address, or data. | | | | | | | | | | | | | | | |
| 1 | 0 | The write-enable is asserted 0.5 clock cycles after any command, address, or data. | | | | | | | | | | | | | | | |
| 1 | 1 | The write-enable is asserted 1 clock cycle after any command, address, or data. | | | | | | | | | | | | | | | |

Table 10-8. OR_n—FCM Field Descriptions (continued)

| Bits | Name | Description | | | | | | | | | | | | | | | |
|-------|------|---|------|-----|---------|---|---|---|---|---|--|---|---|---|---|---|---|
| 24 | CHT | <p>Command hold time. Determines the $\overline{\text{LFWE}}$ negation prior to the command, address, or data change when the external memory access is handled by the FCM.</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>CHT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The write-enable is negated 0.5 clock cycles before any command, address, or data change.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The write-enable is negated 1 clock cycle before any command, address, or data change.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The write-enable is negated 1.5 clock cycles before any command, address, or data change.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The write-enable is negated 2 clock cycles before any command, address, or data change.</td> </tr> </tbody> </table> | TRLX | CHT | Meaning | 0 | 0 | The write-enable is negated 0.5 clock cycles before any command, address, or data change. | 0 | 1 | The write-enable is negated 1 clock cycle before any command, address, or data change. | 1 | 0 | The write-enable is negated 1.5 clock cycles before any command, address, or data change. | 1 | 1 | The write-enable is negated 2 clock cycles before any command, address, or data change. |
| TRLX | CHT | Meaning | | | | | | | | | | | | | | | |
| 0 | 0 | The write-enable is negated 0.5 clock cycles before any command, address, or data change. | | | | | | | | | | | | | | | |
| 0 | 1 | The write-enable is negated 1 clock cycle before any command, address, or data change. | | | | | | | | | | | | | | | |
| 1 | 0 | The write-enable is negated 1.5 clock cycles before any command, address, or data change. | | | | | | | | | | | | | | | |
| 1 | 1 | The write-enable is negated 2 clock cycles before any command, address, or data change. | | | | | | | | | | | | | | | |
| 25–27 | SCY | <p>Cycle length in bus clocks. Determines the following:</p> <ul style="list-style-type: none"> the number of wait states inserted in command, address, or data transfer bus cycles, when the FCM handles the external memory access. Thus it is the main parameter for determining cycle length. The total cycle length depends on other timing attribute settings. the delay between command/address writes and data write cycles, or the delay between write cycles and read cycles from NAND Flash EEPROM. A delay of $4 \times (2 + \text{SCY})$ clock cycles ($\text{TRLX} = 0$) or $8 \times (2 + \text{SCY})$ clock cycles ($\text{TRLX} = 1$) is inserted between the last write and the first data transfer to/from NAND Flash devices. the delay between a command write and the first sample point of the $\text{RDY}/\overline{\text{BSY}}$ pin (connected to $\overline{\text{LFRB}}$). $\overline{\text{LFRB}}$ is not sampled until $8 \times (2 + \text{SCY})$ clock cycles ($\text{TRLX} = 0$) or $16 \times (2 + \text{SCY})$ clock cycles ($\text{TRLX} = 1$) have elapsed following the command. <p>000 No extra wait states 001 1 bus clock cycle wait state ... 111 7 bus clock cycle wait states</p> | | | | | | | | | | | | | | | |
| 28 | RST | <p>Read setup time. Determines the delay of $\overline{\text{LFRE}}$ assertion relative to sampling of read data when the external memory access is handled by the FCM.</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>RST</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The read-enable is asserted 0.75 clock cycles prior to any wait states.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The read-enable is asserted 1 clock cycle prior to any wait states.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The read-enable is asserted 0.5 clock cycles prior to any wait states.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The read-enable is asserted 1 clock cycle prior to any wait states.</td> </tr> </tbody> </table> | TRLX | RST | Meaning | 0 | 0 | The read-enable is asserted 0.75 clock cycles prior to any wait states. | 0 | 1 | The read-enable is asserted 1 clock cycle prior to any wait states. | 1 | 0 | The read-enable is asserted 0.5 clock cycles prior to any wait states. | 1 | 1 | The read-enable is asserted 1 clock cycle prior to any wait states. |
| TRLX | RST | Meaning | | | | | | | | | | | | | | | |
| 0 | 0 | The read-enable is asserted 0.75 clock cycles prior to any wait states. | | | | | | | | | | | | | | | |
| 0 | 1 | The read-enable is asserted 1 clock cycle prior to any wait states. | | | | | | | | | | | | | | | |
| 1 | 0 | The read-enable is asserted 0.5 clock cycles prior to any wait states. | | | | | | | | | | | | | | | |
| 1 | 1 | The read-enable is asserted 1 clock cycle prior to any wait states. | | | | | | | | | | | | | | | |
| 29 | TRLX | <p>Timing relaxed. Modifies the settings of timing parameters for slow memories.</p> <p>0 Normal timing is generated by the FCM.</p> <p>1 Relaxed timing on the following parameters:</p> <ul style="list-style-type: none"> Doubles the number of clock cycles between $\overline{\text{LCSn}}$ assertion and commands. Doubles the number of wait states specified by SCY, providing up to 14 wait states. Works in conjunction with CST and RST to extend command/address/data setup times. Adds one clock cycle to the command/address/data hold times. Works in conjunction with CBT to extend the wait time for read/busy status sampling by 16 clock cycles. Works in conjunction with EHTR to double hold time on read accesses. | | | | | | | | | | | | | | | |

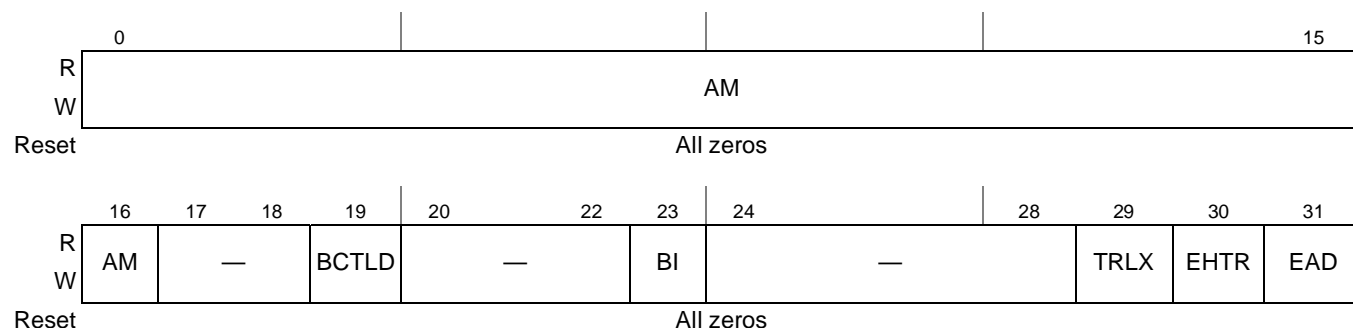
Table 10-8. OR_n—FCM Field Descriptions (continued)

| Bits | Name | Description | | | | | | | | | | | | | | | |
|------|------|---|------|-----------------------------------|---------|---|---|---------------------------------|---|---|-----------------------------------|---|---|-----------------------------------|---|---|-----------------------------------|
| 30 | EHTR | Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. | | | | | | | | | | | | | | | |
| | | <table border="1"> <thead> <tr> <th>TRLX</th> <th>EHTR</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>1 idle clock cycle is inserted.</td> </tr> <tr> <td>0</td> <td>1</td> <td>2 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>0</td> <td>4 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>1</td> <td>8 idle clock cycles are inserted.</td> </tr> </tbody> </table> | TRLX | EHTR | Meaning | 0 | 0 | 1 idle clock cycle is inserted. | 0 | 1 | 2 idle clock cycles are inserted. | 1 | 0 | 4 idle clock cycles are inserted. | 1 | 1 | 8 idle clock cycles are inserted. |
| | | TRLX | EHTR | Meaning | | | | | | | | | | | | | |
| | | 0 | 0 | 1 idle clock cycle is inserted. | | | | | | | | | | | | | |
| | | 0 | 1 | 2 idle clock cycles are inserted. | | | | | | | | | | | | | |
| 1 | 0 | 4 idle clock cycles are inserted. | | | | | | | | | | | | | | | |
| 1 | 1 | 8 idle clock cycles are inserted. | | | | | | | | | | | | | | | |
| 31 | — | Reserved | | | | | | | | | | | | | | | |

10.3.1.2.4 Option Registers (OR_n)—UPM Mode

Figure 10-5 shows the bit fields for OR_n when the corresponding BR_n[MSEL] selects a UPM machine.

Offset OR0: 0x004 Access: Read/Write
 OR1: 0x00C
 OR2: 0x014
 OR3: 0x01C



¹ Refer to Table 10-5 for the OR0 reset value. All other option registers have all bits cleared.

Figure 10-5. Option Registers (OR_n) in UPM Mode

Table 10-9 describes BR_n fields for UPM mode.

Table 10-9. OR_n—UPM Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0–16 | AM | UPM address mask. Masks corresponding BR _n bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked. 1 The corresponding address bits are used in the comparison with address pins. |
| 17–18 | — | Reserved |
| 19 | BCTLD | Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank. |

Table 10-9. OR_n—UPM Field Descriptions (continued)

| Bits | Name | Description | | | | | | | | | | | | | | | |
|-------|------|---|------|------|---------|---|---|---|---|---|---------------------------------|---|---|-----------------------------------|---|---|-----------------------------------|
| 20–22 | — | Reserved | | | | | | | | | | | | | | | |
| 23 | BI | Burst inhibit. Indicates if this memory bank supports burst accesses. 0 The bank supports burst accesses. 1 The bank does not support burst accesses. The selected UPM executes burst accesses as a series of single accesses. | | | | | | | | | | | | | | | |
| 24–28 | — | Reserved | | | | | | | | | | | | | | | |
| 29 | TRLX | Timing relaxed. Works in conjunction with EHTR to extend hold time on read accesses. | | | | | | | | | | | | | | | |
| 30 | EHTR | Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>TRLX</th> <th>EHTR</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The memory controller generates normal timing. No additional cycles are inserted.</td> </tr> <tr> <td>0</td> <td>1</td> <td>1 idle clock cycle is inserted.</td> </tr> <tr> <td>1</td> <td>0</td> <td>4 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>1</td> <td>8 idle clock cycles are inserted.</td> </tr> </tbody> </table> | TRLX | EHTR | Meaning | 0 | 0 | The memory controller generates normal timing. No additional cycles are inserted. | 0 | 1 | 1 idle clock cycle is inserted. | 1 | 0 | 4 idle clock cycles are inserted. | 1 | 1 | 8 idle clock cycles are inserted. |
| TRLX | EHTR | Meaning | | | | | | | | | | | | | | | |
| 0 | 0 | The memory controller generates normal timing. No additional cycles are inserted. | | | | | | | | | | | | | | | |
| 0 | 1 | 1 idle clock cycle is inserted. | | | | | | | | | | | | | | | |
| 1 | 0 | 4 idle clock cycles are inserted. | | | | | | | | | | | | | | | |
| 1 | 1 | 8 idle clock cycles are inserted. | | | | | | | | | | | | | | | |
| 31 | EAD | External address latch delay. Allow extra bus clock cycles when using external address latch (LALE). 0 No additional bus clock cycles (LALE asserted for one bus clock cycle only) 1 Extra bus clock cycles are added (LALE is asserted for the number of bus clock cycles specified by LCRR[EADC]). | | | | | | | | | | | | | | | |

10.3.1.3 UPM Memory Address Register (MAR)

Figure 10-6 shows the fields of the UPM memory address register (MAR).



Figure 10-6. UPM Memory Address Register (MAR)

Table 10-10 describes the MAR fields.

Table 10-10. MAR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–31 | A | Address that can be output to the address signals under control of the AMX bits in the UPM RAM word. |

10.3.1.4 UPM Mode Registers (MxMR)

The UPM machine mode registers (MAMR, MBMR, and MCMR), shown in [Figure 10-7](#), contain the configuration for the three UPMs.

Offset MAMR: 0x070
 MBMR: 0x074
 MCMR: 0x078

Access: Read/Write




Figure 10-7. UPM Mode Registers (MxMR)

[Table 10-11](#) describes UPM mode fields.

Table 10-11. MxMR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0 | — | Reserved |
| 1 | RFEN | Refresh enable. Indicates that the UPM needs refresh services. This bit must be set for UPMA (refresh executor) if refresh services are required on any UPM assigned chip selects. If MAMR[RFEN] = 0, no refresh services can be provided, even if UPMB and/or UPMC have their RFEN bit set. 0 Refresh services are not required 1 Refresh services are required |
| 2–3 | OP | Command opcode. Determines the command executed by the UPM _n when a memory access hits a UPM assigned bank. 00 Normal operation 01 Write to UPM array. On the next memory access that hits a UPM assigned bank, write the contents of the MDR into the RAM location pointed to by MAD. After the access, MAD is automatically incremented. 10 Read from UPM array. On the next memory access that hits a UPM assigned bank, read the contents of the RAM location pointed to by MAD into the MDR. After the access, MAD is automatically incremented. 11 Run pattern. On the next memory access that hits a UPM assigned bank, run the pattern written in the RAM array. The pattern run starts at the location pointed to by MAD and continues until the LAST bit is set in the RAM word. |
| 4 | UWPL | LUPWAIT polarity active low. Sets the polarity of the LUPWAIT pin when in UPM mode. 0 LUPWAIT is active high. 1 LUPWAIT is active low. |

Table 10-11. MxMR Field Descriptions (continued)

| Bits | Name | Description | | | | | | | | | | | | | | |
|-------|----------------------------|--|-----------|----------------------------|---------------------------------|--|-----------|-----------|---|----------------|------|------|---|-----------------|------|------|
| 5–7 | AM | <p>Address multiplex size. Determines how the address of the current memory cycle can be output on the address pins. This field is needed when interfacing with devices requiring row and column addresses multiplexed on the same pins. See Section 10.4.4.4.7, “Address Multiplexing (AMX),” for more information.</p> <p>000 Internal transaction address a[8:23] driven on [10:25]; LAD[0:15] driven low. 001 Internal transaction address a[7:22] driven on [10:25]; LAD[0:15] driven low. 010 Internal transaction address a[6:21] driven on [10:25]; LAD[0:15] driven low. 011 Internal transaction address a[5:20] driven on [10:25]; LAD[0:15] driven low. 100 Internal transaction address a[4:19] driven on [10:25]; LAD[0:15] driven low. 101 Internal transaction address a[3:18] driven on [10:25]; LAD[0:15] driven low. 110 Reserved 111 Reserved</p> | | | | | | | | | | | | | | |
| 8–9 | DS | <p>Disable timer period. Guarantees a minimum time between accesses to the same memory bank controlled by UPMn. The disable timer is turned on by the TODT bit in the RAM array word, and when expired, the UPMn allows the machine access to handle a memory pattern to the same bank. Accesses to a different bank by the same UPMn is also allowed. To avoid conflicts between successive accesses to different banks, the minimum pattern in the RAM array for a request serviced, should not be shorter than the period established by DS.</p> <p>00 1-bus clock cycle disable period 01 2-bus clock cycle disable period 10 3-bus clock cycle disable period 11 4-bus clock cycle disable period</p> | | | | | | | | | | | | | | |
| 10–12 | G0CL | <p>General line 0 control. Determines which logical address line can be output to the LGPL0 pin when the UPMn is selected to control the memory access.</p> <p>000 A12 001 A11 010 A10 011 A9 100 A8 101 A7 110 A6 111 A5</p> | | | | | | | | | | | | | | |
| 13 | GPL4 | <p>LGPL4 output line disable. Determines how the LGPL4/LUPWAIT pin is controlled by the corresponding bits in the UPMn array. See Table 10-38.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th rowspan="2">Value</th> <th rowspan="2">LGPL4/LUPWAIT Pin Function</th> <th colspan="2">Interpretation of UPM Word Bits</th> </tr> <tr> <th>G4T1/DLT3</th> <th>G4T3/WAEN</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>LGPL4 (output)</td> <td>G4T1</td> <td>G4T3</td> </tr> <tr> <td>1</td> <td>LUPWAIT (input)</td> <td>DLT3</td> <td>WAEN</td> </tr> </tbody> </table> | Value | LGPL4/LUPWAIT Pin Function | Interpretation of UPM Word Bits | | G4T1/DLT3 | G4T3/WAEN | 0 | LGPL4 (output) | G4T1 | G4T3 | 1 | LUPWAIT (input) | DLT3 | WAEN |
| Value | LGPL4/LUPWAIT Pin Function | Interpretation of UPM Word Bits | | | | | | | | | | | | | | |
| | | G4T1/DLT3 | G4T3/WAEN | | | | | | | | | | | | | |
| 0 | LGPL4 (output) | G4T1 | G4T3 | | | | | | | | | | | | | |
| 1 | LUPWAIT (input) | DLT3 | WAEN | | | | | | | | | | | | | |
| 14–17 | RLF | <p>Read loop field. Determines the number of times a loop defined in the UPMn will be executed for a burst- or single-beat read pattern or when MxMR[OP] = 11 (RUN command)</p> <p>0000 16 0001 1 0010 2 0011 3 ... 1110 14 1111 15</p> | | | | | | | | | | | | | | |

Table 10-11. MxMR Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 18–21 | WLF | Write loop field. Determines the number of times a loop defined in the UPM n will be executed for a burst- or single-beat write pattern. 0000 16 0001 1 0010 2 0011 3 ... 1110 14 1111 15 |
| 22–25 | TLF | Refresh loop field. Determines the number of times a loop defined in the UPM n will be executed for a refresh service pattern. 0000 16 0001 1 0010 2 0011 3 ... 1110 14 1111 15 |
| 26–31 | MAD | Machine address. RAM address pointer for the command executed. This field is incremented by 1, each time the UPM is accessed and the OP field is set to WRITE or READ. Address range is 64 words per UPM n . |

10.3.1.5 Memory Refresh Timer Prescaler Register (MRTPR)

The refresh timer prescaler register (MRTPR), shown in [Figure 10-8](#), is used to divide the csbclk (or twice csb_clk (if RCWL[LBCM] is set) to provide the UPM refresh timers clock.




Figure 10-8. Memory Refresh Timer Prescaler Register (MRTPR)

[Table 10-12](#) describes MRTPR fields.

Table 10-12. MRTPR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–7 | PTP | Refresh timers prescaler. Determines the period of the refresh timers input clock. The csb clock (or twice csb_clk (if RCWL[LBCM] is set) is divided by PTP except when the value is 00000_0000, which represents the maximum divider of 256. |
| 8–31 | — | Reserved |

10.3.1.6 UPM/FCM Data Register (MDR)

The memory data register (MDR), shown in [Figure 10-9](#) and [Figure 10-10](#), contains data written to or read from the RAM array for UPM read or write commands. MDR also contains data written to or read from

an external NAND Flash EEPROM for FCM write address, write data, and read status commands. MDR must be set up before issuing a write command to the UPM, or before issuing a FCM operation sequence that uses MDR to source address or data bytes.




Figure 10-9. UPM Data Register in UPM Mode (MDR)



Figure 10-10. FCM Data Register in FCM Mode (MDR)

Table 10-13 describes MDR[D].

Table 10-13. MDR Field Description

| Bits | Name | Description |
|-------|------|---|
| 0–31 | D | In UPM mode, D is the data to be read or written into the RAM array when a write or read command is supplied to the UPM (MxMR[OP] = 01 or MxMR[OP] = 10). |
| 0–7 | AS3 | In FCM mode, AS3 is the fourth byte of address sent by a custom address write operation, or the fourth byte of data read from a read status operation. |
| 8–15 | AS2 | In FCM mode, AS2 is the third byte of address sent by a custom address write operation, or the third byte of data read from a read status operation. |
| 16–23 | AS1 | In FCM mode, AS1 is the second byte of address sent by a custom address write operation, or the second byte of data read from a read status operation. |
| 24–31 | AS0 | In FCM mode, AS0 is the first byte of address sent by a custom address write operation, or the first byte of data read from a read status operation. |

10.3.1.7 Special Operation Initiation Register (LSOR)

The special operation initiation register (LSOR), shown in Figure 10-11, is used by software to trigger a special operation on the indicated bank. Writing to LSOR activates a special operation on bank LSOR[BANK] provided that the bank is valid and controlled by a memory controller whose mode OP field is set to a value other than “normal operation.” If eLBC is currently busy with a memory transaction, writing LSOR completes immediately, but the special operation request is queued until eLBC can service it. To avoid race conditions between software and a busy eLBC, registers that affect currently running special operation and LSOR must not be rewritten before a pending special operation has been completed. The UPM and FCM have different indications of when such special operations are completed. The

behavior of eLBC is unpredictable if special operation modes are altered between LSOR being written and the relevant memory controller completing that access.

UPM special operation modes are set in registers MxMR[OP], see [Section 10.3.1.4, “UPM Mode Registers \(MxMR\).”](#) FCM special operation modes are set in FMR[OP], see [Section 10.3.1.16, “Flash Mode Register \(FMR\).”](#) Writing LSOR has the same effect as setting a special controller mode and performing a dummy access to a bank associated with the controller in question, but use of LSOR avoids changing settings for the address space occupied by the bank. More details of special operation sequences appear in [Section 10.4.4.2.1, “UPM Programming Example \(Two Sequential Writes to the RAM Array\).”](#)




Figure 10-11. Special Operation Initiation Register (LSOR)

Table 10-14 describes LSOR.

Table 10-14. LSOR Field Description

| Bits | Name | Description |
|-------|------|---|
| 0–28 | — | Reserved |
| 29–31 | BANK | Bank on which a special operation is initiated. If the bank identified by BANK is marked valid (BRn[V] set) and the bank is controlled by a memory controller whose current mode OP is non-zero—or a special operation—eLBC will request the special operation to be activated on the selected bank when this field is written. Otherwise, writing this field has no effect. 000 Bank 0 is triggered for special operation ... 011 Bank 3 is triggered for special operation 100–111 Reserved |

10.3.1.8 UPM Refresh Timer (LURT)

The UPM refresh timer (LURT), shown in [Figure 10-12](#), generates a refresh request for all valid banks that selected a UPM machine and are refresh-enabled (MxMR[RFEN] = 1). Each time the timer expires, a qualified bank generates a refresh request using the selected UPM. The qualified banks rotate their requests.




Figure 10-12. UPM Refresh Timer (LURT)

Table 10-16 describes LTESR fields.

Table 10-16. LTESR Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0 | BM | Bus monitor time-out 0 No local bus monitor time-out occurred. 1 Local bus monitor time-out occurred. No data beat was acknowledged on the bus within $LBCR[BMT] \times LBCR[BMTPS]$ bus clock cycles from the start of a transaction. |
| 1 | FCT | FCM command time-out 0 No FCM command time-out occurred. 1 A CW0, CW1, CW2, or CW3 command issued to FCM timed-out with respect to the timer configured by FMR[CWTO]. |
| 2 | PAR | ECC error for FCM mode 0 No local bus ECC error 1 U ECC error (FCM). LTEATR[PB] indicates the block that caused the error and LTEATR[BNK] indicates which memory controller bank was accessed. |
| 3–4 | — | Reserved |
| 5 | WP | Write protect error 0 No write protect error occurred. 1 A write was attempted to a local bus memory region that was defined as read-only in the memory controller. Usually, in this case, a bus monitor time-out will occur (as the cycle is not automatically terminated). |
| 6–11 | — | Reserved |
| 12 | CS | Chip select error 0 No chip select error occurred. 1 A transaction was sent to the eLBC that did not hit any memory bank. |
| 13–30 | — | Reserved |
| 31 | CC | FCM command completion event 0 No FCM operation in progress, or operation pending. 1 FCM operation has completed, allowing software to continue processing of results. |

10.3.1.10 Transfer Error Check Disable Register (LTEDR)

The transfer error check disable register (LTEDR), shown in [Figure 10-14](#), is used to disable error/event checking. Note that control of error/event checking is independent of control of reporting of errors/events (LTEIR) through the interrupt mechanism.

Offset 0x0B4

Access: Read/Write




Figure 10-14. Transfer Error Check Disable Register (LTEDR)

[Table 10-17](#) describes LTEDR fields.

Table 10-17. LTEDR Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0 | BMD | Bus monitor disable 0 Bus monitor is enabled. 1 Bus monitor is disabled, but internal bus time-outs can still occur. |
| 1 | FCTD | FCM command time-out disable 0 FCM command timer is enabled. 1 FCM command time-out is disabled, but internal FCM command timer can terminate command waits. |
| 2 | PARD | ECC error checking disabled. 0 ECC error checking is enabled. 1 ECC error checking is disabled. |
| 3–4 | — | Reserved |
| 5 | WPD | Write protect error checking disable. 0 Write protect error checking is enabled. 1 Write protect error checking is disabled. |
| 6–11 | — | Reserved |
| 12 | CSD | Chip select error checking disable. 0 Chip select error checking is enabled. 1 Chip select error checking is disabled. |
| 13–30 | — | Reserved |
| 31 | CCD | FCM command completion checking disable. 0 Command completion checking is enabled. 1 Command completion checking is disabled. |

10.3.1.11 Transfer Error Interrupt Enable Register (LTEIR)

The transfer error interrupt enable register (LTEIR), shown in [Figure 10-15](#), is used to send or block error/event reporting through the eLBC internal interrupt mechanism. Software should clear pending errors/events in LTESR before enabling interrupts. After an interrupt has occurred, clearing relevant LTESR error/event bits negates the interrupt.




Figure 10-15. Transfer Error Interrupt Enable Register (LTEIR)

[Table 10-18](#) describes LTEIR fields.

Table 10-18. LTEIR Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0 | BMI | Bus monitor error interrupt enable. 0 Bus monitor error reporting is disabled. 1 Bus monitor error reporting is enabled. |
| 1 | FCTI | FCM command time-out interrupt enable. 0 FCM command time-out error reporting is disabled. 1 FCM command time-out error reporting is enabled. |
| 2 | PARI | ECC error interrupt enable. 0 ECC error reporting is disabled. 1 ECC error reporting is enabled. |
| 3–4 | — | Reserved |
| 5 | WPI | Write protect error interrupt enable. 0 Write protect error reporting is disabled. 1 Write protect error reporting is enabled. |
| 6–11 | — | Reserved |
| 12 | CSI | Chip select error interrupt enable. 0 Chip select error reporting is disabled. 1 Chip select error reporting is enabled. |
| 13–30 | — | Reserved |
| 31 | CCI | FCM command completion Event interrupt enable. 0 Command completion reporting is disabled. 1 Command completion reporting is enabled. |

10.3.1.12 Transfer Error Attributes Register (LTEATR)

The transfer error attributes register (LTEATR) captures source attributes of an error/event. [Figure 10-16](#) shows the LTEATR. After LTEATR[V] has been set, software must clear this bit to allow LTESR, LTEATR, and LTEAR to update following any subsequent events/errors.

NOTE

LTEATR may not capture accurate information for errors that occur when an FCM special operation is in progress.




Figure 10-16. Transfer Error Attributes Register (LTEATR)

[Table 10-19](#) describes LTEATR fields.

Table 10-19. LTEATR Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0–2 | — | Reserved |
| 3 | RWB | Transaction type for the error: 0 The transaction for the error was a write transaction. 1 The transaction for the error was a read transaction. |
| 4–10 | — | Reserved |
| 11–15 | SRCID | Captures the source of the transaction when this information is provided on the internal interface to the eLBC. The coding of the source ID debug information is the same as the coding of AEATR[MSTR_ID] (see Section 6.2.7, "Arbiter Event Attributes Register (AEATR).") |
| 16–19 | PB | Error on block for FCM. For FCM, there are at most four 512-byte page blocks (for a large page device) checked by ECC. A bit is set for the 512-byte block that had a non-correctable ECC error on read (bit 16 represents block 0, the first 512 bytes of a page; if ORx[PGS] = 0, bits 17–19 are always 0). |
| 20–23 | BNK | Memory controller bank. There is one error status bit per memory controller bank (bit 20 represents bank 0). A bit is set for the local bus memory controller bank that had an error. |
| 24–30 | — | Reserved |
| 31 | V | Error attribute capture is valid. Indicates that the captured error information is valid. 0 Captured error attributes and address are not valid. 1 Captured error attributes and address are valid. |

Table 10-21 describes LBCR fields.

Table 10-21. LBCR Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0 | LDIS | Local bus disable 0 Local bus is enabled. 1 Local bus is disabled. No internal transactions will be acknowledged. |
| 1–7 | — | Reserved |
| 8–9 | BCTLC | Defines the use of LBCTL 00 LBCTL is used as $\overline{W/R}$ control for GPCM or UPM accesses (buffer control). 01 LBCTL is used as \overline{LOE} for GPCM accesses only. 10 LBCTL is used as \overline{LWE} for GPCM accesses only. 11 Reserved |
| 10 | AHD | Address hold disable. Removes part of the hold time for LAD with respect to LALE in order to lengthen the LALE pulse. 0 During address phases on the local bus, the LALE signal negates one platform clock period prior to the address being invalidated. For instance, at 133 MHz, this provides 15 ns of additional address hold time at the external address latch. Running the csb at a lower frequency improves the hold time. 1 During address phases on the local bus, the LALE signal negates 0.5 platform clock period prior to the address being invalidated. This halves the address hold time, but extends the latch enable duration. This may be necessary for very high frequency designs. |
| 11–15 | — | Reserved. Reads to bit 13 return 1. |
| 16–23 | BMT | Bus monitor timing. Defines the bus monitor time-out period. Clearing BMT (reset value) selects the maximum count of bus clock cycles. For non-zero values of BMT, the number of LCLK clock cycles to count down before a time-out error is generated is given by: bus cycles = BMT × PS, where PS is set according to LBCR[BMTPS]. The value of BMT × PS must not be less than 40 bus cycles for reliable operation. |
| 24–27 | — | Reserved |
| 28–31 | BMTPS | Bus monitor timer prescale. Defines the multiplier, PS, to scale LBCR[BMT] for determining bus time-outs. 0000 PS = 8 0001 PS = 16 0010 PS = 32 0011 PS = 64 0100 PS = 128 0101 PS = 256 0110 PS = 512 0111 PS = 1024 1000 PS = 2048 1001 PS = 4096 1010 PS = 8192 1011 PS = 16,384 1100 PS = 32,768 1101 PS = 65,536 1110 PS = 131,072 1111 PS = 262,144 |

10.3.1.16 Flash Mode Register (FMR)

The local bus Flash mode register (FMR), shown in Figure 10-20, controls global operation of the FCM.




Figure 10-20. Flash Mode Register

¹ Bit R (field BOOT) is set if power-on-reset configuration selects FCM as the boot ROM target.

Table 10-23 describes FMR fields.

Table 10-23. FMR Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–19 | CWTO | Command wait time-out. For FCM commands that wait on <u>LFRB</u> being sampled high (CW0, CW1, RBW, and RSW), FCM pauses execution of the instruction sequence until either <u>LFRB</u> is sampled high, or a timer controlled by CTO expires, whichever occurs first. The time-out in the latter case is as follows: 0000 256 cycles of LCLK 0001 512 cycles of LCLK 0010 1024 cycles of LCLK 0011 2048 cycles of LCLK 0100 4096 cycles of LCLK 0101 8192 cycles of LCLK 0110 16,384 cycles of LCLK 0111 32,768 cycles of LCLK 1000 65,536 cycles of LCLK 1001 131,072 cycles of LCLK 1010 262,144 cycles of LCLK 1011 524,288 cycles of LCLK 1100 1,048,576 cycles of LCLK 1101 2,097,152 cycles of LCLK 1110 4,194,304 cycles of LCLK 1111 8,388,608 cycles of LCLK |
| 20 | BOOT | Flash auto-boot load mode. During system boot from NAND Flash EEPROM, this bit remains set to alter the use of the FCM buffer RAM. Software should clear BOOT once FCM is to be restored to normal operation. Setting BOOT without auto-boot in progress only alters the mapping of the buffer RAM. 0 FCM is operating in normal functional mode, with an 8 Kbyte FCM buffer RAM. 1 eLBC has been configured—either from reset or by a special operation OP = 01—to autoload a 4-Kbyte boot block into the FCM buffer RAM, which maps only the 4 Kbytes of NAND Flash main data region comprising the boot block. Any access to the buffer RAM is delayed until the entire boot block has been loaded. |
| 21–22 | — | Reserved |

Table 10-23. FMR Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 23 | ECCM | <p>ECC mode. When hardware checking and/or generation of error correcting codes (ECC) is enabled (that is, when BRη[DECC] is 01 or 10, and full page transfers are specified with FBCR[BC] = 0), ECCM sets the ECC block size and position of the ECC code word(s) in the NAND Flash spare region for both checking and generation functions. The format of the ECC code word conforms with the Samsung/Toshiba spare region assignment specifications.</p> <p>0 ECC is checked/calculated over 512-Byte blocks. A 24-bit ECC is assigned to spare region bytes at offsets $(N \times 16) + 6$ through $(N \times 16) + 8$ for spare region N, $N = 0-3$.</p> <p>1 ECC is checked/calculated over 512-Byte blocks. A 24-bit ECC is assigned to spare region bytes at offsets $(N \times 16) + 8$ through $(N \times 16) + 10$ for spare region N, $N = 0-3$.</p> |
| 24–25 | — | Reserved |
| 26–27 | AL | <p>Address length. AL sets the number of address bytes issued during page address (PA) operations. However, the number of address bytes issued for column address (CA) operations is determined by the device page size (for ORη[PGS] = 0, 1 CA byte is issued; for ORη[PGS] = 1, 2 CA bytes are issued).</p> <p>00 2 bytes are issued for page addresses, thus a total of 3 (ORη[PGS] = 0) or 4 (ORη[PGS] = 1) address bytes are issued for a {CA,PA} sequence</p> <p>01 3 bytes are issued for page addresses, thus a total of 4 (ORη[PGS] = 0) or 5 (ORη[PGS] = 1) address bytes are issued for a {CA,PA} sequence</p> <p>10 4 bytes are issued for page addresses, thus a total of 5 (ORη[PGS] = 0) or 6 (ORη[PGS] = 1) address bytes are issued for a {CA,PA} sequence</p> <p>11 —</p> |
| 28–29 | — | Reserved |
| 30–31 | OP | <p>Flash operation. For OP not equal to 00, a special operation is triggered on the next write to LSOR or dummy access to a bank controlled by FCM. Once a special operation has commenced, OP is automatically reset to 00 by FCM. Individual blocks may be temporarily unlocked for erase and reprogramming operations.</p> <p>00 Normal operation. All read and write accesses to banks controlled by FCM access the shared FCM buffer RAM. No bus activity is caused by this operation.</p> <p>01 Simulate auto-boot block loading, and set FMR[BOOT]. Boot block loading occurs from the bank triggered on the special operation, therefore the appropriate bank configuration must be initialized prior to issuing this operation.</p> <p>10 Execute the command sequence contained in FIR, but with write protection enabled (pin $\overline{\text{LFWP}}$ asserted low) so that all Flash blocks are protected from accidental erasure and reprogramming.</p> <p>11 Execute the command sequence contained in FIR, but permit the single block identified by FBAR[BLK] to be erased or reprogrammed, with pin $\overline{\text{LFWP}}$ remaining high during the access.</p> |

10.3.1.17 Flash Instruction Register (FIR)

The local bus Flash instruction register (FIR), shown in [Figure 10-21](#), holds a sequence of up to eight instructions for issue by the FCM. Setting FMR[OP] non-zero and writing LSOR or accessing a bank controlled by FCM causes FCM to read FIR 4 bits at a time, starting at bit 0 and continuing with adjacent 4-bit opcodes, until only NOP opcodes remain. The programmed instruction sequence of OP0, OP1, ..., OP7 is performed on the activated bank, using the data buffer addressed by FPAR. If LTEIR[CCI] = 1 and LTEDR[CCD] = 0, eLBC will generate an interrupt once the entire sequence has completed, and software should examine LTEATR and clear its V bit.

Software must not alter the contents of the addressed FCM buffer, FIR, MDR, FCR, FBAR, FPAR, or FBCR while an operation is in progress—or eLBC will behave unpredictably—but software can freely modify the contents of any currently unused FCM RAM buffer in preparation for the next operation.




Figure 10-21. Flash Instruction Register

Table 10-24 describes FIR fields.

Table 10-24. FIR Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–3 | OP0 | FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. |
| 4–7 | OP1 | 0000 NOP—No-operation and end of operation sequence |
| | | 0001 CA—Issue current column address as set in FPAR, with length set by ORx[PGS] |
| 8–11 | OP2 | 0010 PA—Issue current block+page address as set in FBAR and FPAR, with length set by FMR[AL] |
| | | 0011 UA—Issue user-defined address byte from next AS field in MDR |
| 12–15 | OP3 | 0100 CM0—Issue command from FCR[CMD0] |
| | | 0101 CM1—Issue command from FCR[CMD1] |
| 16–19 | OP4 | 0110 CM2—Issue command from FCR[CMD2] |
| | | 0111 CM3—Issue command from FCR[CMD3] |
| 20–23 | OP5 | 1000 WB—Write FBCR bytes of data from current FCM buffer to Flash device |
| 24–27 | OP6 | 1001 WS—Write one byte (8b port) of data from next AS field of MDR to Flash device |
| | | 1010 RB—Read FBCR bytes of data from Flash device into current FCM RAM buffer |
| 28–31 | OP7 | 1011 RS—Read one byte (8b port) of data from Flash device into next AS field of MDR |
| | | 1100 CW0—Wait for LFRB to return high or time-out, then issue command from FCR[CMD0] |
| | | 1101 CW1—Wait for LFRB to return high or time-out, then issue command from FCR[CMD1] |
| | | 1110 RBW—Wait for LFRB to return high or time-out, then read FBCR bytes of data from Flash device into current FCM RAM buffer |
| | | 1111 RSW—Wait for LFRB to return high or time-out, then read one byte (8b port) of data from Flash device into next AS field of MDR |

10.3.1.18 Flash Command Register (FCR)

The local bus Flash command register (FCR), shown in Figure 10-22, holds up to four NAND Flash EEPROM command bytes that may be referenced by opcodes in FIR during FCM operation. The values of the commands should follow the manufacturer’s datasheet for the relevant NAND Flash device.




Figure 10-22. Flash Command Register




Figure 10-25. Flash Page Address Register, Large Page Device (ORx[PGS] = 1)

Table 10-27 describes FPAR fields for small page devices.

Table 10-27. FPAR Field Descriptions, Small Page Device (ORx[PGS] = 0)

| Bits | Name | Description |
|-------|------|--|
| 0–16 | — | Reserved |
| 17–21 | PI | Page index. PI indexes the page in NAND Flash EEPROM at the current block defined by FBAR, and locates the corresponding transfer buffer in the FCM buffer RAM. The 3 LSBs of PI index one of the eight 1 Kbyte buffers in the FCM buffer RAM as follows: 000 The page is transferred to/from FCM buffer 0, address offsets 0x0000–0x03FF 001 The page is transferred to/from FCM buffer 1, address offsets 0x0400–0x07FF 010 The page is transferred to/from FCM buffer 2, address offsets 0x0800–0x0BFF 011 The page is transferred to/from FCM buffer 3, address offsets 0x0C00–0x0FFF 100 The page is transferred to/from FCM buffer 4, address offsets 0x1000–0x13FF 101 The page is transferred to/from FCM buffer 5, address offsets 0x1400–0x17FF 110 The page is transferred to/from FCM buffer 6, address offsets 0x1800–0x1BFF 111 The page is transferred to/from FCM buffer 7, address offsets 0x1C00–0x1FFF |
| 22 | MS | Main/spare region locator. In the case that FBCR[BC] = 0, MS is treated as 0. 0 Data is transferred to/from the main region of the FCM buffer; that is, the first 512 bytes of the buffer are used as the starting address. 1 Data is transferred to/from the spare region of the FCM buffer; that is, the second 512 bytes of the buffer are used as the starting address, but only an initial 16 bytes of spare region are defined. |
| 23–31 | CI | Column index. CI indexes the first byte to transfer to/from the main or spare region of the NAND Flash EEPROM and corresponding transfer buffer. In the case that FBCR[BC] = 0, CI is treated as 0. For MS = 0, CI can range 0x000–0x1FF; for MS = 1, CI can range 0x000–0x00F. |

Table 10-28 describes FPAR fields for large page devices.

Table 10-28. FPAR Field Descriptions, Large Page Device (ORx[PGS] = 1)

| Bits | Name | Description |
|-------|------|---|
| 0–13 | — | Reserved |
| 14–19 | PI | Page index. PA indexes the page in NAND Flash EEPROM at the current block defined by FBAR, and locates the corresponding transfer buffer in the FCM buffer RAM. The LSB of PI indexes one of the two 4 Kbyte buffers in the FCM buffer RAM as follows: 0 The page is transferred to/from FCM buffer 0, address offsets 0x0000–0x0FFF 1 The page is transferred to/from FCM buffer 1, address offsets 0x1000–0x1FFF |

Table 10-28. FPAR Field Descriptions, Large Page Device (ORx[PGS] = 1) (continued)

| Bits | Name | Description |
|-------|------|---|
| 20 | MS | Main/spare region locator. In the case that FBCR[BC] = 0, MS is treated as 0. 0 Data is transferred to/from the main region of the FCM buffer; that is, the first 2048 bytes of the buffer are used as the starting address. 1 Data is transferred to/from the spare region of the FCM buffer; that is, the second 2048 bytes of the buffer are used as the starting address, but only an initial 64 bytes of spare region are defined. |
| 21–31 | CI | Column index. CI indexes the first byte to transfer to/from the main or spare region of the NAND Flash EEPROM and corresponding transfer buffer. In the case that FBCR[BC] = 0, CI is treated as 0. For MS = 0, CI can range 0x000–0x7FF; for MS = 1, CI can range 0x000–0x03F. |

10.3.1.21 Flash Byte Count Register (FBCR)

The local bus Flash byte count register (FBCR), shown in [Figure 10-26](#), defines the size of FCM block transfers for reads and writes to the NAND Flash EEPROM.




Figure 10-26. Flash Byte Count Register

[Table 10-29](#) describes FBCR fields.

Table 10-29. FBCR Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–19 | — | Reserved |
| 20–31 | BC | Byte count determines how many bytes are transferred by the FCM during data read (RB) or data write (WB) opcodes. The first byte accessed in the NAND Flash EEPROM is located by the FPAR register, and successive bytes are transferred until either BC bytes have been counted, or the end of the spare region of the currently addressed Flash page has been reached. If BC = 0, an entire Flash page and its spare region will be transferred by FCM, in which case FPAR[MS] and FPAR[CI] are treated as zero regardless of their values. BC = 0 is the only setting that permits FCM to generate and check ECC. |

10.4 Functional Description

The eLBC allows the implementation of memory systems with very specific timing requirements.

- The GPCM provides interfacing for simpler, lower-performance memories and memory-mapped devices. It has inherently lower performance because it does not support bursting. For this reason, GPCM-controlled banks are used primarily for boot-loading from NVRAM or NOR Flash, and access to low-performance memory-mapped peripherals.
- The FCM interfaces the eLBC to NAND Flash EEPROMs with 8-bit data bus. The FCM has an automatic boot-loading feature that allows the CPU to boot from high density EEPROM, loading the boot block into 4 Kbytes of RAM for execution of the first level boot code. Following boot,

FCM provides a flexible instruction sequencer that allows a user-defined command, address, and data transfer sequence of up to 8 steps to be executed against a memory-mapped buffer RAM. Programmable set-up time, hold time, and wait states permit the FCM to maximize the performance of NAND Flash block transfers, which can proceed in parallel with software processing of the multiple RAM buffers. A single-pass ECC engine in the FCM permits zero-overhead error checking, reporting, and correction in both boot blocks and page data transfers if enabled.

- The UPM supports refresh timers, address multiplexing of the external bus, and generation of programmable control signals for row address and column address strobes, to allow for a minimal glue logic interface to DRAMs, burstable SRAMs, and almost any other kind of peripheral with asynchronous timing or single data rate clocking. The UPM can be used to generate flexible, user-defined timing patterns for control signals that govern a memory device. These patterns define how the external control signals behave during a read, write, burst-read, or burst-write access. Refresh timers are also available to periodically initiate user-defined refresh patterns.

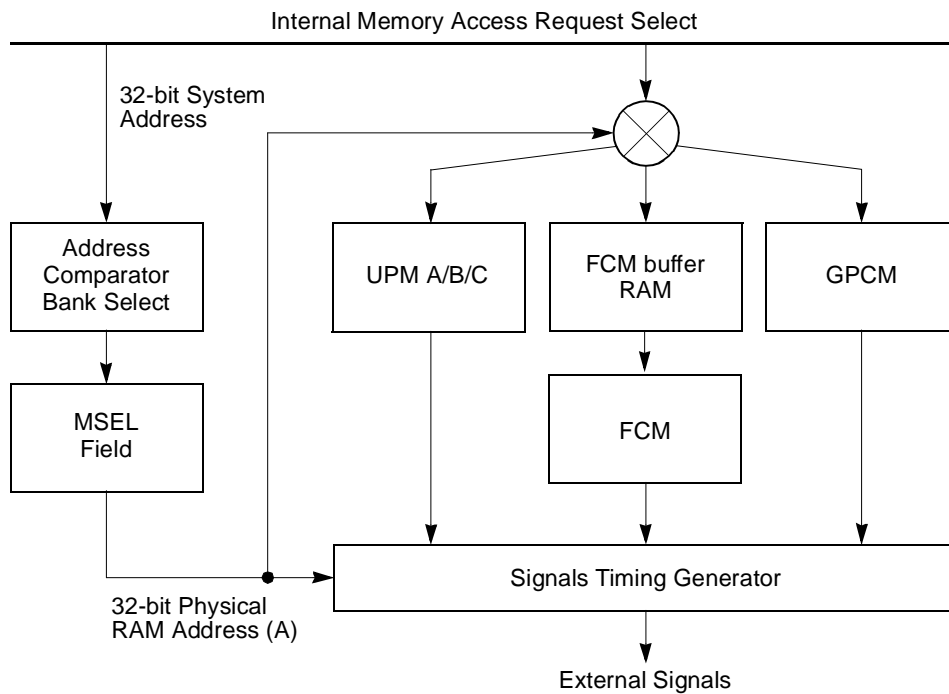


Figure 10-27. Basic Operation of Memory Controllers in the eLBC

Each memory bank (chip select) can be assigned to any of these three types of machines through the machine select bits of the base register for that bank ($BR_n[MSEL]$), as illustrated in [Figure 10-27](#). If a bank match occurs, the corresponding machine (GPCM, FCM, or UPM) then takes ownership of the external signals that control the access and maintains control until the transaction ends.

NOTE

Different machines (FCM/GPCM/UPM) share the address, data, and control signals. While the eLBC is servicing a transaction, subsequent transactions are queued until the current transaction has completed.

10.4.1 Basic Architecture

The following subsections describe the basic architecture of the eLBC.

10.4.1.1 Address and Address Space Checking

The defined base addresses are written to the BR_n registers, while the corresponding address masks are written to the OR_n registers. Each time a local bus access is requested, the internal transaction address is compared with each bank. Addresses are decoded by comparing the 17 MSBs of the address, masked by $OR_n[AM]$, with the base address for each bank ($BR_n[BA]$). If a match is found on a memory controller bank, the attributes defined in the BR_n and OR_n for that bank are used to control the memory access. If a match is found in more than one bank, the lowest-numbered bank handles the memory access (that is, bank 0 has priority over bank 1).

10.4.1.2 External Address Latch Enable Signal (LALE)

The local bus uses a multiplexed address/data bus. Therefore the eLBC must distinguish between address and data phases, which take place on the same bus (LAD pins). The LALE signal, when asserted, signifies an address phase during which the eLBC drives the memory address on the LAD pins. An external address latch uses this signal to capture the address and provide it to the address pins of the memory or peripheral device. When LALE is negated, LAD then serves as the (bi-directional) data bus for the access. Any address phase initiates the assertion of LALE, which has a programmable duration of between 1 and 4 bus clock cycles.

To ensure adequate hold time on the external address latch, LALE negates earlier than the address changes on LAD during address phases. By default, LALE negates earlier by 2 platform clock period. For example, if the platform clock is operating at 66.6 MHz, then 3 ns of address hold time is introduced. However, at higher frequencies, the duration of the shortened LALE pulse may not meet the minimum latch enable pulse width specifications of some latches. In such cases, setting $LBCR[AHD] = 1$ increases the LALE pulse width by 1 platform clock cycle, but decreases the address hold time by the same amount. If both longer hold time and longer LALE pulse duration are needed, then the address phase can be extended using the $OR_n[EAD]$ and $LCRR[EADC]$ fields, and the $LBCR[AHD]$ bit can be left at 0. However, this will add latency to all address tenures.

The frequency of LALE assertion varies across the three memory controllers:

- For GPCM, every assertion of $\overline{LCS_n}$ is considered an independent access, and accordingly, LALE asserts prior to each such access. For example, GPCM driving an 8-bit port would assert LALE and $\overline{LCS_n}$ 32 times in order to satisfy a 32-byte cache line transfer.
- For FCM, LALE asserts prior to each multi-command operation sequence, but LALE can be ignored on NAND Flash EEPROM accesses as the signal does *not* enable address latching in such devices. The value on the LAD and LA pins during LALE assertion is driven low-impedance, but otherwise not defined for FCM banks.
- In the case of UPM, the frequency of LALE assertion depends on how the UPM RAM is programmed. UPM single accesses typically assert LALE once, upon commencement, but it is possible to program UPM to assert LALE several times, and to change the values of LAN with and without LALE being involved.

In general, when using the GPCM controller it is not necessary to use LA if a sufficiently wide latch is used to capture the entire address during LALE phases. The UPMs may require LA if the eLBC is generating its own burst address sequence.

To illustrate how a large transaction is handled by the eLBC, Figure 10-28 shows eLBC signals for the GPCM performing a 32-byte write starting at address 0x5420. Note that during each of the 32 assertions of LALE, LA[21:25] exactly mirror LAD[27:31], but during data phases, only LAD[0:7] is driven with valid data.

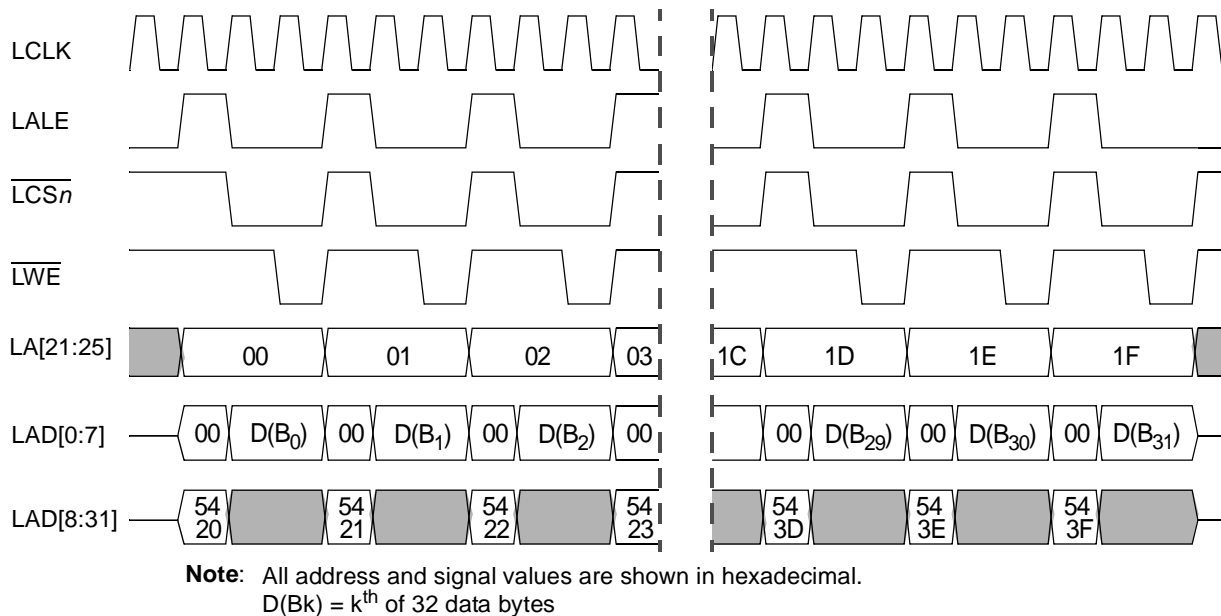


Figure 10-28. Example of 8-Bit GPCM Writing 32 Bytes to Address 0x5420 (LCRR[PBYP] = 0)

If the RCW is loaded by the eLBC, LALE may remain at an unknown value for up to 8 cycles after PORESET negation. Thus, LALE should be ignored for 8 CLKIN/PCI_SYNC_IN cycles after PORESET negation. In general, it is recommended that a latch be implemented for this adjustment and not a state machine triggered by LALE.

If the RCW is not loaded by the eLBC (for example, I2C or hard-coded options are used), then LALE is at an unknown value until the PLL is locked and should be ignored until the negation of HRESET.

10.4.1.3 Data Transfer Acknowledge (TA)

The three memory controllers in the eLBC generate an internal transfer acknowledge signal, TA, to allow data on LAD to be either sampled (for reads) or changed (on writes). The data sampling/data change always occurs at the end of the bus cycle in which the eLBC asserts TA internally. In eLBC debug mode, TA is also visible externally on the LDVAL pin. The GPCM controller automatically generates TA according to the timing parameters programmed for them in the option and mode registers; FCM generates TA whenever data read and write instructions are executed out of register FIR; a UPM generates TA only when a UPM pattern has the UTA RAM word bit set. Figure 10-29 shows LALE, TA (internal), and \overline{LCSn} .

Note that TA and LALE are never asserted together, and that for the duration of LALE, $\overline{\text{LCS}}_n$ (or any other control signal) remains negated or frozen.




Figure 10-29. Basic eLBC Bus Cycle with LALE, TA, and $\overline{\text{LCS}}_n$

10.4.1.4 Data Buffer Control (LBCTL)

The memory controller provides a data buffer control signal for the local bus (LBCTL). This signal is activated when a GPCM-, FCM-, or UPM-controlled bank is accessed. LBCTL can be disabled by setting $\text{OR}_n[\text{BCTLD}]$. LBCTL can be further configured by $\text{LBCR}[\text{BCTLC}]$ to act as an extra $\overline{\text{LWE}}$ or an extra $\overline{\text{LOE}}$ signal when in GPCM mode.

If LBCTL is configured as a data buffer control ($\text{LBCR}[\text{BCTLC}] = 00$), the signal is asserted (high) on the rising edge of the bus clock on the first cycle of the memory controller operation, coincident with LALE. If the access is a write, LBCTL remains high for the whole duration. However, if the access is a read, LBCTL is negated (low) with the negation of LALE so that the memory device is able to drive the bus. If back-to-back read accesses are pending, LBCTL is asserted (high) one bus clock cycle before the next transaction starts (that is, one bus clock cycle before LALE) to allow a whole bus cycle for the bus to turn around before the next address is driven.

10.4.1.5 Bus Monitor

A bus monitor is provided to ensure that each bus cycle is terminated within a reasonable (user defined) period. When a transaction starts, the bus monitor starts counting down from the time-out value ($\text{LBCR}[\text{BMT}] \times \text{LBCR}[\text{BMTPS}]$) until a data beat is acknowledged on the bus. It then reloads the time-out value and resumes the countdown until the data tenure completes and then idles if there is no pending transaction. Setting $\text{LTEDR}[\text{BMD}]$ disables bus monitor error checking (that is, the $\text{LTESR}[\text{BM}]$ bit is not set by a bus monitor time-out); however, the bus monitor is still active and can generate a UPM exception (as noted in [Section 10.4.4.1.4, “Exception Requests,”](#)) or terminate a GPCM access.

It is very important to ensure that the value of $\text{LBCR}[\text{BMT}]$ is not set too low; otherwise spurious bus time-outs may occur during normal operation—resulting in incomplete data transfers. Accordingly, the time-out value represented by the $\text{LBCR}[\text{BMT}]$, $\text{LBCR}[\text{BMTPS}]$ pair must not be set below 40 bus cycles for time-out under any circumstances.

NOTE

When the FCM is in the middle of a long transaction (such as NAND erase, write, and so on), another transaction on the GPCM or UPM will trigger the bus monitor to start, even though the GPCM or UPM is waiting for the FCM to finish. If the bus monitor times out, it could corrupt the current NAND flash operation as well as terminate the GPCM or UPM operation. To avoid such cases, it is recommended that the bus monitor timeout be programmed to its maximum setting of $LBCR[BMT] = 0$ and $LBCR[BMTPS] = 0xF$.

10.4.2 General-Purpose Chip-Select Machine (GPCM)

The GPCM allows a minimal glue logic and flexible interface to SRAM, EPROM, FEPRM, ROM devices, and external peripherals. The GPCM contains two basic configuration register groups— BR_n and OR_n .

Figure 10-30 shows a simple connection between an 8-bit port size SRAM device and the eLBC in GPCM mode. Byte-write enable signals (\overline{LWE}) are available for each byte written to memory. Also, the output enable signal (\overline{LOE}) is provided to minimize external glue logic. On system reset, a global (boot) chip-select is available that provides a boot ROM chip-select (LCS_0) prior to the system being fully configured.

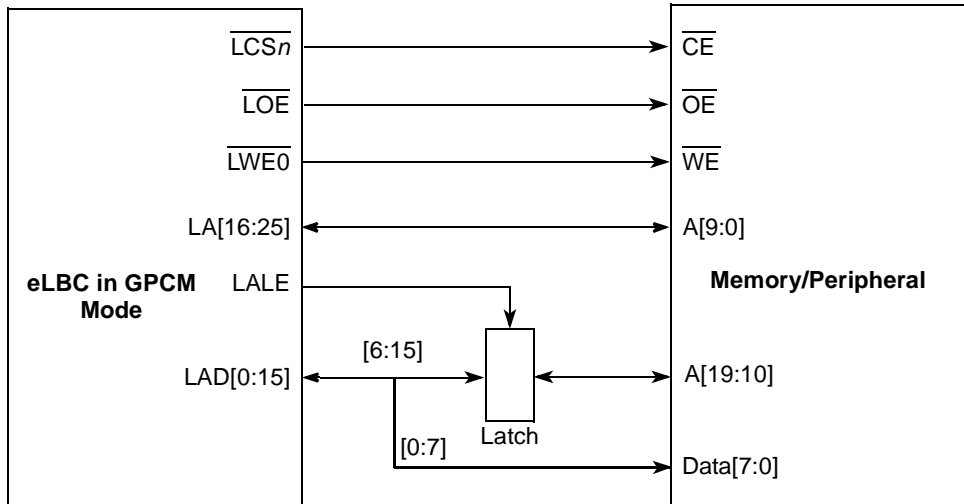


Figure 10-30. Enhanced Local Bus to GPCM Device Interface

Figure 10-31 shows \overline{LCS} as defined by the setup time required between the address lines and \overline{CE} . The user can configure $OR_n[ACS]$ to specify \overline{LCS} to meet this requirement. Generally, the attributes for the

memory cycle are taken from OR_n . These attributes include the CSNT, ACS, XACS, SCY, TRLX, EHTR, and SETA fields.





Figure 10-31. GPCM Basic Read Timing (XACS = 0, ACS = 1x, TRLX = 0, CLKDIV = 4,8)

10.4.2.1 GPCM Read Signal Timing

The basic GPCM read timing parameters that may be set by the OR_n attributes are shown in Figure 10-32. The read access cycle commences upon latching of the memory address (LALE negated), and concludes when LBCTL returns high to turn the local bus around for a subsequent address phase. Read data is captured by eLBC on the falling edge of TA. \overline{LOE} and \overline{LCSn} negate high simultaneously, in some cases before the end of the read access to provide additional hold time for the external memory.



Notes:
 t_{RC} = Read cycle time. t_{CSRP} = Read chip-select assertion period.
 t_{ARCS} = Address valid to read chip-select time. t_{OEN} = Output enable negated time.
 t_{AOE} = Address valid to output enable time.

Figure 10-32. GPCM General Read Timing Parameters

Table 10-30 lists the signal timing parameters for a GPCM read access as the option register attributes are varied.

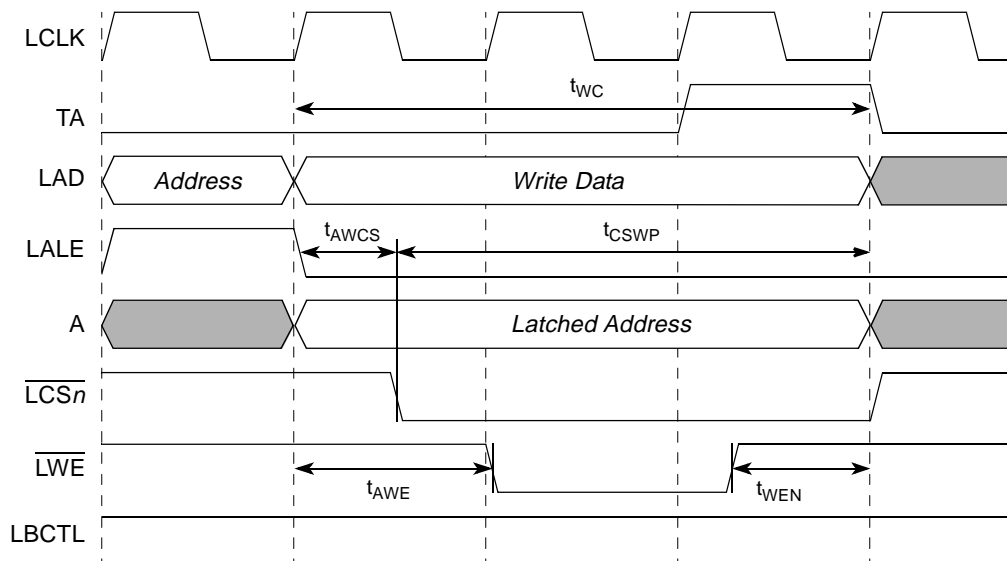
Table 10-30. GPCM Read Control Signal Timing

| Option Register Attributes | | | | Signal Timing (LCLK clock cycles) ¹ | | | | |
|----------------------------|------|------|-----|--|----------------------------|------------------|------------------|-----------------|
| TRLX | EHTR | XACS | ACS | t _{ARCS} | t _{CSRP} | t _{AOE} | t _{OEN} | t _{RC} |
| 0 | 0 | 0 | 0X | 0 | 2 + SCY | 1 | 0 | 2 + SCY |
| 0 | 0 | 0 | 10 | ¼ (½) | 1¾ + SCY (2+SCY) | 1 | 0 | 2 + SCY |
| 0 | 0 | 0 | 11 | ½ | 1½ + SCY | 1 | 0 | 2 + SCY |
| 0 | 0 | 1 | 0X | 0 | 2 + SCY | 1 | 0 | 2 + SCY |
| 0 | 0 | 1 | 10 | 1 | 1 + SCY | 1 | 0 | 2 + SCY |
| 0 | 0 | 1 | 11 | 2 | 1 + SCY | 2 | 0 | 3 + SCY |
| 0 | 1 | 0 | 0X | 0 | 2 + SCY | 1 | 1 | 3 + SCY |
| 0 | 1 | 0 | 10 | ¼ (½) | 1¾ + SCY (1½+SCY) | 1 | 1 | 3 + SCY |
| 0 | 1 | 0 | 11 | ½ | 1½ + SCY | 1 | 1 | 3 + SCY |
| 0 | 1 | 1 | 0X | 0 | 2 + SCY | 1 | 1 | 3 + SCY |
| 0 | 1 | 1 | 10 | 1 | 1 + SCY | 1 | 1 | 3 + SCY |
| 0 | 1 | 1 | 11 | 2 | 1 + SCY | 2 | 1 | 4 + SCY |
| 1 | 0 | 0 | 0X | 0 | 2 + 2 × SCY | 1 | 4 | 6 + 2 × SCY |
| 1 | 0 | 0 | 10 | 1¼ (1½) | 1¾ + 2 × SCY (1½+2×SCY) | 2 | 4 | 7 + 2 × SCY |
| 1 | 0 | 0 | 11 | 1½ | 1½ + 2 × SCY | 2 | 4 | 7 + 2 × SCY |
| 1 | 0 | 1 | 0X | 0 | 2 + 2 × SCY | 1 | 4 | 6 + 2 × SCY |
| 1 | 0 | 1 | 10 | 2 | 1 + 2 × SCY | 2 | 4 | 7 + 2 × SCY |
| 1 | 0 | 1 | 11 | 3 | 1 + 2 × SCY | 3 | 4 | 8 + 2 × SCY |
| 1 | 1 | 0 | 0X | 0 | 2 + 2 × SCY | 1 | 8 | 10 + 2 × SCY |
| 1 | 1 | 0 | 10 | 1¼ (1½) | 1¾ + 2 × SCY (1½+2×SCY) | 2 | 8 | 11 + 2 × SCY |
| 1 | 1 | 0 | 11 | 1½ | 1½ + 2 × SCY | 2 | 8 | 11 + 2 × SCY |
| 1 | 1 | 1 | 0X | 0 | 2 + 2 × SCY | 1 | 8 | 10 + 2 × SCY |
| 1 | 1 | 1 | 10 | 2 | 1 + 2 × SCY | 2 | 8 | 11 + 2 × SCY |
| 1 | 1 | 1 | 11 | 3 | 1 + 2 × SCY | 3 | 8 | 12 + 2 × SCY |

¹

10.4.2.2 GPCM Write Signal Timing

The basic GPCM write timing parameters that may be set by the OR_n attributes are shown in Figure 10-33. The write access cycle commences upon latching of the memory address (LALE negated), and concludes when \overline{LCSn} returns high. LBCTL remains stable for the entire cycle to drive data onto any secondary data bus. Write data becomes invalid following the falling edge of TA. \overline{LWE} may, in some cases, negate high before the end of the write access to provide additional hold time for the external memory.



Notes:
 t_{WC} = Write cycle time.
 t_{AWCS} = Address valid to write chip-select time.
 t_{AWE} = Address valid to write enable time.
 t_{CSWP} = Write chip-select assertion period.
 t_{WEN} = Write enable negated time wrt chip-select.

Figure 10-33. GPCM General Write Timing Parameters

Table 10-31 lists the signal timing parameters for a GPCM write access as the option register attributes are varied.

Table 10-31. GPCM Write Control Signal Timing

| Option Register Attributes | | | | Signal Timing (LCLK clock cycles) ¹ | | | | |
|----------------------------|------|-----|------|--|---------------------------------|-----------|----------------------|----------|
| TRLX | XACS | ACS | CSNT | t_{AWCS} | t_{CSWP} | t_{AWE} | t_{WEN} | t_{WC} |
| 0 | 0 | 00 | 0 | 0 | 2 + SCY | 1 | 0 | 2 + SCY |
| 0 | 0 | 10 | 0 | $\frac{1}{4}$ ($\frac{1}{2}$) | $1\frac{3}{4}$ + SCY (2+SCY) | 1 | 0 | 2 + SCY |
| 0 | 0 | 11 | 0 | $\frac{1}{2}$ | $1\frac{1}{2}$ + SCY | 1 | 0 | 2 + SCY |
| 0 | 1 | 00 | 0 | 0 | 2 + SCY | 1 | 0 | 2 + SCY |
| 0 | 1 | 10 | 0 | 1 | 1 + SCY | 1 | 0 | 2 + SCY |
| 0 | 1 | 11 | 0 | 2 | 1 + SCY | 2 | 0 | 3 + SCY |
| 0 | 0 | 00 | 1 | 0 | 2 + SCY | 1 | $\frac{1}{4}$ (0) | 2 + SCY |

Table 10-31. GPCM Write Control Signal Timing (continued)

| Option Register Attributes | | | | Signal Timing (LCLK clock cycles) ¹ | | | | |
|----------------------------|------|-----|------|--|--|------------------|----------------------|--|
| TRLX | XACS | ACS | CSNT | t _{AWCS} | t _{CSWP} | t _{AWE} | t _{WEN} | t _{wc} |
| 0 | 0 | 10 | 1 | $\frac{1}{4}$ ($\frac{1}{2}$) | $1\frac{1}{2} + \text{SCY}$ | 1 | 0 | $1\frac{3}{4} + \text{SCY}$ ($1\frac{1}{2} + \text{SCY}$) |
| 0 | 0 | 11 | 1 | $\frac{1}{2}$ | $1\frac{1}{4} + \text{SCY}$ ($1 + \text{SCY}$) | 1 | 0 | $1\frac{3}{4} + \text{SCY}$ ($1\frac{1}{2} + \text{SCY}$) |
| 0 | 1 | 00 | 1 | 0 | $2 + \text{SCY}$ | 1 | $\frac{1}{4}$ (0) | $2 + \text{SCY}$ |
| 0 | 1 | 10 | 1 | 1 | $\frac{3}{4} + \text{SCY}$ ($\frac{1}{2} + \text{SCY}$) | 1 | 0 | $1\frac{3}{4} + \text{SCY}$ ($1\frac{1}{2} + \text{SCY}$) |
| 0 | 1 | 11 | 1 | 2 | $\frac{3}{4} + \text{SCY}$ ($\frac{1}{2} + \text{SCY}$) | 2 | 0 | $2\frac{3}{4} + \text{SCY}$ ($2\frac{1}{2} + \text{SCY}$) |
| 1 | 0 | 00 | 0 | 0 | $2 + 2 \times \text{SCY}$ | 1 | 0 | $2 + 2 \times \text{SCY}$ |
| 1 | 0 | 10 | 0 | $\frac{1}{4}$ ($\frac{1}{2}$) | $1\frac{3}{4} + 2 \times \text{SCY}$ ($2 + 2 \times \text{SCY}$) | 2 | 0 | $3 + 2 \times \text{SCY}$ |
| 1 | 0 | 11 | 0 | $\frac{1}{2}$ | $1\frac{1}{2} + 2 \times \text{SCY}$ | 2 | 0 | $3 + 2 \times \text{SCY}$ |
| 1 | 1 | 00 | 0 | 0 | $2 + 2 \times \text{SCY}$ | 1 | 0 | $2 + 2 \times \text{SCY}$ |
| 1 | 1 | 10 | 0 | 2 | $1 + 2 \times \text{SCY}$ | 2 | 0 | $3 + 2 \times \text{SCY}$ |
| 1 | 1 | 11 | 0 | 3 | $1 + 2 \times \text{SCY}$ | 3 | 0 | $4 + 2 \times \text{SCY}$ |
| 1 | 0 | 00 | 1 | 0 | $3 + 2 \times \text{SCY}$ | 1 | $\frac{1}{4}$ (1) | $3 + 2 \times \text{SCY}$ |
| 1 | 0 | 10 | 1 | $\frac{1}{4}$ ($\frac{1}{2}$) | $1\frac{1}{2} + 2 \times \text{SCY}$ | 2 | 0 | $2\frac{3}{4} + 2 \times \text{SCY}$ ($2\frac{1}{2} + 2 \times \text{SCY}$) |
| 1 | 0 | 11 | 1 | $\frac{1}{2}$ | $1\frac{1}{4} + 2 \times \text{SCY}$ ($1 + 2 \times \text{SCY}$) | 2 | 0 | $2\frac{3}{4} + 2 \times \text{SCY}$ ($2\frac{1}{2} + 2 \times \text{SCY}$) |
| 1 | 1 | 00 | 1 | 0 | $3 + 2 \times \text{SCY}$ | 1 | $\frac{1}{4}$ (1) | $3 + 2 \times \text{SCY}$ |
| 1 | 1 | 10 | 1 | 2 | $\frac{3}{4} + 2 \times \text{SCY}$ ($\frac{1}{2} + 2 \times \text{SCY}$) | 2 | 0 | $2\frac{3}{4} + 2 \times \text{SCY}$ ($2\frac{1}{2} + 2 \times \text{SCY}$) |
| 1 | 1 | 11 | 1 | 3 | $\frac{3}{4} + 2 \times \text{SCY}$ ($\frac{1}{2} + 2 \times \text{SCY}$) | 3 | 0 | $3\frac{3}{4} + 2 \times \text{SCY}$ ($3\frac{1}{2} + 2 \times \text{SCY}$) |

¹ Times in parentheses are specific for the case LCRR[CLKDIV] = 2; other times apply to all CLKDIV values.

10.4.2.3 Chip-Select Assertion Timing

The banks selected to work with the GPCM support an option to drive the $\overline{\text{LCS}}_n$ signal with different timings (with respect to the external address/data bus). $\overline{\text{LCS}}_n$ can be driven in any of the following ways:

- Simultaneous with the latched memory address. (This refers to the externally latched address and not the address timing on LAD. That is, the chip select does not assert during LALE).
- One quarter of a clock cycle later (for LCRR[CLKDIV] = 4, 8).

- One half of a clock cycle later (for LCRR[CLKDIV] = 2, 4, or 8).
- One clock cycle later (for LCRR[CLKDIV] = 4), when OR_n[XACS] = 1.
- Two clock cycles later (for LCRR[CLKDIV] = 2, 4, or 8), when OR_n[XACS] = 1.
- Three clock cycles later (for LCRR[CLKDIV] = 2, 4, or 8), when OR_n[XACS] = 1 and OR_n[TRLX] = 1.

The timing diagram in [Figure 10-31](#) shows two chip-select assertion timings for the case LCRR[CLKDIV] = 4 or 8. If LCRR[CLKDIV] = 2, $\overline{\text{LCS}}_n$ asserts identically for OR_n[ACS] = 10 or 11.

10.4.2.3.1 Programmable Wait State Configuration

The GPCM supports internal generation of transfer acknowledge. It allows between zero and 30 wait states to be added to an access by programming OR_n[SCY] and OR_n[TRLX]. Internal generation of transfer acknowledge is enabled if OR_n[SETA] = 0. If $\overline{\text{LGTA}}$ is asserted externally two bus clock cycles or more before the wait state counter has expired (to allow for synchronization latency), the current memory cycle is terminated by $\overline{\text{LGTA}}$; otherwise it is terminated by the expiration of the wait state counter. Regardless of the setting of OR_n[SETA], wait states prolong the assertion duration of both $\overline{\text{LOE}}$ and $\overline{\text{LWE}}_n$ in the same manner. When TRLX = 1, the number of wait states inserted by the memory controller is doubled from OR_n[SCY] cycles to 2 × OR_n[SCY] cycles, allowing a maximum of 30 wait states.

10.4.2.3.2 Chip-Select and Write Enable Negation Timing

[Figure 10-30](#) shows a basic connection between the local bus and a static memory device. In this case, $\overline{\text{LCS}}_n$ is connected directly to $\overline{\text{CE}}$ of the memory device. The $\overline{\text{LWE}}[0:1]$ signals are connected to the respective $\overline{\text{WE}}[1:0]$ signals on the memory device where each $\overline{\text{LWE}}[0:1]$ signal corresponds to a different data byte.




Figure 10-34. GPCM Basic Write Timing
(XACS = 0, ACS = 00, CSNT = 1, SCY = 1, TRLX = 0, CLKDIV = 4, 8)

As [Figure 10-34](#) shows, the timing for $\overline{\text{LCS}}_n$ is the same as for the latched address. The strobes for the transaction are supplied by $\overline{\text{LOE}}$ or $\overline{\text{LWE}}_n$, depending on the transaction direction—read or write (write

case shown in the figure). $OR_n[CSNT]$, along with $OR_n[TRLX]$, control the timing for the appropriate strobe negation in write cycles. When this attribute is asserted, the strobe is negated one quarter of a clock before the normal case provided that $LCRR[CLDIV] = 4$ or 8 . For example, when $ACS = 00$ and $CSNT = 1$, \overline{LWEn} is negated one quarter of a clock earlier, as shown in [Figure 10-34](#). If $LCRR[CLDIV] = 2$, \overline{LWEn} is negated either coincident with \overline{LCSn} or one cycle earlier.

1. \overline{LCSn} is affected by $CSNT$ and $TRLX$ only if $ACS[0]$ is non zero. However, \overline{LWEn} is affected independent of ACS .
2. When $CSNT$ attribute is asserted, the strobe is negated one quarter of a clock before the normal case provided that $LCRR[CLDIV] = 4$ or 8 .
3. $TRLX = 1$ in conjunction with $CSNT = 1$, negates the \overline{LCSn} and \overline{LWEn} $1 + 1/4$ cycle earlier if $LCRR[CLKDIV] = 4$ or 8 .

If $LCRR[CLKDIV] = 2$, \overline{LCSn} and \overline{LWEn} are negated either normally or one cycle earlier if $TRLX = 1$.

For example, when $ACS = 00$, $CSNT = 1$ and $TRLX = 0$, \overline{LWEn} is negated one quarter of a clock earlier and \overline{LCSn} is negated normally as shown in [Figure 10-34](#).

10.4.2.3.3 Relaxed Timing

$OR_x[TRLX]$ is provided for memory systems that require more relaxed timing between signals. Setting $TRLX = 1$ has the following effect on timing:

- An additional bus cycle is added between the address and control signals (but only if ACS is not equal to 00).
- The number of wait states specified by SCY is doubled, providing up to 30 wait states.
- The extended hold time on read accesses ($EHTR$) is extended further.
- \overline{LCSn} signals are negated one cycle earlier during writes (but only if ACS is not equal to 00).
- $\overline{LWE}[0:1]$ signals are negated one cycle earlier during writes.

[Figure 10-35](#) and [Figure 10-36](#) show relaxed timing read and write transactions. The effect of $LCRR[CLKDIV] = 2$ for these examples is only to delay the assertion of \overline{LCSn} in the $ACS = 10$ case to

the ACS = 11 case. The example in Figure 10-36 also shows address and data multiplexing on LAD for a pair of writes issued consecutively.




Figure 10-35. GPCM Relaxed Timing Back-to-Back Reads
 (XACS = 0, ACS = 1x, SCY = 1, CSNT = 0, TRLX = 1, EHTR = 0, CLKDIV = 4, 8)




Figure 10-36. GPCM Relaxed Timing Back-to-Back Writes
 (XACS = 0, ACS = 1x, SCY = 0, CSNT = 0, TRLX = 1, CLKDIV = 4, 8)

When TRLX and CSNT are set in a write access, the $\overline{\text{LWE}}[0:1]$ strobe signals are negated one clock earlier than in the normal case, as shown in Figure 10-37 and Figure 10-38. If $\text{ACS} \neq 00$, $\overline{\text{LCSn}}$ is also negated one clock earlier.




Figure 10-37. GPCM Relaxed Timing Write
 (XACS = 0, ACS = 10, SCY = 0, CSNT = 1, TRLX = 1, CLKDIV = 4, 8)




Figure 10-38. GPCM Relaxed Timing Write
 (XACS = 0, ACS = 00, SCY = 1, CSNT = 1, TRLX = 1, CLKDIV = 4, 8)

10.4.2.3.4 Output Enable (\overline{LOE}) Timing

The timing of the \overline{LOE} is affected only by $TRLX$. It always asserts and negates on the rising edge of the bus clock. \overline{LOE} asserts either on the rising edge of the bus clock after \overline{LCSn} is asserted or coinciding with \overline{LCSn} (if $XACS = 1$ and $ACS = 10$ or $ACS = 11$). Accordingly, assertion of \overline{LOE} can be delayed (along with the assertion of \overline{LCSn}) by programming $TRLX = 1$. \overline{LOE} negates on the rising clock edge coinciding with \overline{LCSn} negation

10.4.2.3.5 Extended Hold Time on Read Accesses

Slow memory devices that take a long time to disable their data bus drivers on read accesses should choose some combination of $ORn[TRLX,EHTR]$. Any access following a read access to the slower memory bank is delayed by the number of clock cycles specified in [Table 10-7](#) in addition to any existing bus turnaround cycle. The final bus turnaround cycle is automatically inserted by the eLBC for reads, regardless of the setting of $ORn[EHTR]$.




Figure 10-39. GPCM Read Followed by Read ($TRLX = 0$, $EHTR = 0$, Fastest Timing)

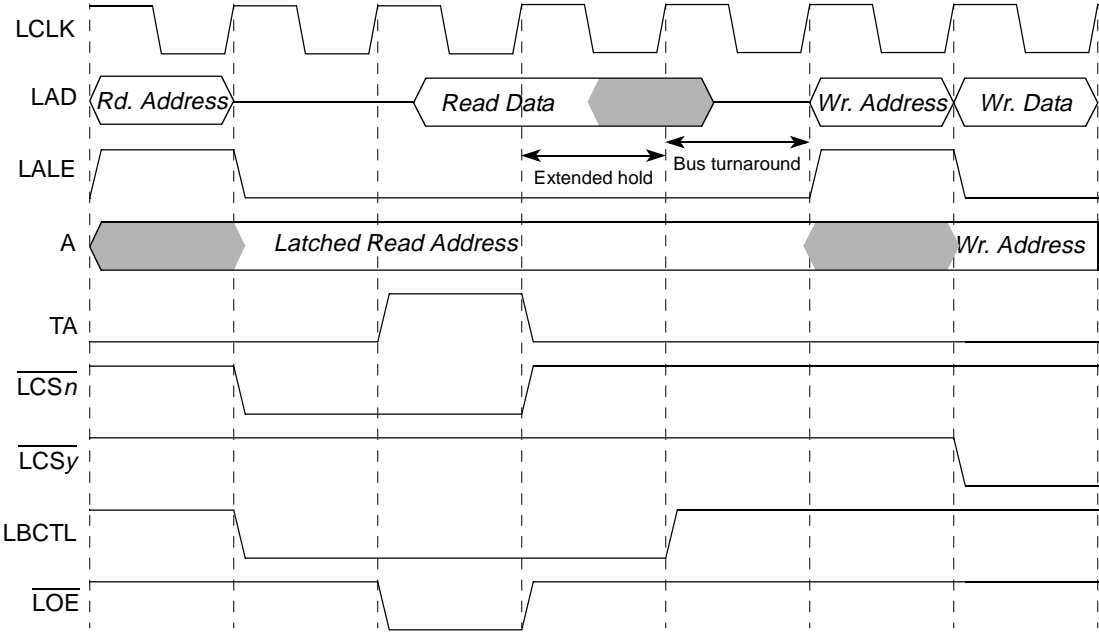


Figure 10-40. GPCM Read Followed by Write
(TRLX = 0, EHTR = 1, One-Cycle Extended Hold Time on Reads)

10.4.2.4 External Access Termination (\overline{LGTA})

External access termination is supported by the GPCM using the asynchronous \overline{LGTA} input signal, which is synchronized and sampled internally by the local bus. If, during assertion of \overline{LCSn} , the sampled \overline{LGTA} signal is asserted, it is converted to an internal generation of transfer acknowledge, which terminates the current GPCM access (regardless of the setting of $ORn[SETA]$). \overline{LGTA} should be asserted for at least one bus cycle to be effective. Note that because \overline{LGTA} is synchronized, bus termination occurs two cycles after \overline{LGTA} assertion, so in case of read cycle, the device still must drive data as long as \overline{LOE} is asserted.

The user selects whether transfer acknowledge is generated internally or externally (\overline{LGTA}) by programming $ORn[SETA]$. Asserting \overline{LGTA} always terminates an access, even if $ORn[SETA] = 0$

(internal transfer acknowledge generation), but it is the only means by which an access can be terminated if $OR_n[SETA] = 1$.

In PLL bypass mode, the timing of \overline{LGTA} is illustrated by the example in Figure 10-41.




Figure 10-41. External Termination of GPCM Access (PLL Bypass Mode)

10.4.2.5 GPCM Boot Chip-Select Operation

Boot chip-select operation allows address decoding for a boot ROM before system initialization. $\overline{LCS0}$ is the boot chip-select output; its operation differs from other external chip-select outputs after a system reset. When the core begins accessing memory after system reset, $\overline{LCS0}$ is asserted for every local bus access until BR0 or OR0 is reconfigured.

The boot chip-select also provides a programmable port size, which is configured during reset. The boot chip-select does not provide write protection. $\overline{LCS0}$ operates this way until the first write to OR0 and it can be used as any other chip-select register after the preferred address range is loaded into BR0. After the

first write to OR0, the boot chip-select can be restarted only with a hardware reset. Table 10-32 describes the initial values of the boot bank in the memory controller.

Table 10-32. Boot Bank Field Values after Reset for GPCM as Boot Controller

| Register | Field | Setting |
|----------|-------|--------------------------|
| BR0 | BA | 0000_0000_0000_0000_0 |
| | PS | <i>From RCWH[ROMLOC]</i> |
| | DECC | 00 |
| | WP | 0 |
| | MSEL | 000 |
| | — | — |
| | V | 1 |
| OR0 | AM | 0000_0000_0000_0000_0 |
| | BCTLD | 0 |
| | CSNT | 1 |
| | ACS | 11 |
| | XACS | 1 |
| | SCY | 1111 |
| | SETA | 0 |
| | TRLX | 1 |
| | EHTR | 1 |
| | EAD | 1 |

10.4.3 Flash Control Machine (FCM)

The FCM provides a glueless interface to parallel-bus NAND Flash EEPROM devices. The FCM contains three basic configuration register groups—BR_n, OR_n, and FMR.

Figure 10-42 shows a simple connection between an 8-bit port size NAND Flash EEPROM and the eLBC in FCM mode. Commands, address bytes, and data are all transferred on LAD[0:7]¹, with $\overline{\text{LFW}}\overline{\text{E}}$ asserted for transfers written to the device, or $\overline{\text{LFR}}\overline{\text{E}}$ asserted for transfers read from the device. eLBC signals LFCLE and LFALE determine whether writes are of type command (only LFCLE asserted), address (only LFALE asserted), or write data (neither LFCLE nor LFALE asserted). The NAND Flash RDY/ $\overline{\text{BSY}}$ pin is normally open-drain, and should be pulled high by a 4.7-K Ω resistor. On system reset, a global (boot)

1. Note bit numbering reversal: LAD[0] (msb) connects to Flash IO[7], while LAD[7] (lsb) connects to IO[0].

chip-select is available that provides a boot ROM chip-select ($\overline{\text{LCS0}}$) prior to the system being fully configured.

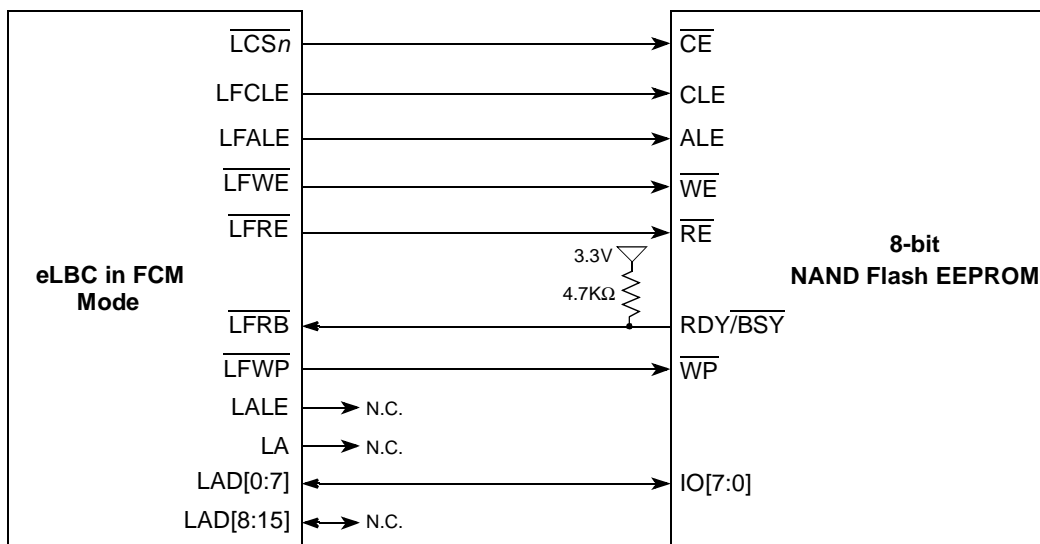


Figure 10-42. Local Bus to 8-Bit FCM Device Interface

Basic read access timing for FCM is shown in Figure 10-43. Although LCLK is shown for reference, NAND Flash EEPROMs do not make use of the clock.




Figure 10-43. FCM Basic Page Read Timing
(PGS = 1, CSCT = 0, CST = 0, CHT = 1, RST = 1, SCY = 0, TRLX = 0, EHTR = 1)

Following the assertion of LALE, FCM asserts \overline{LCSn} to commence a command sequence to the Flash device. After a delay of t_{CSCT} , the first command can be written to the device on assertion of \overline{LFWE} , followed by any parameters (typically address bytes and data), and concluded with a secondary command. In many cases, the second command initiates a long-running operation inside the Flash device, which pulls the wired-OR pin \overline{LFRB} low to indicate that the device is busy. Since in Figure 10-43 FCM is now expecting a read response, it takes LBCTL low to turnaround any bus transceivers that are present. Upon \overline{LFRB} indicating ready status, FCM asserts \overline{LFRE} repeatedly to recover bytes of read data, and the bytes are stored in eLBC’s FCM buffer RAM while an ECC is optionally computed on the bytes transferred. Finally, FCM negates \overline{LCSn} and delays eLBC by t_{EHTR} before any subsequent memory access occurs.

10.4.3.1 FCM Buffer RAM

Read and write accesses to eLBC banks controlled by FCM do not access attached NAND Flash EEPROMs directly. Rather, these accesses read and write the FCM buffer RAM—a single, shared 8-Kbyte space internal to eLBC and mapped by the base address of every FCM bank. Even though each FCM-controlled bank will have a different base address to differentiate it, all accesses to such banks will access the same buffer space. External eLBC signals, such as LALE and \overline{LCSn} , will not assert upon accesses to the buffer RAM. The FCM buffer RAM is logically divided into two or more buffers,

depending on the setting of $OR_n[PGS]$, with different buffers being accessible concurrently by software and FCM.

To perform a page read operation from a NAND Flash device, software initializes the FCM command, mode, and address registers, before issuing a special operation (FMR[OP] set non-zero) to a particular FCM-controlled bank. FCM will execute the sequence of op-codes held in FIR, reading data from the Flash device into the shared buffer RAM. While this read is taking place, software is free to access any data stored in other, currently inactive buffers of the FCM buffer RAM through reads or writes to any bank controlled by FCM. If command completion interrupts are enabled, an interrupt will be generated once FCM has completed the read. When FCM has completed its last command, software can switch to the newly read buffer and issue further commands.

To perform a page write operation, software first prepares data to be written in a fresh buffer. Then, the FCM command, mode, and address registers are initialized, and a special operation (FMR[OP] set non-zero) is issued to a particular FCM-controlled bank. FCM will execute the sequence of op-codes held in FIR, writing data from shared buffer RAM to the Flash device. To ensure that the device is enabled for programming, software must initialize $FMR[OP] = 11$, which prevents assertion of \overline{LFWP} during the write. While this write is taking place, software is free to access any data stored in other, currently inactive buffers of the FCM buffer RAM through reads or writes to any bank controlled by FCM. When FCM has completed its last command, software can re-use the previously written buffer and issue further commands.

See [Section 10.4.3.4.2, “Boot Block Loading into the FCM Buffer RAM,”](#) for a description of the shared buffer RAM layout during boot.

10.4.3.1.1 Buffer Layout and Page Mapping for Small-Page NAND Flash Devices

The FCM buffer space is divided into eight 1-Kbyte buffers for small-page devices ($OR_n[PGS] = 0$), mapped as shown in [Figure 10-44](#). Each page in a small-page NAND Flash comprises 528 bytes, where 512 bytes appear as main region data, and 16 bytes appear as spare region data. The EEPROM's page numbered P is associated with buffer number $(P \bmod 8)$, where $P = FPAR[PI]$. Since the bank size set by $OR_n[AM]$ will be greater than 8 Kbytes, an identical image of the FCM buffer RAM appears replicated every 8 Kbytes throughout the bank address space. It is recommended that the bank size be set to 32 Kbytes, which covers a single NAND Flash block for small-page devices.

For FCM commands, register FPAR sets the page address and, therefore, also the buffer number. In the case that $FBCR[BC] = 0$, FCM transfers an entire page, comprising the 512-byte main region followed by the 16-byte spare region; the 496-byte reserved region is not accessed, and remains undefined for software. However, for commands given a specific byte count in $FBCR[BC]$, $FPAR[MS]$ locates the starting address in either the main region ($MS = 0$) or the spare region ($MS = 1$). Where different eLBC banks control both

small and large-page devices, a large-page 4-Kbyte buffer must be assigned to either the first 4 or last 4 small-page buffers.




Figure 10-44. FCM Buffer RAM Memory Map for Small-Page (512-byte page) NAND Flash Devices

10.4.3.1.2 Buffer Layout and Page Mapping for Large-Page NAND Flash Devices

The FCM buffer space is divided into two 4 Kbyte buffers for large-page devices ($ORn[PGS] = 1$), mapped as shown in [Figure 10-45](#). Each page in a large-page NAND Flash comprises 2112 bytes, where 2048 bytes appear as main region data, and 64 bytes appear as spare region data. The EEPROM’s page numbered P is associated with buffer number $(P \bmod 2)$, where $P = FPAR[PI]$. Since the bank size set by $ORn[AM]$ will be greater than 8 Kbytes, an identical image of the FCM buffer RAM appears replicated every 8 Kbytes throughout the bank address space. It is recommended that the bank size be set to 256 Kbytes, which covers a single NAND Flash block for large-page devices.

For FCM commands, register FPAR sets the page address and, therefore, also the buffer number. In the case that $FBCR[BC] = 0$, FCM transfers an entire page, comprising the 2048-byte main region followed by the 64-byte spare region; the 1984-byte reserved region is not accessed, and remains undefined for software. However, for commands given a specific byte count in $FBCR[BC]$, $FPAR[MS]$ locates the starting address in either the main region ($MS = 0$) or the spare region ($MS = 1$). Where different eLBC

banks control both small and large-page devices, a large-page 4 Kbyte buffer must be assigned to either the first 4 or last 4 small-page buffers.




Figure 10-45. FCM Buffer RAM Memory Map for Large-Page (2-Kbyte page) NAND Flash Devices

10.4.3.1.3 Error Correcting Codes and the Spare Region

The FCM’s ECC engine makes use of data in the NAND Flash spare region to store pre-computed ECC code words. ECC is calculated in a single pass over blocks of 512 bytes of data in the main region. The setting of FMR[ECCM] determines the location of the 24-bit ECC in the spare region.

The basic ECC algorithm is depicted in Figure 10-46. The stream of data bytes is considered to form a matrix having 8 columns (corresponding with the device bus IO[7:0] or IO[15:8]) and 512 rows (corresponding with each byte in the ECC block).

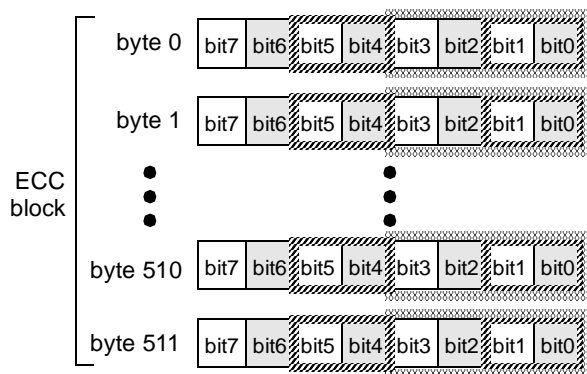


Figure 10-46. FCM ECC Calculation

The placement of ECC code words in relation to FMR[ECCM] is shown in Figure 10-47. For small-page devices, only a single 512-byte main region is ECC-protected. For large-page devices, there are four

adjacent main regions, and each has a 16-byte spare region—of which only one is shown in the figure. If eLBC is configured to generate ECC ($BR_n[DECC] = 10$), FCM will substitute on full-page write transfers the three code word bytes in place of the spare region data originally provided at the locations shown in [Figure 10-47](#). Transfers shorter than a full page, however, require software to prepare the appropriate ECC in the spare region. Similarly, FCM can check and correct bit errors on full-page reads if $BR_n[DECC] = 01$ or 10 . A correctable error is a single bit error in any 512-byte block of main region data, as judged by comparison of a regenerated ECC with the ECC retrieved from the spare region, or a single bit error in the retrieved ECC only. Bit errors in the main region are corrected before FCM completes its final read transfer and signals an event in $LTESR[CC]$. Errors that appear more complex (two or more bits in error per 512-byte block) are not corrected, but are flagged as parity errors by FCM. The bit vector in $LTEATR[PB]$ can be checked to determine which 512-byte blocks in a large-page NAND Flash main region were found to be non-correctable.

| ECCM | Byte 0 | Byte 511 | Other Mains | Spare 0 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|------|-------------|----------|-------------|---------|-----|-----|-----|-----|-----|-----|----|----|----|----|----|
| 0 | Main Region | | | — | EC0 | EC1 | EC2 | — | | | | | | | |
| 1 | Main Region | | | — | | | | EC0 | EC1 | EC2 | — | | | | |

Figure 10-47. ECC Placement in NAND Flash Spare Regions in Relation to FMR[ECCM]

10.4.3.2 Programming FCM

FCM has a fully general command and data transfer sequencer that caters for both common and specific/proprietary NAND Flash command sequences. The command sequencer reads a program out of the FIR register, which can hold up to 8 instructions, each represented by a 4-bit op-code, as illustrated in [Figure 10-48](#). The first instruction executed is read from $FIR[OP0]$, the next is read from $FIR[OP1]$, and likewise to subsequent instructions, ending at $FIR[OP7]$ or until the only instructions remaining are NOPs. If FIR contains nothing but NOP instructions, FCM will not assert \overline{LCSn} , otherwise, \overline{LCSn} is asserted prior to the first instruction and remains asserted until the last instruction has completed. If $LTESR[CC]$ is enabled, completion of the last instruction will trigger a command completion event interrupt from eLBC.

Prior to executing a sequence, necessary operands for the instructions will need to be set in the FMR, FCR, MDR, FBCR, FBAR, and FPAR registers. The AS0–AS3 address and data pointers associated with FCM’s use of MDR all reset to select AS0 at the start of the instruction sequence. A complete list of op-codes can be found in [Section 10.3.1.17, “Flash Instruction Register \(FIR\).”](#)




Figure 10-48. FCM Instruction Sequencer Mechanism

10.4.3.2.1 FCM Command Instructions

There are two kinds of command instruction:

- Commands that issue immediately—CM0, CM1, CM2, and CM3. These commands write a single command byte by asserting LFCLE and $\overline{\text{LFWE}}$ while driving an 8-bit command onto LAD[0:7]. Op-code CM n sources its command byte from field FCR[CMD n], therefore up to four different commands can be issued in any FCM instruction sequence.
- Commands that wait for $\overline{\text{LFRB}}$ to be sampled high (EEPROM in ready state) before issuing—CW0, and CW1. These commands first poll the $\overline{\text{LFRB}}$ pin, waiting for it to go high, before writing a single command byte onto LAD[0:7], sourced from FCR[CMD n] for op-code CW n . It is necessary to use CW n op-codes whenever the EEPROM is expected to be in a busy state (such as following a page read, block erase, or program operation) and therefore initially unresponsive to commands. To avoid deadlock in cases where the device is already available, FCM does not expect a transition on $\overline{\text{LFRB}}$. Rather, FCM waits for $8 \times (2 + \text{OR}_n[\text{SCY}])$ clock cycles (when $\text{OR}_n[\text{TRLX}] = 0$) or $16 \times (2 + \text{OR}_n[\text{SCY}])$ clock cycles (when $\text{OR}_n[\text{TRLX}] = 1$) before sampling the level of $\overline{\text{LFRB}}$. If the level of $\overline{\text{LFRB}}$ does not return high before a time-out set by FMR[CWTO] occurs, FCM proceeds to issue the command normally, and a FCT event is issued to LTESR.

The manufacturer’s datasheet should be consulted to determine values for programming into the FCR register, and whether a given command in the sequence is expected to initiate busy device behavior.

10.4.3.2.2 FCM No-Operation Instruction

A NOP instruction that appears in FIR ahead of the last instruction is executed with the timing of a regular command instruction, but neither LFCLE nor $\overline{\text{LFW}}\overline{\text{E}}$ are asserted. Thus a NOP instruction may be used to insert a pause matching the time taken for a regular command write.

10.4.3.2.3 FCM Address Instructions

Address instructions are used to issue addresses to the NAND Flash EEPROM. A complete device address is formed from a sequence of one or more bytes, each written onto LAD[0:7] with LFALE and $\overline{\text{LFW}}\overline{\text{E}}$ asserted together. There are three kinds of address generation provided:

- Column address—CA. A column address comprises one byte ($\text{OR}_n[\text{PGS}] = 0$) or two bytes ($\text{OR}_n[\text{PGS}] = 1$) locating the starting byte or word to be transferred in the next page read or write sequence. FPAR[CI] sets the value of the column index provided that FBCR[BC] is non-zero. In the case that FBCR[BC] = 0, a column index of zero is issued to the device, regardless of the value in FPAR[CI].
- Page address—PA. A page address comprises 2, 3, or 4 bytes, depending on the setting of FMR[AL], and locates the data page in the NAND Flash address space. The complete page address is the concatenation of the block index, read from FBAR[BLK], with the page-in-block index, read from FPAR[PI]. The page address length set in FMR[AL] should correspond with the size of EEPROM being accessed. Similarly, the block index in FBAR[BLK] must not exceed the maximum block index for the device, as most devices require reserved address bits to be written as zero.
- User-defined address—UA. This instruction allows the FCM to write a user-defined address byte, which is read from the next AS field in MDR, starting at MDR[AS0]. Each subsequent UA instruction reads an adjacent AS field in MDR, until all four AS bytes (MDR[AS0], MDR[AS1], MDR[AS2], MDR[AS3]) have been sent; a fifth and any following UA instructions send zero as the address byte. Note that each UA instruction advances the MDR pointer for writes by one byte, and therefore a mix of UA and WS instructions can consume adjacent bytes from MDR.

10.4.3.2.4 FCM Data Read Instructions

Data read instructions assert $\overline{\text{LFR}}\overline{\text{E}}$ repeatedly to transfer one or more bytes of read data from the NAND Flash EEPROM. Data read instructions are distinguished by their data destination:

- Read data to buffer RAM immediately—RB. This instruction reads FBCR[BC] bytes of data into the current FCM RAM buffer addressed by FPAR. If FBCR[BC] = 0, an entire page (including spare region) is transferred in a burst, starting at the page boundary, and the ECC calculation is checked against the ECC stored in the spare region. Correctable ECC errors are corrected; other errors may cause an interrupt if enabled. If the value of FBCR[BC] takes the read pointer beyond the end of the spare region in the buffer, FCM discards any excess bytes read.
- Read data/status to MDR immediately—RS. This instruction asserts $\overline{\text{LFR}}\overline{\text{E}}$ exactly once to read one byte (8-bit port size) of data into the next AS field of MDR. Reads beyond the fourth byte of MDR are discarded. The MDR read pointer is independent of the MDR write pointer used by UA and WS instructions.

- Read data to buffer RAM once waited on ready—RBW. This instruction first polls the $\overline{\text{LFRB}}$ pin, waiting for it to go high, before proceeding with a read to buffer as described for the RB instruction. Sampling and time-outs for polling the $\overline{\text{LFRB}}$ pin follow the behavior of CW_n instructions.
- Read data/status to MDR once waited on ready—RSW. This instruction first polls the $\overline{\text{LFRB}}$ pin, waiting for it to go high, before proceeding with a status read to MDR as described for the RS instruction. Sampling and time-outs for polling the $\overline{\text{LFRB}}$ pin follow the behavior of CW_n instructions.

10.4.3.2.5 FCM Data Write Instructions

Data write instructions assert $\overline{\text{LFW\!E}}$ repeatedly (with LFCLE and LFALE both negated) to transfer one or more bytes of write data to the NAND Flash EEPROM. Data write instructions are distinguished by their data source:

- Write data from FCM buffer RAM—WB. This instruction writes $\text{FBCR}[\text{BC}]$ bytes of data from the current FCM RAM buffer addressed by FPAR . If $\text{FBCR}[\text{BC}] = 0$, an entire page (including spare region) is transferred in a burst, starting at the page boundary, and the ECC calculation is stored in the and spare region in accordance with the setting of $\text{FMR}[\text{ECCM}]$. If the value of $\text{FBCR}[\text{BC}]$ takes the write pointer beyond the end of the spare region in the buffer, the value of data written by FCM is undefined.
- Write data/status from MDR—WS. This instruction asserts $\overline{\text{LFW\!E}}$ exactly once to write one byte (8-bit port size) of data taken from the next AS field of MDR. Attempts to write beyond four bytes of MDR has the effect of writing zeros. The MDR write pointer is independent of the MDR read pointer used by RS and RSW instructions.

10.4.3.3 FCM Signal Timing

If $\text{BR}_n[\text{MSEL}]$ selects the FCM, the attributes for the memory cycle are taken from OR_n . These attributes include the CSCT , CST , CHT , RST , SCY , TRLX , and EHTR fields.

10.4.3.3.1 FCM Chip-Select Timing

The timing of $\overline{\text{LCS}}_n$ assertion in FCM mode is illustrated by the timing diagram in [Figure 10-43](#). $\overline{\text{LCS}}_n$ is asserted immediately following LALE negation, and remains asserted until the last instruction in FIR has completed. The delay, t_{CSCT} , between $\overline{\text{LCS}}_n$ assertion and commencement of the first NAND Flash instruction is controlled by $\text{OR}_n[\text{CSCT}]$ and $\text{OR}_n[\text{TRLX}]$, as shown in [Table 10-33](#). $\text{OR}_n[\text{CSCT}]$ should be set in accordance with the NAND Flash EEPROM chip-select to $\overline{\text{WE}}$ set-up time specification.

Table 10-33. FCM Chip-Select to First Command Timing

| $\text{OR}_n[\text{TRLX}]$ | $\text{OR}_n[\text{CSCT}]$ | $\overline{\text{LCS}}_n$ to First Command Delay |
|----------------------------|----------------------------|--|
| 0 | 0 | 1 LCLK clock cycle |
| 0 | 1 | 4 LCLK clock cycles |
| 1 | 0 | 2 LCLK clock cycles |
| 1 | 1 | 8 LCLK clock cycles |

10.4.3.3.2 FCM Command, Address, and Write Data Timing

The FCM command (CM0–CM3, CW0, CW1), address (CA, PA, UA), and data write (WB, WS) instructions all share the same basic timing attributes. Assertion of $\overline{\text{LFWE}}$ initiates transfer via LAD[0:7], and the options in OR_n for FCM mode establish the set-up, hold, and wait state timings with respect to $\overline{\text{LFWE}}$, as shown in Figure 10-49.




Figure 10-49. Timing of FCM Command/Address and Write Data Cycles
 (for $\text{TRLX} = 0$, $\text{CHT} = 0$, $\text{CST} = 1$, $\text{SCY} = 1$, $\text{CLKDIV} = 4 \times \text{N}$)

The timing parameters are summarized in Table 10-34.

Table 10-34. FCM Command, Address, and Write Data Timing Parameters

| Option Register Attributes | | | Timing Parameter (LCLK Clock Cycles) ¹ | | | | | |
|----------------------------|-----|-----|---|------------------|-----------------------|-------------------------------------|---------------------------|-----------------------------|
| TRLX | CHT | CST | t_{CST} | t_{CHT} | t_{WS} | t_{WP} | t_{WC} | t_{ADL} |
| 0 | 0 | 0 | 0 | $\frac{1}{2}$ | SCY | $1\frac{1}{2} + \text{SCY}$ | $2 + \text{SCY}$ | $4 \times (2 + \text{SCY})$ |
| 0 | 0 | 1 | $\frac{1}{4}$ | $\frac{1}{2}$ | SCY | $1\frac{1}{4} + \text{SCY}$ | $2 + \text{SCY}$ | $4 \times (2 + \text{SCY})$ |
| 0 | 1 | 0 | 0 | 1 | SCY | $1 + \text{SCY}$ | $2 + \text{SCY}$ | $4 \times (2 + \text{SCY})$ |
| 0 | 1 | 1 | $\frac{1}{4}$ | 1 | SCY | $\frac{3}{4} + \text{SCY}$ | $2 + \text{SCY}$ | $4 \times (2 + \text{SCY})$ |
| 1 | 0 | 0 | $\frac{1}{2}$ | $1\frac{1}{2}$ | $2 \times \text{SCY}$ | $1 + 2 \times \text{SCY}$ | $3 + 2 \times \text{SCY}$ | $8 \times (2 + \text{SCY})$ |
| 1 | 0 | 1 | 1 | $1\frac{1}{2}$ | $2 \times \text{SCY}$ | $\frac{1}{2} + 2 \times \text{SCY}$ | $3 + 2 \times \text{SCY}$ | $8 \times (2 + \text{SCY})$ |
| 1 | 1 | 0 | $\frac{1}{2}$ | 2 | $2 \times \text{SCY}$ | $\frac{1}{2} + 2 \times \text{SCY}$ | $3 + 2 \times \text{SCY}$ | $8 \times (2 + \text{SCY})$ |
| 1 | 1 | 1 | 1 | 2 | $2 \times \text{SCY}$ | $2 \times \text{SCY}$ | $3 + 2 \times \text{SCY}$ | $8 \times (2 + \text{SCY})$ |

¹ In the parameters, SCY refers to a delay of $\text{OR}_n[\text{SCY}]$ clock cycles.

An example of minimum delay command timing appears in [Figure 10-50](#). Note that the set-up, wait-state, and hold timing of command, address, and write data cycles with respect to $\overline{\text{LFW}}\overline{\text{E}}$ assertion are all identical, and that the minimum cycle extends for two LCLK clock cycles.




Figure 10-50. Example of FCM Command and Address Timing with Minimum Delay Parameters (for $\text{TRLX} = 0$, $\text{CHT} = 0$, $\text{CST} = 0$, $\text{SCY} = 0$, $\text{CLKDIV} = 4 \cdot \text{N}$)

An example of relaxed command timing is shown in [Figure 10-51](#).




Figure 10-51. Example of FCM Command and Address Timing with Relaxed Parameters (for $\text{TRLX} = 1$, $\text{CHT} = 0$, $\text{CST} = 1$, $\text{SCY} = 2$, $\text{CLKDIV} = 4 \cdot \text{N}$)

10.4.3.3.3 FCM Ready/Busy Timing

Instructions CW0, CW1, RBW, and RSW force FCM to observe the state of the $\overline{\text{LFR}}\overline{\text{B}}$ pin, which may be driven low by a long-latency NAND Flash operation, such as a page read. Following the issue of such commands, FCM waits as shown in [Figure 10-52](#) before sampling the state of $\overline{\text{LFR}}\overline{\text{B}}$. This guards against observing $\overline{\text{LFR}}\overline{\text{B}}$ before it has been properly driven low by the device, but does not preclude $\overline{\text{LFR}}\overline{\text{B}}$ from

remaining high after a command. In addition, FCM samples and compares the state of $\overline{\text{LFRB}}$ on two consecutive cycles of LCLK to filter out noise on this signal as it rises to the ready state ($\overline{\text{LFRB}} = 1$).

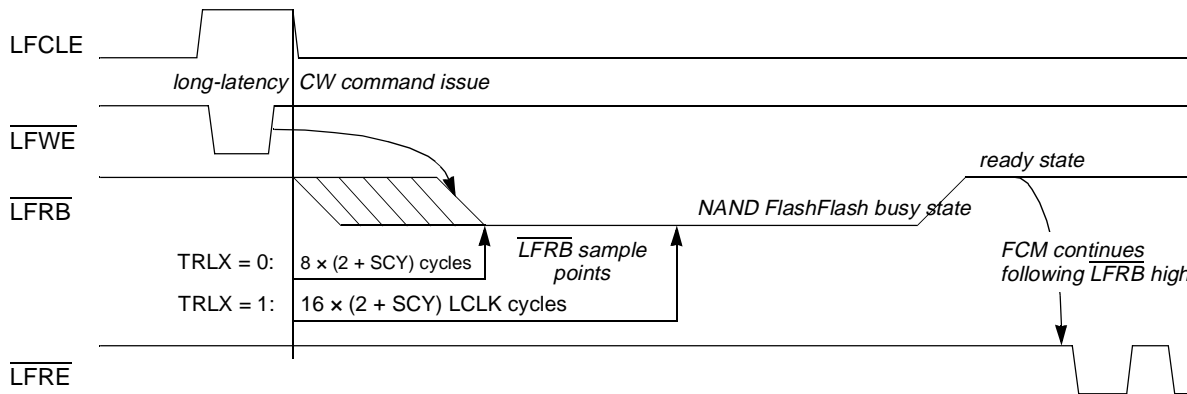
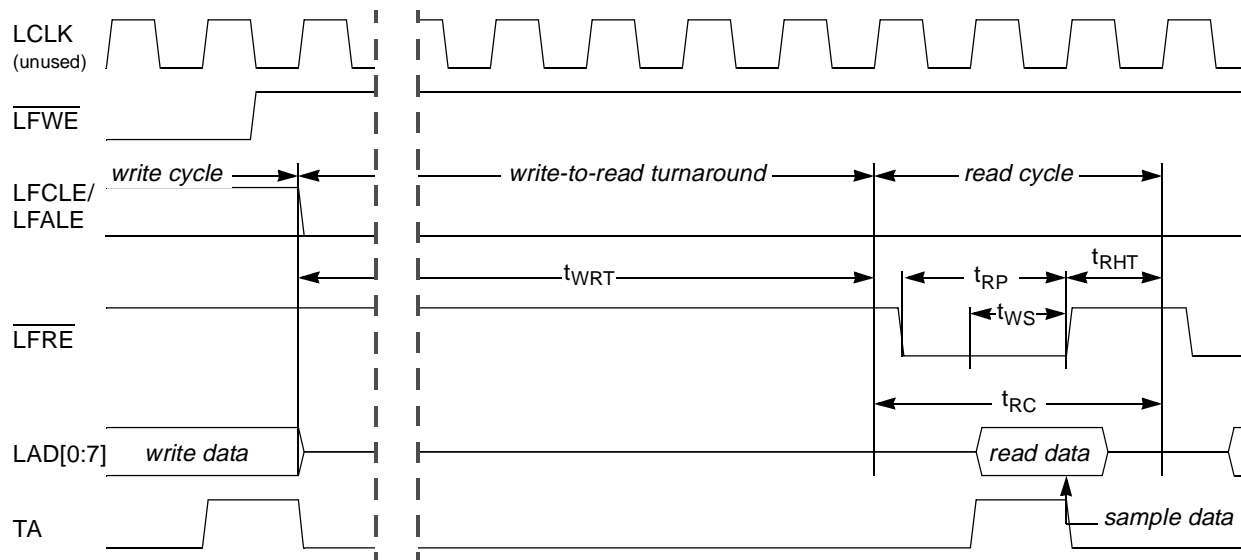


Figure 10-52. FCM Delay Prior to Sampling $\overline{\text{LFRB}}$ State

10.4.3.3.4 FCM Read Data Timing

The timing for read data transfers is shown in Figure 10-53. Upon assertion of $\overline{\text{LFRĒ}}$, the Flash device will enable its output drivers and drive valid read data while $\overline{\text{LFRĒ}}$ is held low. FCM samples read data on the rising edge of $\overline{\text{LFRĒ}}$, which follows an optional number of wait states. Note that FCM will delay the first read if a RBW or RSW instruction is issued, in which case $\overline{\text{LFRB}}$ sample timing takes effect (see Section 10.4.3.3.3, “FCM Ready/Busy Timing”).



Notes: t_{RP} = $\overline{\text{LFRĒ}}$ pulse time, read period. t_{WS} = Read wait state time.
 t_{RHT} = $\overline{\text{LFRĒ}}$ hold time. t_{RC} = Read data cycle time.
 t_{WRT} = Write to read turnaround time.

Figure 10-53. FCM Read Data Timing
 (for $\text{TRLX} = 0$, $\text{RST} = 0$, $\text{SCY} = 1$, $\text{CLKDIV} = 4 \cdot N$)

The timing parameters are summarized in [Table 10-35](#).

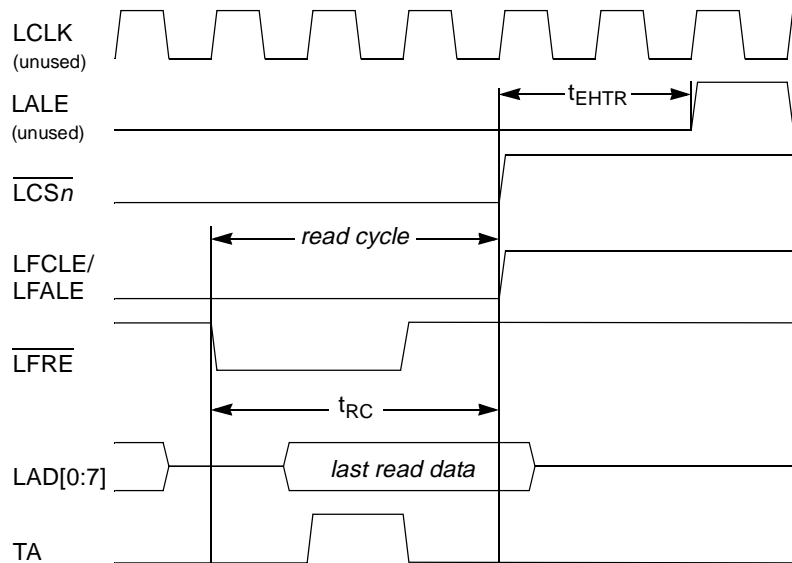
Table 10-35. FCM Read Data Timing Parameters

| Option Register Attributes | | Timing Parameter (LCLK Clock Cycles) ¹ | | | | |
|----------------------------|-----|---|------------------|-----------------|-----------------|------------------|
| TRLX | RST | t _{RP} | t _{RHT} | t _{WS} | t _{RC} | t _{WRT} |
| 0 | 0 | ¾ + SCY | 1 | SCY | 2 + SCY | 4 × (2 + SCY) |
| 0 | 1 | 1 + SCY | 1 | SCY | 2 + SCY | 4 × (2 + SCY) |
| 1 | 0 | ½ + 2 × SCY | 2 | 2 × SCY | 3 + 2 × SCY | 8 × (2 + SCY) |
| 1 | 1 | 1 + 2 × SCY | 2 | 2 × SCY | 3 + 2 × SCY | 8 × (2 + SCY) |

¹ In the parameters, SCY refers to a delay of OR_n[SCY] clock cycles.

10.4.3.3.5 FCM Extended Read Hold Timing

Allowance for slow output driver turn-off when reading NAND Flash EEPROMs is made via setting of OR_n[EHTR] and OR_n[TRLX]. The extended read data hold time, shown at t_{EHTR} in [Figure 10-43](#) and [Figure 10-54](#), is a delay inserted by FCM between the last data read and another eLBC memory access (requiring LALE assertion). \overline{LCSn} is negated during t_{EHTR} to allow external devices and bus transceivers time to disable their drivers.



Notes: t_{RC} = Read data cycle time.
t_{EHTR} = Extended read data hold time.

Figure 10-54. FCM Read Data Timing with Extended Hold Time
(for TRLX = 0, EHTR = 1, RST = 1, SCY = 1, CLKDIV = 4*N)

10.4.3.4 FCM Boot Chip-Select Operation

Boot chip-select operation allows address decoding for a boot ROM before system initialization. $\overline{LCS0}$ is the boot chip-select output; its operation differs from other external chip-select outputs after a system reset.

When the core begins accessing memory after system reset, $\overline{\text{LCS0}}$ is asserted initially to load a 4-Kbyte boot block into the FCM buffer RAM, but core instruction fetches occur from the buffer RAM.

10.4.3.4.1 FCM Bank 0 Reset Initialization

The boot chip-select also provides a programmable port size, which is configured during reset. The boot chip-select does not provide write protection. $\overline{\text{LCS0}}$ operates this way until the first write to OR0 and it can be used as any other chip-select register after the preferred address range is loaded into BR0. After the first write to OR0, the boot chip-select can be restarted only with a hardware reset. Table 10-36 describes the initial values of the boot bank in the memory controller.

Table 10-36. Boot Bank Field Values after Reset for FCM as Boot Controller

| Register | Field | Setting |
|----------|-------|-----------------------|
| BR0 | BA | 0000_0000_0000_0000_0 |
| | PS | From 01 |
| | DECC | 00 |
| | WP | 0 |
| | MSEL | 001 |
| | V | 0 |
| OR0 | AM | 0000_0000_0000_0000_0 |
| | BCTLD | 0 |
| | PGS | From RCWH[ROMLOC] |
| | CSCT | 1 |
| | CST | 1 |
| | CHT | 1 |
| | RST | 1 |
| | SCY | From por_cfg_scy[1:3] |
| | TRLX | 1 |
| | EHTR | 1 |

10.4.3.4.2 Boot Block Loading into the FCM Buffer RAM

If FCM is selected as the boot ROM controller from power-on-reset configuration, eLBC will automatically load from bank 0 a single 4 Kbyte page of boot code into the FCM buffer RAM during HRESET. The CPU can execute boot code directly from the FCM buffer RAM, but must ensure that any further data read from the NAND Flash EEPROM is transferred under software control in order to continue the bootstrap process.

Since OR0[AM] is initially cleared during reset, all CPU fetches to eLBC will access the FCM buffer RAM, which appears in the memory map as a 4-Kbyte RAM. No NAND Flash spare regions are mapped

during boot, therefore only 4 Kbytes of contiguous, main region data, loaded from the first pages of the boot block, are accessible in eLBC bank 0, as indicated in [Figure 10-55](#).




Figure 10-55. FCM Buffer RAM Memory Map During Boot Loading

The process for booting is as follows:

1. Following negation of $\overline{\text{PORESET}}$, eLBC is released from reset and commences automatic boot block loading if FCM is selected as the boot ROM location. Small-page or large-page, 8-bit NAND Flash devices can be used for boot loading when enabled with $\overline{\text{LCS0}}$. eLBC drives $\overline{\text{LFWP}}$ low during boot accesses to prevent accidental erasure of the NAND Flash boot ROM.
2. FCM starts searching for a valid boot block at block index 0.
3. FCM reads the spare regions of the first two pages of the current block, checking the bad block indication (BI) bytes to validate the block for reading. BI bytes must all hold the value 0xFF for the page to be considered readable.
 - For small-page devices, BI is a single byte read from spare region byte offset 5.
 - For large-page devices, BI is a single byte read from spare region byte offset 0.

If either of the first two pages of the current block are marked invalid, then the boot block index is incremented by 1, and FCM repeats step 3. eLBC will continue searching for a bootable block indefinitely, therefore at least one block must be marked valid for boot loading to proceed. At the conclusion of the boot block search, the value of $\text{FBAR}[\text{BLK}]$ points to the boot block.

4. If ECC checking is enabled, the FCM recovers from the spare region the stored ECC for each 512-byte block of boot data. The boot block must be prepared with ECC protection. During ECC generation, software should use $\text{FMR}[\text{ECCM}] = 0$ for small-page devices, and $\text{FMR}[\text{ECCM}] = 1$ for large-page devices.

If RCW initialization is required, the first 64 bytes of the boot block must be prepared in accordance with the layout described in [Section 4.3.3.1.1, “Local Bus Controller Setting.”](#)

5. FCM performs a sequence of random-access page reads, reading entire pages from the boot block until 4 Kbytes have been saved to the FCM buffer RAM. If ECC checking is enabled, the ECC of each 512-byte region is verified and single-bit errors are corrected if possible. If FCM is unable to correct ECC errors, eLBC halts the boot process and signals an unrecoverable error by asserting the $\overline{\text{hreset_req}}$ signal.

- The CPU now commences fetching instructions, in random order, from the FCM buffer RAM. This first-level boot loader typically copies a secondary boot loader into system memory, and continues booting from there. Boot software must clear FMR[BOOT] to enable normal operation of FCM.

10.4.4 User-Programmable Machines (UPMs)

UPMs are flexible interfaces that connect to a wide range of memory devices. At the heart of each UPM is an internal RAM array that specifies the logical value driven on the external memory control signals (\overline{LCSn} , $\overline{LBS}[0:1]$ and $\overline{LGPL}[0:5]$) for a given clock cycle. Each word in the RAM array provides bits that allow a memory access to be controlled with a resolution of up to one quarter of the external bus clock period on the byte-select and chip-select lines. A gap of 2 dead LCLK cycles is present on the UPM interface between UPM transactions.

NOTE

If the $\overline{LGPL4}/\overline{LGTA}/\overline{LFRB}/\overline{LUPWAIT}$ signal is used as both an input and an output, a weak pull-up is required. Refer to the hardware specification for details regarding termination options.

Figure 10-56 shows the basic operation of each UPM.




Figure 10-56. User-Programmable Machine Functional Block Diagram

The following events initiate a UPM cycle:

- Any internal device requests an external memory access to an address space mapped to a chip-select serviced by the UPM
- A UPM refresh timer expires and requests a transaction, such as a DRAM refresh
- A bus monitor time-out error during a normal UPM cycle redirects the UPM to execute an exception sequence

The RAM array contains 64 words of 32-bits each. The signal timing generator loads the RAM word from the RAM array to drive the general-purpose lines, byte-selects, and chip-selects. If the UPM reads a RAM word with WAEN set, the external LUPWAIT signal is sampled and synchronized by the memory controller and the current request is frozen.

10.4.4.1 UPM Requests

A special pattern location in the RAM array is associated with each of the possible UPM requests. An internal device's request for a memory access initiates one of the following patterns (MxMR[OP] = 00):

- Read single-beat pattern (RSS)
- Read burst cycle pattern (RBS)
- Write single-beat pattern (WSS)
- Write burst cycle pattern (WBS)

A UPM refresh timer request pattern initiates a refresh timer pattern (RTS).

An exception (caused by a bus monitor time-out error) occurring while another UPM pattern is running initiates an exception condition pattern (EXS).

Figure 10-57 and Table 10-37 show the start addresses of these patterns in the UPM RAM, according to cycle type. RUN commands (MxMR[OP] = 11), however, can initiate patterns starting at any of the 64 UPM RAM words.

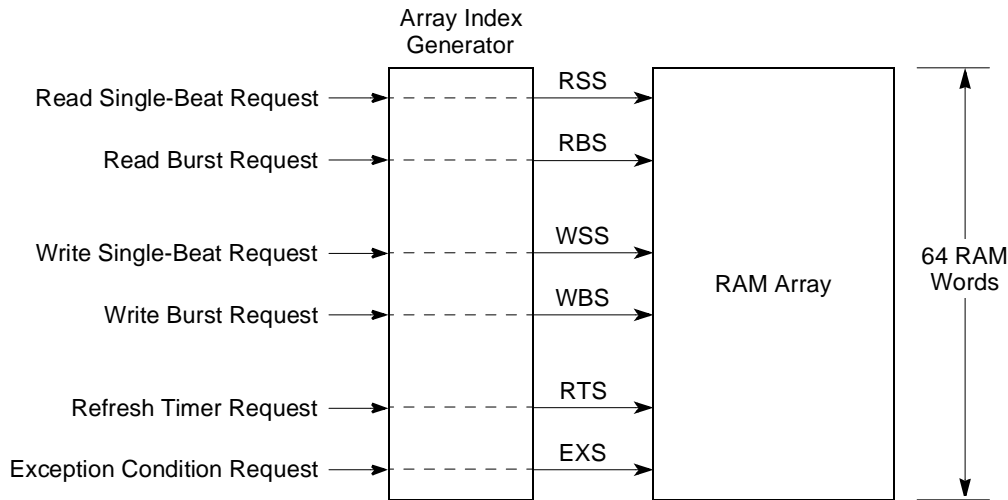


Figure 10-57. RAM Array Indexing

Table 10-37. UPM Routines Start Addresses

| UPM Routine | Routine Start Address |
|-------------------------|-----------------------|
| Read single-beat (RSS) | 0x00 |
| Read burst (RBS) | 0x08 |
| Write single-beat (WSS) | 0x18 |
| Write burst (WBS) | 0x20 |

Table 10-37. UPM Routines Start Addresses (continued)

| UPM Routine | Routine Start Address |
|---------------------------|-----------------------|
| Refresh timer (RTS) | 0x30 |
| Exception condition (EXS) | 0x3C |

10.4.4.1.1 Memory Access Requests

The user must ensure that the UPM is appropriately initialized before a request occurs.

The UPM supports two types of memory reads and writes:


- A single-beat transfer transfers one operand consisting of up to a single word (dependent on port size). A single-beat cycle starts with one transfer start and ends with one transfer acknowledge.
- A burst transfer transfers exactly 4 double words regardless of port size. For 32-bit accesses, the burst cycle starts with one transfer start but ends after eight transfer acknowledges, whereas an 8-bit device requires 32 transfer acknowledges.

The user must ensure that patterns for single-beat transfers contain one, and only one, transfer acknowledge (UTA bit in RAM word set high) and for a burst transfer, contain the exact number of transfer acknowledges required.

Any transfers that do not naturally fit single or burst transfers are synthesized as a series of single transfers. These accesses are treated by the UPM as back-to-back, single-beat transfers. Burst transfers can also be inhibited by setting $OR_n[BI]$. Burst performance can be achieved by ensuring that UPM transactions are 32-byte aligned with a transaction size being some multiple of 32-bytes, which is a natural fit for cache-line transfers, for example.

10.4.4.1.2 UPM Refresh Timer Requests

Each UPM contains a refresh timer that can be programmed to generate refresh service requests of a particular pattern in the RAM array. [Figure 10-58](#) shows the clock division hardware associated with memory refresh timer request generation. The UPM refresh timer register (LURT) defines the period for the timers associated with all three UPMs.


Figure 10-58. Memory Refresh Timer Request Block Diagram

By default, all local bus refreshes are performed using the refresh pattern of UPMA. This means that if refresh is required, $MAMR[RFEN]$ must be set. It also means that only one refresh routine should be programmed and be placed in UPMA, which serves as the refresh executor. Any banks assigned to a UPM are provided with the common UPMA refresh pattern if the $RFEN$ bit of the corresponding UPM is set, concurrently. UPMA assigned banks, therefore, always receive refresh services when $MAMR[RFEN]$ is set, while UPMB and UPMC assigned banks also receive (the same) refresh services if the corresponding $MxMR[RFEN]$ bits are set. In this scenario, more than one chip select may assert at the same time, as refresh pattern runs for all banks assigned to UPM with $RFEN$ bit set.

10.4.4.1.3 Software Requests—RUN Command

Software can start a request to the UPM by issuing a RUN command to the UPM. Some memory devices have their own signal handshaking protocol to put them into special modes, such as self-refresh mode.

For these special cycles, the user creates a special RAM pattern that can be stored in any unused areas in the UPM RAM. Then a RUN command is used to run the cycle. The UPM runs the pattern beginning at the specified RAM location until it encounters a RAM word with its LAST bit set. The RUN command is issued by setting $MxMR[OP] = 11$ and accessing UPM n memory region with any write transaction that hits the corresponding UPM machine. $MxMR[MAD]$ determines the starting address in the RAM array for the pattern.

Note that transfer acknowledges (UTA bit in the RAM word) are ignored for software (RUN command) requests, and hence the LAD signals remain high-impedance unless the normal initial LALE occurs or the RUN pattern causes assertion of LALE to occur on changes to the RAM word AMX field.

10.4.4.1.4 Exception Requests

When the eLBC under UPM control initiates an access to a memory device and an exception occurs (bus monitor time-out), the UPM provides a mechanism by which memory control signals can meet the device's timing requirements without losing data. The mechanism is the exception pattern that defines how the UPM negates its signals in a controlled manner.

10.4.4.2 Programming the UPMs

The UPM is a micro sequencer that requires microinstructions or RAM words to generate signal timings for different memory cycles. Follow these steps to program the UPMs:

1. Set up BR n and OR n registers.
2. Write patterns into the RAM array.
3. Program MRTPR, LURT and MAMR[RFEN] if refresh is required.
4. Program MxMR.

Patterns are written to the RAM array by setting $MxMR[OP] = 01$ and accessing the UPM with any write transaction that hits the relevant chip select. The entire array is thus programmed by an alternating series of writes: to MDR (RAM word to be written) each time followed by a read from MDR and then followed by a (dummy) write transaction to the relevant UPM assigned bank. A read from MDR is required to ensure that the MDR update has occurred prior to the (dummy) write transaction.

RAM array contents may also be read for debug purposes, for example, by alternating dummy read transactions, each time followed by reads of MDR (when $MxMR[OP] = 10$).

NOTE

$MxMR/MDR$ registers should not be updated while dummy read/write access is still in progress. If the $MxMR[MAD]$ is incremented then the previous dummy transaction is already completed.

In order to enforce proper ordering between updates to the MxMR/MDR register and the dummy accesses to the UPM memory region, two rules must be followed:

- Since the result of any update to the MxMR/MDR register must be in effect before the dummy read or write to the UPM region, a write to MxMR/MDR should be followed immediately by a read of MxMR/MDR.
- The UPM memory region should have the same MMU settings as the memory region containing the MxMR configuration register; both should be mapped by the MMU as cache-inhibited and guarded. This prevents the CPU from re-ordering a read of the UPM memory around the read of MxMR. Once the programming of the UPM array is complete the MMU setting for the associated address range can be set to the proper mode for normal operation, such as cacheable and copyback.

For proper signalling, the following guidelines must be followed while programming UPM RAM words:

- For UPM reads, program UTA and LAST in the same or consecutive RAM words.
- For UPM burst reads, program last UTA and LAST in the same or consecutive RAM words.
- For UPM writes, program UTA and LAST in the same RAM word.
- For UPM burst writes, program last UTA and LAST in the same RAM word.

10.4.4.2.1 UPM Programming Example (Two Sequential Writes to the RAM Array)

The following example further illustrates the steps required to perform two writes to the RAM array at non-sequential addresses assuming that the relevant BR n and OR n registers have been previously set up:

1. Program MxMR for the first write (with the desired RAM array address).
2. Write pattern/data to MDR to ensure that the MxMR has already been updated with the desired configuration.
3. Read MDR to ensure that the MDR has already been updated with the desired pattern. (Or, read MxMR register if step 2 is not performed.)
4. Perform a dummy write transaction.
5. Read/check MxMR[MAD]. If incremented, the previous dummy write transaction is completed; proceed to step 6. Repeat step 5 until incremented.
6. Program MxMR for the second write with the desired RAM array address.
7. Write pattern/data to MDR to ensure that the MxMR has already been updated with the desired configuration.
8. Read MDR to ensure that the MDR has already been updated with the desired pattern.
9. Perform a dummy write transaction.
10. Read/check MxMR[MAD]. If incremented, the previous dummy write transaction is completed.

Note that if steps 1 and 2 or steps 6 and 7 are reversed, step 3 or 8 (as appropriate) is replaced by the following:

- Read MxMR to ensure that the MxMR has already been updated with the desired configuration.

10.4.4.2.2 UPM Programming Example (Two Sequential Reads from the RAM Array)

RAM array contents may also be read for debug purposes, for example, by alternating dummy read transactions, each time followed by reads of MDR (MxMR[OP] = 0b10). The following example further illustrates the steps required to perform two reads from the RAM array at non-sequential addresses assuming that the relevant BR_n and OR_n registers have been previously set up:

1. Program MxMR for the first read with the desired RAM array address.
2. Read MxMR to ensure that the MxMR has already been updated with the desired configuration, such as RAM array address.
3. Perform a dummy read transaction.
4. Read/check MxMR[MAD]. If incremented, the previous dummy read transaction is completed; proceed to step 5. Repeat step 4 until incremented.
5. Read MDR.
6. Program MxMR for the second read with the desired RAM array address.
7. Read MxMR to ensure that the MxMR has already been updated with the desired configuration, such as RAM array address.
8. Perform a dummy read transaction.
9. Read/check MxMR[MAD]. If incremented, the previous dummy read transaction is completed; proceed to step 10. Repeat step 9 until incremented.
10. Read MDR.

10.4.4.3 UPM Signal Timing

RAM word fields specify the value of the various external signals at a granularity of up to four values for each bus clock cycle. The signal timing generator causes external signals to behave according to timing specified in the current RAM word. For LCRR[CLKDIV] = 4 or 8, each bit in the RAM word relating to \overline{LCS}_n and \overline{LBS} timing specifies the value of the corresponding external signal at each quarter phase of the bus clock. If LCRR[CLKDIV] = 2, the external signal can change value only on each half phase of the bus clock. If the RAM word in this case (LCRR[CLKDIV] = 2) specifies a quarter phase signal change, the signal timing generator interprets this as a half cycle change.

The division of UPM bus cycles into phases is shown in [Figure 10-59](#) and [Figure 10-60](#). If LCRR[CLKDIV] = 2, the bus cycle comprises only two active phases, T1 and T3, which correspond with the first and second halves of the bus clock cycle, respectively. However, if LCRR[CLKDIV] = 4 or 8, four phases, T1–T4, define four quarters of the bus clock cycle. Because T2 and T4 are inactive when LCRR[CLKDIV] = 2, UPM ignores signal timing programmed for assertion in either of these phases in the case LCRR[CLKDIV] = 2.




Figure 10-59. UPM Clock Scheme for LCRR[CLKDIV] = 2




Figure 10-60. UPM Clock Scheme for LCRR[CLKDIV] = 4 or 8

10.4.4.4 RAM Array

The RAM array for each UPM is 64 locations deep and 32 bits wide, as shown in Figure 10-61. The signals at the bottom of the figure are UPM outputs. The selected \overline{LCS}_n is for the bank that matches the current address. The selected \overline{LBS} is for the byte lanes read or written by the access.

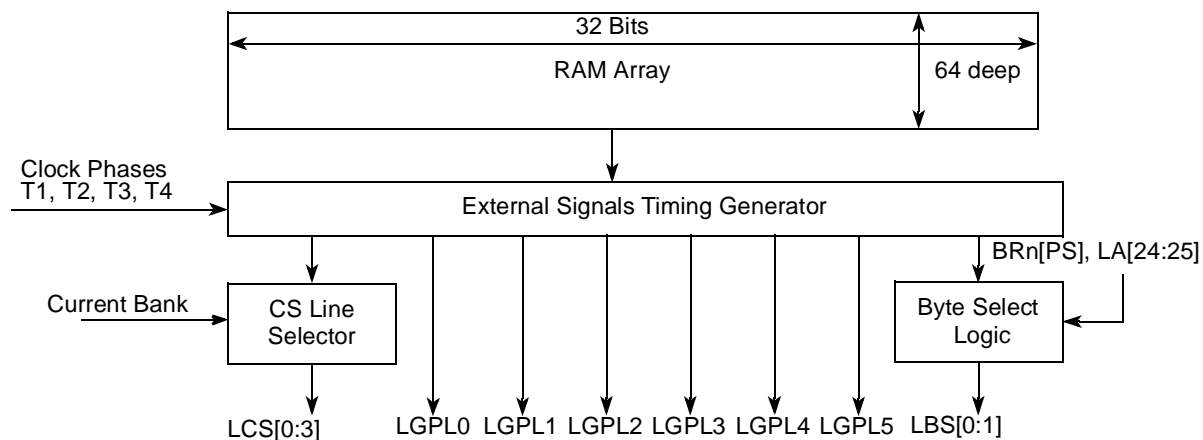


Figure 10-61. RAM Array and Signal Generation

10.4.4.4.1 RAM Words

The RAM word is a 32-bit microinstruction stored in one of 64 locations in the RAM array. It specifies timing for external signals controlled by the UPM. Figure 10-62 shows the RAM word fields. When

LCRR[CLKDIV] = 4 or 8, the CST_n and BST_n bits determine the state of UPM signals \overline{LCS}_n and $\overline{LBS}[0:1]$ at each quarter phase of the bus clock. When LCRR[CLKDIV] = 2, CST2 and CST4 are ignored and the external has the values defined by CST1 and CST3 but extended to half the clock cycle in duration. The same interpretation occurs for the BST_n bits when LCRR[CLKDIV] = 2.

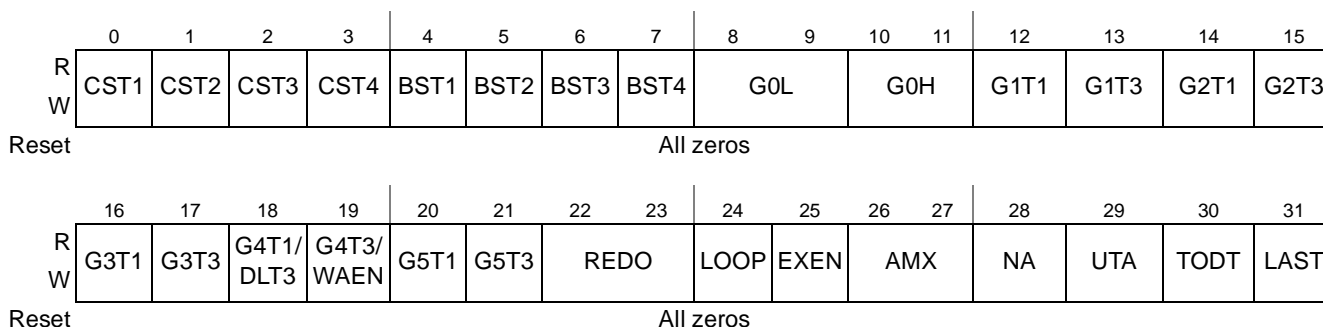


Figure 10-62. RAM Word Fields

Table 10-38 contains descriptions of the RAM word fields.

Table 10-38. RAM Word Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0 | CST1 | Chip select timing 1. Defines the state (0 or 1) of \overline{LCS}_n during bus clock quarter phase 1 if LCRR[CLKDIV] = 4 or 8. Defines the state (0 or 1) of \overline{LCS}_n during bus clock half phase 1 if LCRR[CLKDIV] = 2. |
| 1 | CST2 | Chip select timing 2. Defines the state (0 or 1) of \overline{LCS}_n during bus clock quarter phase 2 if LCRR[CLKDIV] = 4 or 8. Ignored when LCRR[CLKDIV] = 2. |
| 2 | CST3 | Chip select timing 3. Defines the state (0 or 1) of \overline{LCS}_n during bus clock quarter phase 3 if LCRR[CLKDIV] = 4 or 8. Defines the state (0 or 1) of \overline{LCS}_n during bus clock half phase 2 if LCRR[CLKDIV] = 2. |
| 3 | CST4 | Chip select timing 4. Defines the state (0 or 1) of \overline{LCS}_n during bus clock quarter phase 4 if LCRR[CLKDIV] = 4 or 8. Ignored when LCRR[CLKDIV] = 2. |
| 4 | BST1 | Byte select timing 1. Defines the state (0 or 1) of \overline{LBS} during bus clock quarter phase 1 (LCRR[CLKDIV] = 4 or 8) or bus clock half phase 1 (LCRR[CLKDIV] = 2), in conjunction with BR _n [PS] and LA[24:25]. |
| 5 | BST2 | Byte select timing 2. Defines the state (0 or 1) of \overline{LBS} during bus clock quarter phase 2 (LCRR[CLKDIV] = 4 or 8), in conjunction with BR _n [PS] and LA[24:25]. Ignored when LCRR[CLKDIV] = 2. |
| 6 | BST3 | Byte select timing 3. Defines the state (0 or 1) of \overline{LBS} during bus clock quarter phase 3 (LCRR[CLKDIV] = 4 or 8) or bus clock half phase 2 (LCRR[CLKDIV] = 2), in conjunction with BR _n [PS] and LA[24:25]. |
| 7 | BST4 | Byte select timing 4. Defines the state (0 or 1) of \overline{LBS} during bus clock quarter phase 4 (LCRR[CLKDIV] = 4 or 8), in conjunction with BR _n [PS] and LA[24:25]. Ignored when LCRR[CLKDIV] = 2. |
| 8–9 | G0L | General purpose line 0 lower. Defines the state of LGPL0 during the bus clock quarter phases 1 and 2 (first half phase). 00 Value defined by MxMR[G0CL] 01 Reserved 10 0 11 1 |

Table 10-38. RAM Word Field Descriptions (continued)

| Bits | Name | Description |
|-------|-----------|--|
| 10–11 | G0H | General purpose line 0 higher. Defines the state of LGPL0 during the bus clock quarter phases 3 and 4 (second half phase). 00 Value defined by MxMR[G0CL] 01 Reserved 10 0 11 1 |
| 12 | G1T1 | General purpose line 1 timing 1. Defines the state (0 or 1) of LGPL1 during bus clock quarter phases 1 and 2 (first half phase). |
| 13 | G1T3 | General purpose line 1 timing 3. Defines the state (0 or 1) of LGPL1 during bus clock quarter phases 3 and 4 (second half phase) |
| 14 | G2T1 | General purpose line 2 timing 1. Defines state (0 or 1) of LGPL2 during bus clock quarter phases 1 and 2 (first half phase). |
| 15 | G2T3 | General purpose line 2 timing 3. Defines the state (0 or 1) of LGPL2 during bus clock quarter phases 3 and 4 (second half phase). |
| 16 | G3T1 | General purpose line 3 timing 1. Defines the state (0 or 1) of LGPL3 during bus clock quarter phases 1 and 2 (first half phase). |
| 17 | G3T3 | General purpose line 3 timing 3. Defines the state (0 or 1) of LGPL3 during bus clock quarter phases 3 and 4 (second half phase). |
| 18 | G4T1/DLT3 | General purpose line 4 timing 1/delay time 3. The function of this bit is determined by MxMR[GPL4]. If MxMR[GPL4] = 0 and LGPL4/LUPWAIT pin functions as an output (LGPL4), G4T1/DLT3 defines the state (0 or 1) of LGPL4 during bus clock quarter phases 1 and 2 (first half phase). If MxMR[GPL4] = 1 and LGPL4/LUPWAIT functions as an input (LUPWAIT), if a read burst or single read is executed, G4T1/DLT3 defines the sampling of the data bus as follows: 0 In the current word, the data bus should be sampled at the start of bus clock quarter phase 1 of the next bus clock cycle. 1 In the current word, the data bus should be sampled at the start of bus clock quarter phase 3 of the current bus clock cycle. |
| 19 | G4T3/WAEN | General purpose line 4 timing 3/wait enable. Bit function is determined by MxMR[GPL4]. If MxMR[GPL4] = 0 and LGPL4/LUPWAIT pin functions as an output (LGPL4), G4T3/WAEN defines the state (0 or 1) of LGPL4 during bus clock quarter phases 3 and 4 (second half phase). If MxMR[GPL4] = 1 and LGPL4/LUPWAIT functions as an input (LUPWAIT), G4T3/WAEN is used to enable the wait mechanism: 0 LUPWAIT detection is disabled. 1 LUPWAIT is enabled. If LUPWAIT is detected as being asserted, a freeze in the external signals logical values occurs until LUPWAIT is detected as being negated. |
| 20 | G5T1 | General purpose line 5 timing 1. Defines the state (0 or 1) of LGPL5 during bus clock quarter phases 1 and 2 (first half phase). |
| 21 | G5T3 | General purpose line 5 timing 3. Defines the state (0 or 1) of LGPL5 during bus clock quarter phases 3 and 4 (second half phase). |
| 22–23 | REDO | Redo current RAM word. Defines the number of times to execute the current RAM word. 00 Once (normal operation) 01 Twice 10 Three times 11 Four times |

Table 10-38. RAM Word Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 24 | LOOP | <p>Loop start/end. The first RAM word in the RAM array where LOOP is 1 is recognized as the loop start word. The next RAM word where LOOP is 1 is the loop end word. RAM words between, and including the start and end words, are defined as part of the loop. The number of times the UPM executes this loop is defined in the corresponding loop fields of the MxMR.</p> <p>0 The current RAM word is not the loop start word or loop end word. 1 The current RAM word is the start or end of a loop.</p> <p>Note: AMX must not change values in any RAM word which begins a loop</p> |
| 25 | EXEN | <p>Exception enable. Allows branching to an exception pattern at the exception start address (EXS). When an internal bus monitor time-out exception is recognized and EXEN in the RAM word is set, the UPM branches to the special exception start address (EXS) and begins operating as the pattern defined there specifies.</p> <p>The user should provide an exception pattern to negate signals controlled by the UPM in a controlled fashion. For DRAM control, a handler should negate RAS and CAS to prevent data corruption. If EXEN = 0, exceptions are ignored by UPM (but not by local bus) and execution continues. After the UPM branches to the exception start address, it continues reading until the LAST bit is set in the RAM word.</p> <p>0 The UPM continues executing the remaining RAM words, ignoring any internal bus monitor time-out. 1 The current RAM word allows a branch to the exception pattern after the current cycle if an exception condition is detected.</p> |
| 26–27 | AMX | <p>Address multiplexing. Determines the source of LAD during an LALE phase. Any change in the AMX field initiates a new LALE (address) phase.</p> <p>00 LAD (and/or in conjunction with LA) is the non-multiplexed address. For example, column address. 01 Reserved 10 LAD (and/or in conjunction with LA) is driven with the multiplexed address according to MxMR[AM]. For example, row address. See Section 10.4.4.4.7, “Address Multiplexing (AMX)” for more information. 11 LAD (and/or in conjunction with LA) is driven with the contents of MAR. Used, for example, to initialize a mode.</p> <p>Note: Source ID debug mode is only supported for the AMX = 00 setting. Note: AMX must not change values in any RAM word which begins a loop.</p> |
| 28 | NA | <p>Next burst address. Determines when the address is incremented during a burst access.</p> <p>0 The address increment function is disabled. 1 The address is incremented in the next cycle. In conjunction with the BRn[PS], the increment value of LAN is 1 or 2 for port sizes of 8 and 16 bits, respectively.</p> |
| 29 | UTA | <p>UPM transfer acknowledge. Indicates assertion of transfer acknowledge in the current cycle.</p> <p>0 Transfer acknowledge is not asserted in the current cycle. 1 Transfer acknowledge is asserted in the current cycle.</p> <p>In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.</p> |
| 30 | TODT | <p>Turn-on disable timer. The disable timer associated with each UPM allows a minimum time to be guaranteed between two successive accesses to the same memory bank. This feature is critical when DRAM requires a RAS precharge time. TODT turns the timer on to prevent another UPM access to the same bank until the timer expires. The disable timer period is determined in MxMR[DSn]. The disable timer does not affect memory accesses to different banks. Note that TODT must be set together with LAST, otherwise it is ignored.</p> <p>0 The disable timer is turned off. 1 The disable timer for the current bank is activated preventing a new access to the same bank (when controlled by the UPMs) until the disable timer expires. For example, precharge time.</p> |

Table 10-38. RAM Word Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 31 | LAST | <p>Last word. When LAST is read in a RAM word, the current UPM pattern terminates and control signal timing set in the RAM word is applied to the current (and last) cycle. However, if the disable timer is activated and the next access is to the same bank, execution of the next UPM pattern is held off and the control signal values specified in the last word are extended in duration for the number of clock cycles specified in MxMR[DS_n].</p> <p>0 The UPM continues executing RAM words. 1 Indicates the last RAM word in the program. The service to the UPM request is done after this cycle concludes.</p> <p>In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.</p> |

10.4.4.4.2 Chip-Select Signal Timing (CST_n)

If BR_n[MSEL] of the accessed bank selects a UPM on the currently requested cycle, the UPM manipulates the \overline{LCS}_n for that bank with timing as specified in the UPM RAM word CST_n fields. The selected UPM affects only the assertion and negation of the appropriate \overline{LCS}_n signal. The state of the selected \overline{LCS}_n signal of the corresponding bank depends on the value of each CST_n bit. [Figure 10-63](#) shows how UPMs control \overline{LCS}_n signals.

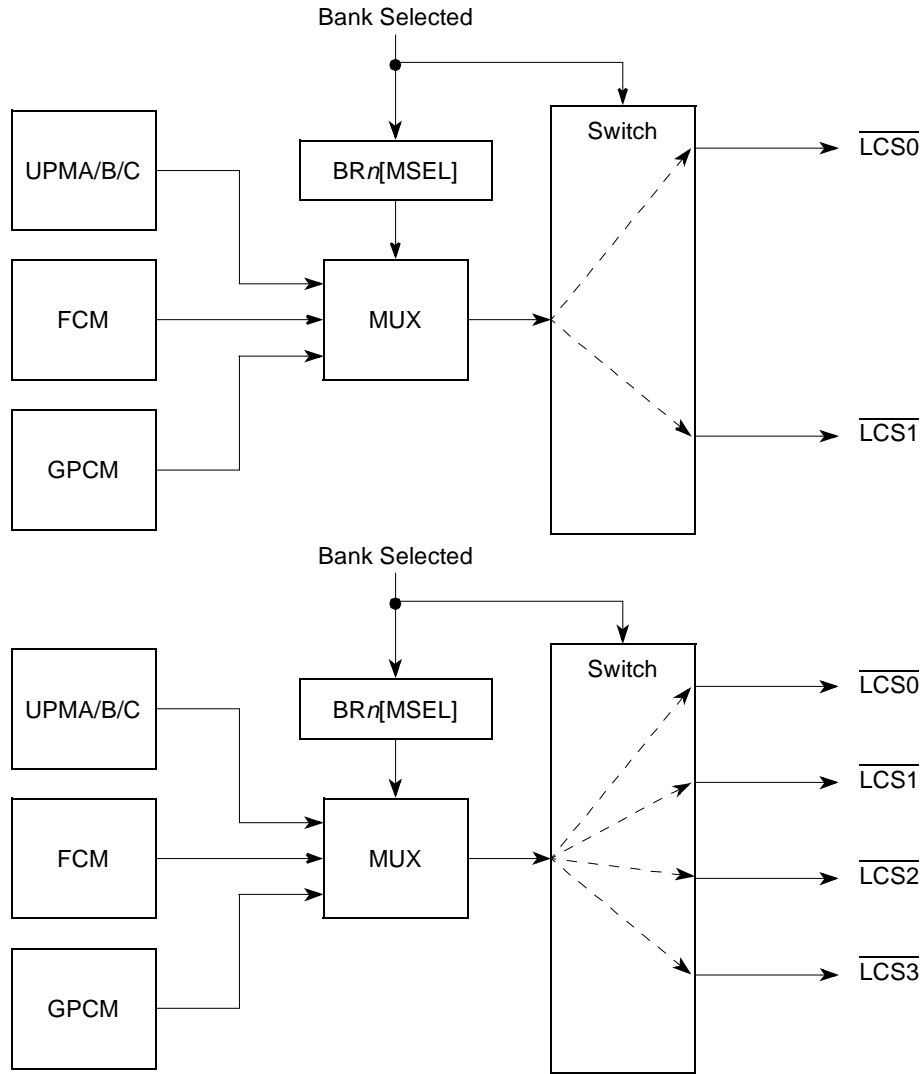


Figure 10-63. \overline{LCSn} Signal Selection

10.4.4.4.3 Byte Select Signal Timing (BST_n)

If $BR_n[MSEL]$ of the accessed memory bank selects a UPM on the currently requested cycle, the selected UPM affects the assertion and negation of the appropriate $\overline{LBS}[0:1]$ signal. The timing of both byte-select signals is specified in the RAM word. However, $\overline{LBS}[0:1]$ are also controlled by the port size of the accessed bank, the number of bytes to transfer, and the address accessed. [Figure 10-64](#) shows how UPMs control $\overline{LBS}[0:1]$.

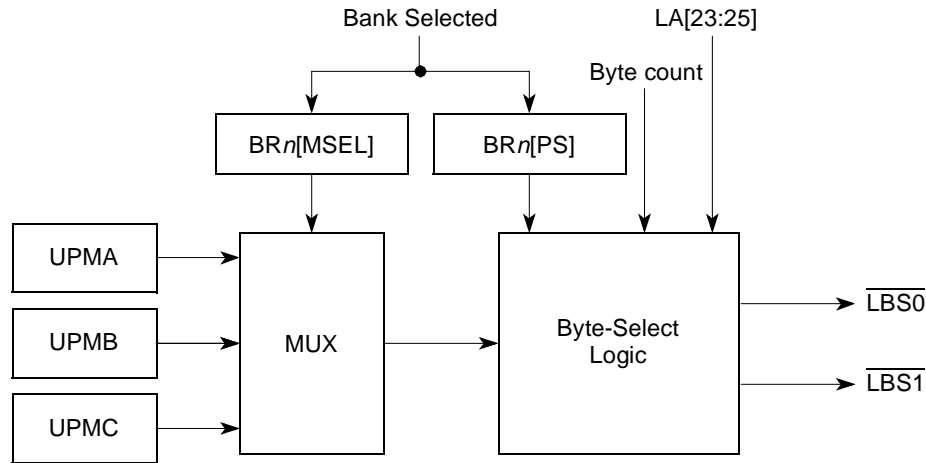


Figure 10-64. $\overline{\text{LBS}}$ Signal Selection

The uppermost byte select ($\overline{\text{LBS0}}$), when asserted, indicates that LAD[0:7] contains valid data during a cycle. Likewise, $\overline{\text{LBS1}}$ indicates that LAD[8:15] contain valid data. For a UPM refresh timer request, all $\overline{\text{LBS}}[0:1]$ signals are asserted/negated by the UPM according to the refresh pattern only. Following any internal bus monitor exception, the $\overline{\text{LBS}}[0:1]$ signals are negated regardless of the exception handling provided by any UPM exception pattern to prevent spurious writes to external RAM.

10.4.4.4.4 General-Purpose Signals ($GnTn$, GO_n)

The general-purpose signals (LGPL[0:5]) each have two bits in the RAM word that define the logical value of the signal to be changed at the rising edge of the bus clock and/or at the falling edge of the bus clock. LGPL0 offers enhancements beyond the other LGPL n lines.

LGPL0 can be controlled by an address line specified in MxMR[G0CL]. To use this feature, G0H and G0L should be set in the RAM word. For example, for a SIMM with multiple banks, this address line can be used to switch between internal memory device banks.

10.4.4.4.5 Loop Control (LOOP)

The LOOP bit in the RAM word specifies the beginning and end of a set of UPM RAM words that are to be repeated. The first time LOOP = 1, the memory controller recognizes it as a loop start word and loads the memory loop counter with the corresponding contents of the loop field shown in Table 10-39. The next RAM word for which LOOP = 1 is recognized as a loop end word. When it is reached, the loop counter is decremented by one.

Continued loop execution depends on the loop counter. If the counter is not zero, the next RAM word executed is the loop start word. Otherwise, the next RAM word executed is the one after the loop end word. Loops can be executed sequentially but cannot be nested. Also, special care must be taken:

- LAST and LOOP must not be set together.
- Loop start word should not have an AMX change with regard to the previous word.

Table 10-39. MxMR Loop Field Use

| Request Serviced | Loop Field |
|-------------------------|------------|
| Read single-beat cycle | RLF |
| Read burst cycle | RLF |
| Write single-beat cycle | WLF |
| Write burst cycle | WLF |
| Refresh timer expired | TLF |
| RUN command | RLF |

10.4.4.4.6 Repeat Execution of Current RAM Word (REDO)

The REDO function is useful for wait-state insertion in a long UPM routine that would otherwise need too many RAM words. Setting the REDO bits of the RAM word to a nonzero value causes the UPM to re-execute the current RAM word up to three more times, as defined in the REDO field of the current RAM word.

Special care must be taken in the following cases:

- When UTA and REDO are set together, TA is asserted the number of times specified by the REDO function.
- When NA and REDO are set together, the address is incremented the number of times specified by the REDO function.
- When LOOP and REDO are set together, the loop mechanism works as usual and the line is repeated according to the REDO function.
- LAST and REDO must not be set together.
- REDO should not be used within the exception routine.

10.4.4.4.7 Address Multiplexing (AMX)

Address lines can be controlled by the user-provided pattern in the UPM. The address multiplex (AMX) bits in the RAM word can choose between driving the transaction address (AMX = 00), driving it according to the multiplexing specified by the MxMR[AM] field (AMX = 10), or driving the contents of MAR (AMX = 11) on the address signals. The next address (NA) bit of the RAM word does not affect LA signals, unless AMX = 00 and chooses the column address for NA = 1.

In all cases, LA[21:25] of the eLBC are driven by the five lsbs of the address selected by AMX, regardless of whether the next address (NA) bit of the RAM word is used to increment the current address. The effect of NA = 1 is visible only when AMX = 00 chooses the column address.

10.4.4.4.8 Data Valid and Data Sample Control (UTA)

When a read access is handled by the UPM, and the UTA bit is 1 (data is to be sampled by the eLBC), the value of the DLT3 bit in the same RAM word, in conjunction with MxMR[GPL4], determines when the data input is sampled by the eLBC as follows:

- If MxMR[GPL4] = 1 (G4T4/DLT3 functions as DLT3) and DLT3 = 1 in the RAM word, data is latched on the falling edge of the bus clock instead of the rising edge. The eLBC samples the data on the next falling edge of the bus clock, which is during the middle of the current bus cycle. This feature should be used only in systems without external synchronous bus devices that require mid-cycle sampling.
- If MxMR[GPL4] = 0 (G4T4/DLT3 functions as G4T4), or if MxMR[GPL4] = 1 but DLT3 = 0 in the RAM word, data is latched on the rising edge of the bus clock, which occurs at the end of the current bus clock cycle (normal operation).

Figure 10-65 shows how data sampling is controlled by the UPM.




Figure 10-65. UPM Read Access Data Sampling

10.4.4.4.9 LGPL[0:5] Signal Negation (LAST)

When the LAST bit is read in a RAM word, the current UPM pattern is terminated at the end of the current cycle. On the next cycle (following LAST) all the UPM signals are negated unconditionally (driven to logic 1), unless there is a back-to-back UPM request pending. In this case, the signal values for the cycle following the one in which the LAST bit was set are taken from the first RAM word of the pending UPM routine.

In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.

10.4.4.4.10 Wait Mechanism (WAEN)

The WAEN bit in the RAM array word can be used to enable the UPM wait mechanism in selected UPM RAM words. If the UPM reads a RAM word with WAEN set, the external LUPWAIT signal is sampled and synchronized by the memory controller as if it were an asynchronous signal. The WAEN bit is ignored if LAST = 1 in the same RAM word.

Synchronization of LUPWAIT starts at the rising edge of the bus clock and takes at least 1 bus cycle to complete. If LUPWAIT is asserted and WAEN = 1 in the current UPM word, the UPM is frozen until LUPWAIT is negated. The value of external signals driven by the UPM remains as indicated in the previous RAM word. When LUPWAIT is negated, the UPM continues normal functions. Note that during WAIT cycles, the UPM does not handle data.

Figure 10-66 shows how the WAEN bit in the word read by the UPM and the LUPWAIT signal are used to hold the UPM in a particular state until LUPWAIT is negated. As the example shows, the LCS_n and LGPL1 states and the WAEN value are frozen until LUPWAIT is recognized as negated. WAEN is typically set before the line that contains $UTA = 1$. Note that if WAEN and NA are both set in the same RAM word, NA causes the burst address to increment once as normal regardless of whether the UPM freezes.

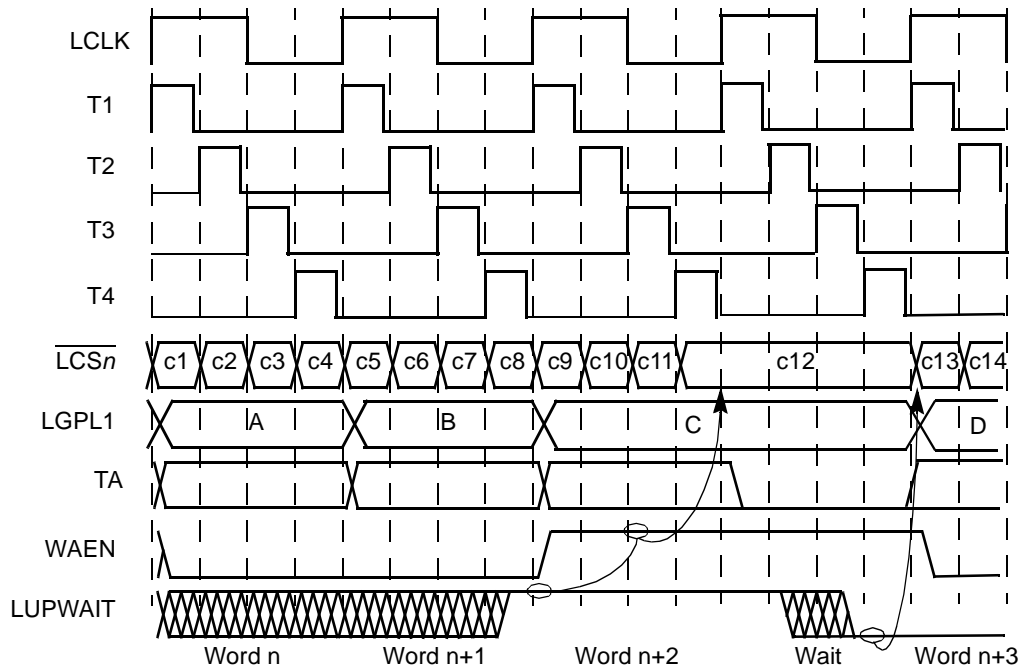


Figure 10-66. Effect of LUPWAIT Signal

If LUPWAIT is to be considered an asynchronous signal, which can be asserted/negated at any time, no UPM RAM word must contain both WAEN = 1 and UTA = 1 simultaneously.

10.4.4.5 Extended Hold Time on Read Accesses

Slow memory devices that take a long time to turn off their data bus drivers on read accesses should choose some non-zero combination of $OR_n[TRLX]$ and $OR_n[EHTR]$. The next accesses after a read access to the slow memory device is delayed by the number of clock cycles specified in the OR_n register in addition to any existing bus turnaround cycle.

10.5 Initialization/Application Information

This section provides information about the following:

- Section 10.5.1, “Interfacing to Peripherals in Different Address Modes”
- Section 10.5.2, “Bus Turnaround”
- Section 10.5.3, “Interface to Different Port-Size Devices”
- Section 10.5.4, “Command Sequence Examples for NAND Flash EEPROM”
- Section 10.5.5, “Interfacing to Fast-Page Mode DRAM Using UPM”
- Section 10.5.6, “Interfacing to ZBT SRAM Using UPM”

10.5.1 Interfacing to Peripherals in Different Address Modes

This section provides guidelines for interfacing to peripherals.

10.5.1.1 Multiplexed Address/Data Bus for 26-Bit Addressing

In this mode, the eLBC is used with port sizes of 8 and 16 bits; address bits A[6:21] will appear on LAD[0:15] (with zero bits on LAD[16:31]) during address phases, while the lower 10 bits of the address, A[22:31], are driven permanently on LA[16:25]. The connection is illustrated in Figure 10-67.

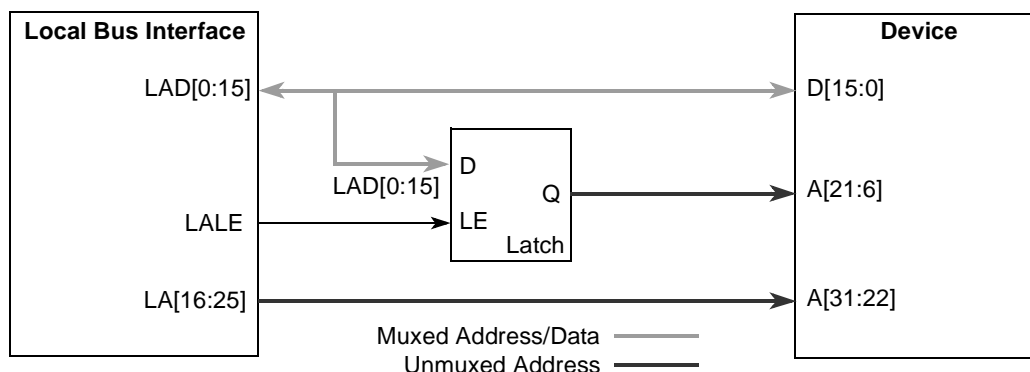


Figure 10-67. Multiplexed Address/Data Bus for 26-Bit Addressing

10.5.1.2 Peripheral Hierarchy on the Local Bus for High Bus Speeds

To achieve the highest possible bus speeds on the local bus, it is recommended to reduce the number of devices connected directly to the bus. For best results, only one bank of synchronous SRAMs should have a direct connection, and a bus demultiplexer should be used to replace separate latch and separate bus transceiver combinations. Figure 10-68 shows an example of such a hierarchy. This section is only a

guideline, and the board designer must simulate the electric characteristics of the scenario to determine the maximum operating frequency.

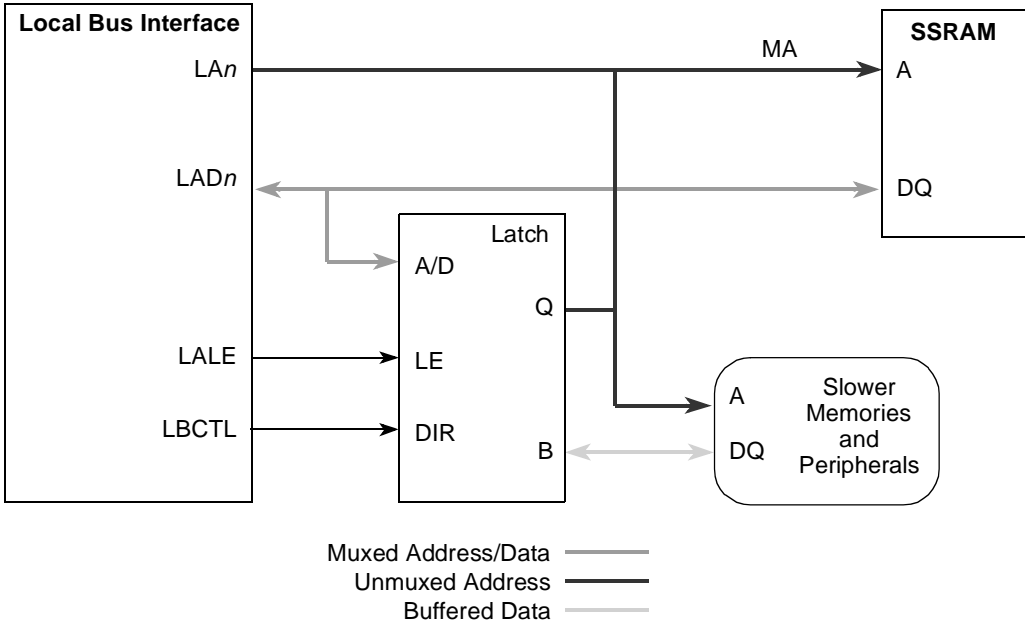


Figure 10-68. Local Bus Peripheral Hierarchy for High Bus Speeds

10.5.1.3 GPCM Timings

In case a system contains a memory hierarchy with high speed synchronous memories (synchronous SRAM) and lower speed asynchronous memories (for example, FLASH EPROM and peripherals) the GPCM-controlled memories should be decoupled by buffers to reduce capacitive loading on the bus. Those buffers have to be taken into account for the timing calculations. Figure 10-69 illustrates GPCM address timings.




Figure 10-69. GPCM Address Timings

To calculate address setup timing for a slower peripheral/memory device, several parameters have to be added: propagation delay for the address latch, propagation delay for the buffer and the address setup for the actual peripheral. Typical values for the two propagation delays are in the order of 3 to 6 ns, so for a 100-MHz bus frequency, \overline{LCS} should arrive on the order of 2 to 3 bus clocks later.

For data timings, only the propagation delay of one buffer plus the actual data setup time has to be considered. See [Figure 10-70](#) for an illustration of GPCM data timings.




Figure 10-70. GPCM Data Timings

10.5.2 Bus Turnaround

Because the local bus uses multiplexed address and data, special consideration must be given to avoid bus contention at bus turnaround. The following cases must be examined:

- Address phase after previous read
- Read data phase after address phase
- Read-modify-write cycle for parity protected memory banks
- UPM cycles with additional address phases

The bus does not change direction for the following cases so they need no special attention:

- Continued burst after the first beat
- Write data phase after address phase
- Address phase after previous write

10.5.2.1 Address Phase after Previous Read

During a read cycle, the memory/peripheral drives the bus and the bus transceiver drives LAD. After the data has been sampled, the output drivers of the external device must be disabled. This can take some time; for slow devices the EHTR feature of the GPCM or the programmability of the UPM should be used to guarantee that those devices have stopped driving the bus when the eLBC memory controller ends the bus cycle.

In this case, after the previous cycle ends, LBCTL goes high and changes the direction of the bus transceiver. The eLBC then inserts a bus turnaround cycle to avoid contention. The external device has now already placed its data signals in high impedance and no bus contention will occur.

10.5.2.2 Read Data Phase after Address Phase

During the address phase, LAD actively drives the address and LBCTL is high, driving the bus transceivers in the same direction as during a write. After the end of the address phase, LBCTL goes low and changes the direction of the bus transceiver. The eLBC places the LAD signals in high impedance after its $t_{dis}(LB)$. The LBCTL will have its new state after $t_{en}(LB)$ and, because this is an asynchronous input,

the transceiver starts to drive those signals after its t_{en} (transceiver) time. The system designer has to ensure, that $[t_{en}(LB) + t_{en}(\text{transceiver})]$ is larger than $t_{dis}(LB)$ to avoid bus contention.

10.5.2.3 Read-Modify-Write Cycle for Parity Protected Memory Banks

Principally, a read-modify-write cycle is a read cycle immediately followed by a write cycle. Because the write cycle will have a new address phase in any case, this basically is the same case as an address phase after a previous read.

10.5.2.4 UPM Cycles with Additional Address Phases

The flexibility of the UPM allows the user to insert additional address phases during read cycles by changing the AMX field, therefore turning around the bus during one pattern. The eLBC automatically inserts a single bus turnaround cycle if the bus (LAD) was previously high impedance for any reason, such as a read, before LALE is driven and LAD is driven with the new address. The turnaround cycle is not inserted on a write, because the bus was already driven to begin with.

However, bus contention could potentially still occur on the far side of a bus transceiver. It is the responsibility of the designer of the UPM pattern to guarantee that enough idle cycles are inserted in the UPM pattern to avoid this.

10.5.3 Interface to Different Port-Size Devices

The eLBC supports 8- and 16-bit data port sizes. However, the bus requires that the portion of the data bus used for a transfer to or from a particular port size be fixed. A 16-bit port must reside on LAD[0–15], and an 8-bit port must reside on LAD[0–7]. The local bus always tries to transfer the maximum amount of data on all bus cycles. [Figure 10-71](#) shows the device connections on the data bus.




Figure 10-71. Interface to Different Port-Size Devices

Table 10-41 lists the bytes required on the data bus for read cycles.

Table 10-41. Data Bus Drive Requirements For Read Cycles

| Transfer Size | Address State ¹ 3 lsbs | Port Size/LAD Data Bus Assignments | | | | | | | |
|---------------|--------------------------------------|------------------------------------|------|-------|-------|-------|------|-------|-------|
| | | 16-Bit | | | | 8-Bit | | | |
| | | 0-7 | 8-15 | 16-23 | 24-31 | 0-7 | 8-15 | 16-23 | 24-31 |
| Byte | 000 | OP0 | — | | | OP0 | | | |
| | 001 | — | OP1 | | | OP1 | | | |
| | 010 | OP2 | — | | | OP2 | | | |
| | 011 | — | OP3 | | | OP3 | | | |
| | 100 | OP4 | — | | | OP4 | | | |
| | 101 | — | OP5 | | | OP5 | | | |
| | 110 | OP6 | — | | | OP6 | | | |
| | 111 | — | OP7 | | | OP7 | | | |
| Half Word | 000 | OP0 | OP1 | | | OP0 | | | |
| | 001 | — | OP1 | | | OP1 | | | |
| | 010 | OP2 | OP3 | | | OP2 | | | |
| | 100 | OP4 | OP5 | | | OP4 | | | |
| | 101 | — | OP5 | | | OP5 | | | |
| | 110 | OP6 | OP7 | | | OP6 | | | |
| Word | 000 | OP0 | OP1 | | | OP0 | | | |
| | 100 | OP4 | OP5 | | | OP4 | | | |

¹ Address state is the calculated address for port size.

10.5.4 Command Sequence Examples for NAND Flash EEPROM

In order to program the eLBC and FCM for executing NAND Flash command sequences, command codes and pause states should be obtained from the relevant NAND Flash device data sheet and programmed into FCM configuration registers. This section illustrates some common sequences for large-page, multi-gigabit NAND Flash EEPROMs; however, details should be verified against manufacturers' specific programming data.

Throughout these examples it is assumed that one or more banks of eLBC has been configured under FCM control (BR_n[MSEL] = 001), with base address, port size, ECC mode, and timing parameters configured in accordance with the device's hardware specifications.

10.5.4.1 NAND Flash Soft Reset Command Sequence Example

An example of configuring FCM to execute a soft reset command to large-page NAND Flash is shown in [Table 10-42](#). This sequence does not require use of the shared FCM buffer RAM. The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled.

Table 10-42. FCM Register Settings for Soft Reset (ORn[PGS] = 1)

| Register | Initial Contents | Description |
|----------|------------------|--|
| FCR | 0xFF00_0000 | CMD0 = 0xFF = reset command; other commands unused |
| FBAR | — | Unused |
| FPAR | — | Unused |
| FBCR | — | Unused |
| MDR | — | Unused |
| FIR | 0x4000_0000 | OP0 = CM0 = command 0; OP1–OP7 = NOP |

10.5.4.2 NAND Flash Read Status Command Sequence Example

An example of configuring FCM to execute a status read command to large-page NAND Flash is shown in [Table 10-43](#). This sequence does not require use of the shared FCM buffer RAM, but reads the NAND Flash status into register MDR[AS0]. The sequence is initiated by writing FMR[OP] = 10 and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled.

Table 10-43. FCM Register Settings for Status Read (ORn[PGS] = 1)

| Register | Initial Contents | Description |
|----------|------------------|---|
| FCR | 0x7000_0000 | CMD0 = 0x70 = read status command; other commands unused |
| FBAR | — | Unused |
| FPAR | — | Unused |
| FBCR | — | Unused |
| MDR | — | Status returned in AS0 |
| FIR | 0x4B00_0000 | OP0 = CM0 = command 0; OP1 = RS = read status to MDR; OP2–OP7 = NOP |

10.5.4.3 NAND Flash Read Identification Command Sequence Example

An example of configuring FCM to execute a status ID command to large-page NAND Flash is shown in [Table 10-44](#). This sequence does not require use of the shared FCM buffer RAM, but uses MDR to set up a dummy address prior to the sequence, and then to receive the first 4 bytes of ID during the sequence. The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. At the

conclusion of the sequence, eLBC will issue a command complete interrupt (LTERSR[CC]) if interrupts are enabled. MDR[AS3–AS0] then can be read to obtain the first 4 bytes of NAND Flash ID.

Table 10-44. FCM Register Settings for ID Read (ORn[PGS] = 1)

| Register | Initial Contents | Description |
|----------|------------------|---|
| FCR | 0x9000_0000 | CMD0 = 0x90 = read ID command; other commands unused |
| FBAR | — | Unused |
| FPAR | — | Unused |
| FBCR | — | Unused |
| MDR | 0x0000_0000 | AS0 = 0x00 = dummy address for read ID command; AS0–AS3 return with first 4 bytes of ID code |
| FIR | 0x43BB_BBB0 | OP0 = CM0 = command 0; OP1 = UA = user address from MDR; OP2–OP6 = RS = read 4 bytes ID into MDR[AS3–AS0]; OP7 = NOP |

10.5.4.4 NAND Flash Page Read Command Sequence Example

An example of configuring FCM to execute a random page read command to large-page NAND Flash is shown in [Table 10-45](#). This sequence reads an entire page (main and spare region) into the shared FCM buffer RAM, checking ECC as it proceeds. The sequence is initiated by writing FMR[OP] = 11, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTERSR[CC]) if interrupts are enabled.

Table 10-45. FCM Register Settings for Page Read (ORn[PGS] = 1)

| Register | Initial Contents | Description |
|----------|--|---|
| FCR | 0x0030_0000 | CMD0 = 0x00 = random read address entry; CMD1 = 0x30 = read page |
| FBAR | block index (for example block 0x0001_0AB4) | BLK locates index of 128-Kbyte block |
| FPAR | page offset (for example 0x0000_5000 locates page 5, buffer 1) | PI locates page index in BLK; PI mod 2 indexes FCM buffer RAM; MS = 0 and CI = 0 |
| FBCR | 0x0000_0000 | BC = 0 to read entire 2112-byte page with ECC check |
| MDR | — | Unused |
| FIR | 0x4125_E000 | OP0 = CM0 = command 0; OP1 = CA = column address; OP2 = PA = page address; OP3 = CM1 = command 1; OP4 = RBW = wait on Flash ready and read data into FCM buffer; OP5–OP7 = NOP |

10.5.4.5 NAND Flash Block Erase Command Sequence Example

An example of configuring FCM to execute a block erase command to large-page NAND Flash is shown in [Table 10-46](#). This sequence does not require use of the shared FCM buffer RAM, but returns with the erase status in MDR[AS0]. The sequence is initiated by writing FMR[OP] = 11, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTER[CC]) if interrupts are enabled.

Note that operations specified by OP3 and OP4 (status read) should never be skipped while erasing a NAND Flash device, because, in case that happens, contention may arise on LGPL4. A possible case is that the next transaction from eLBC may try to use that pin as an output and since the NAND Flash device might already be driving it, contention will occur. In case OP3 and OP4 operations are skipped, it may also happen that a new command is issued to the NAND Flash device even when the device has not yet finished processing the previous request. This may also result in unpredictable behavior.

Table 10-46. FCM Register Settings for Block Erase (ORn[PGS] = 1)

| Register | Initial Contents | Description |
|----------|---|--|
| FCR | 0x6070_D000 | CMD0 = 0x60 = block address entry; CMD1 = 0x70 = read status CMD2 = 0xD0 = erase block; |
| FBAR | Block Index (for example block 0x0001_0AB4) | BLK locates index of 128-Kbyte block |
| FPAR | 0x0000_0000 | PI = 0 to locate block boundary |
| FBCR | — | Unused |
| MDR | — | Returns with AS0 holding erase status |
| FIR | 0x426D_B000 | OP0 = CM0 = command 0; OP1 = PA = page address; OP2 = CM2 = command 2; OP3 = CW1 = wait on Flash ready and issue command 1; OP4 = RS = read erase status into MDR[AS0]; OP5–OP7 = NOP |

10.5.4.6 NAND Flash Program Command Sequence Example

An example of configuring FCM to execute a program command to large-page NAND Flash is shown in [Table 10-47](#). This sequence writes an entire page (main and spare region) from the shared FCM buffer RAM, generating ECC as it proceeds. The shared buffer (buffer 1 for page index 5) must be initialized by software prior to starting the sequence. The sequence is initiated by writing FMR[OP] = 11, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTER[CC]) if interrupts are enabled. The status of the programming operation is returned in MDR[AS0].

Note that operations specified by OP5 and OP6 (status read) should never be skipped while programming a NAND Flash device, because, in case that happens, contention may arise on LGPL4. A possible case is that the next transaction from eLBC may try to use that pin as an output and since the NAND Flash device might already be driving it, contention will occur. In case OP5 and OP6 operations are skipped, it may also


happen that a new command is issued to the NAND Flash device even when the device has not yet finished processing the previous request. This may also result in unpredictable behavior.

Table 10-47. FCM Register Settings for Page Program (OR_n[PGS] = 1)

| Register | Initial Contents | Description |
|----------|---|--|
| FCR | 0x8070_1000 | CMD0 = 0x80 = page address and data entry; CMD1 = 0x70 = read status CMD2 = 0x10 = program page; |
| FBAR | block index (for example, block 0x0001_0AB4) | BLK locates index of 128-Kbyte block |
| FPAR | page offset (for example, 0x0000_5000 locates page 5, buffer 1) | PI locates page index in BLK; PI mod 2 indexes FCM buffer RAM; MS = 0 and CI = 0 |
| FBCR | 0x0000_0000 | BC = 0 to write entire 2112-Byte page with ECC generation |
| MDR | — | Returns with AS0 holding program status |
| FIR | 0x4128_6DB0 | OP0 = CM0 = command 0; OP1 = CA = column address; OP2 = PA = page address; OP3 = WB = write data from buffer; OP4 = CM2 = command 2; OP5 = CW1 = wait on Flash ready and issue command 1; OP6 = RS = read erase status into MDR[AS0]; OP7 = NOP |

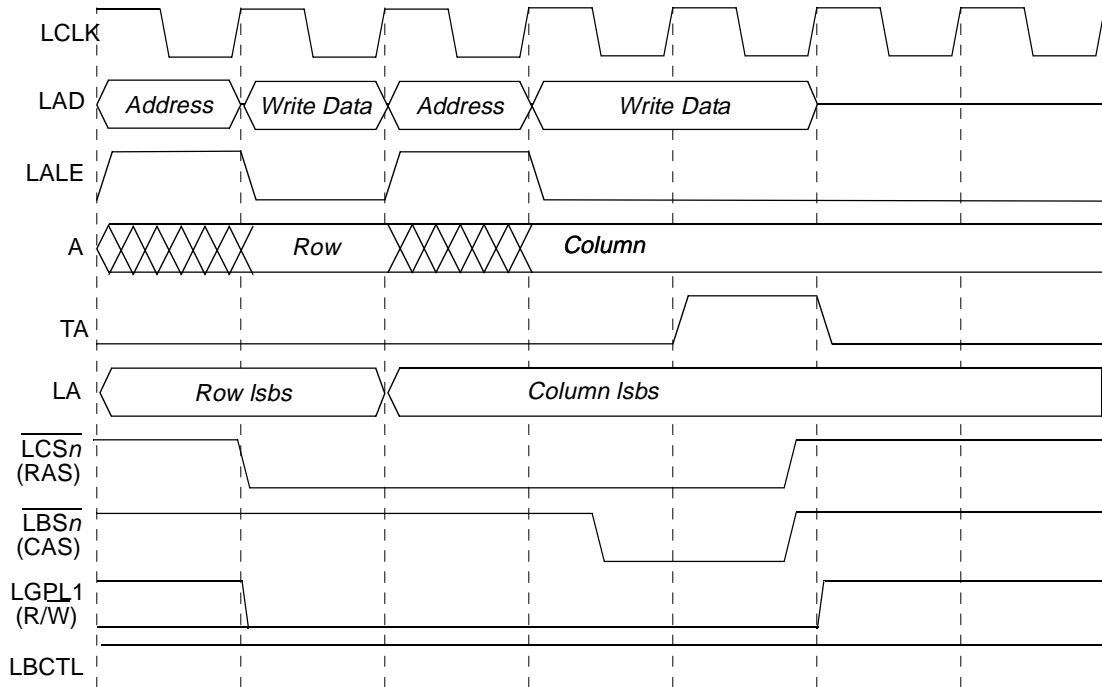
10.5.5 Interfacing to Fast-Page Mode DRAM Using UPM

Connecting the local bus UPM controller to a DRAM device requires a detailed examination of the timing diagrams representing the possible memory cycles that must be performed when accessing this device. This section describes timing diagrams for various UPM configurations for fast-page mode DRAM, with LCRR[CLKDIV] = 4 or 8. The illustrative examples shown in [Figure 10-72–Figure 10-77](#) may not represent the timing necessary for any specific device used with the eLBC. Here, LGPL1 is programmed to drive R/W of the DRAM, although any LGPL_n signal may be used for this purpose.



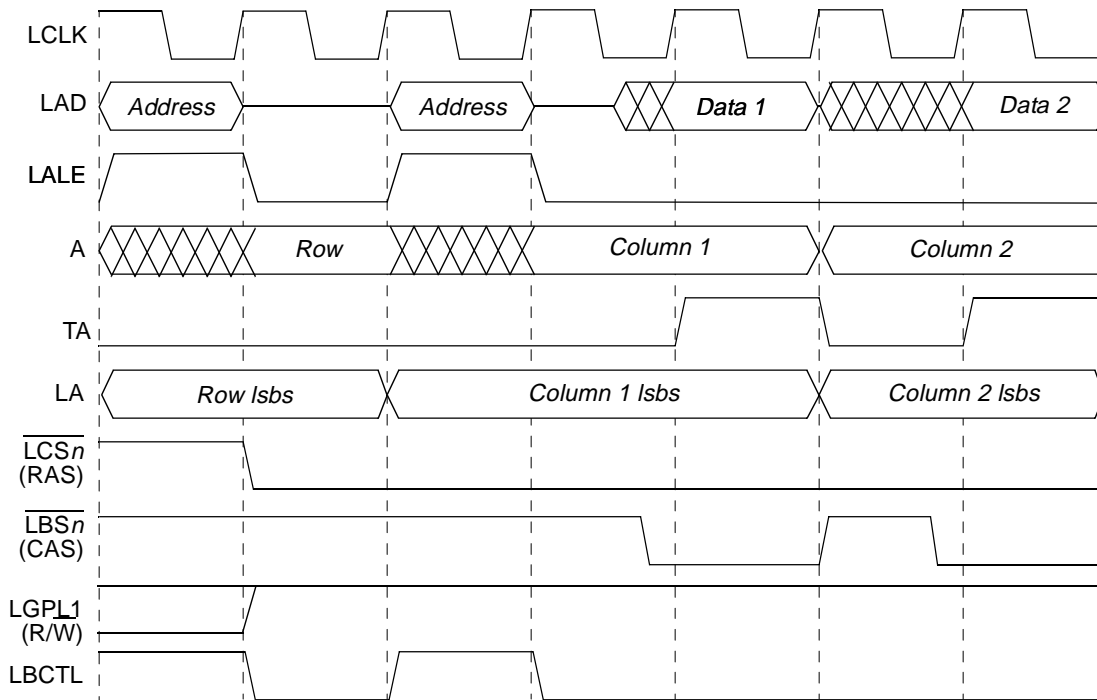
| | | | | | |
|---------|-----|--|---------|---------|--------|
| cst1 | 0 | LALE pause (due to change in AMX) | 0 | 0 | Bit 0 |
| cst2 | 0 | | 0 | 0 | Bit 1 |
| cst3 | 0 | | 0 | 0 | Bit 2 |
| cst4 | 0 | | 0 | 0 | Bit 3 |
| bst1 | 1 | | 1 | 0 | Bit 4 |
| bst2 | 1 | | 0 | 0 | Bit 5 |
| bst3 | 1 | | 0 | 0 | Bit 6 |
| bst4 | 1 | | 0 | 0 | Bit 7 |
| g0i0 | | | | | Bit 8 |
| g0i1 | | | | | Bit 9 |
| g0h0 | | | | | Bit 10 |
| g0h1 | | | | | Bit 11 |
| g1t1 | 1 | | 1 | 1 | Bit 12 |
| g1t3 | 1 | | 1 | 1 | Bit 13 |
| g2t1 | | | | | Bit 14 |
| g2t3 | | | | | Bit 15 |
| g3t1 | | | | | Bit 16 |
| g3t3 | | | | | Bit 17 |
| g4t1 | | | | | Bit 18 |
| g4t3 | | | | | Bit 19 |
| g5t1 | | | | | Bit 20 |
| g5t3 | | | | | Bit 21 |
| redo[0] | | | | | Bit 22 |
| redo[1] | | | | | Bit 23 |
| loop | 0 | | 0 | 0 | Bit 24 |
| exen | 0 | | 0 | 0 | Bit 25 |
| amx0 | 1 | | 0 | 0 | Bit 26 |
| amx1 | 0 | | 0 | 0 | Bit 27 |
| na | 0 | | 0 | 0 | Bit 28 |
| uta | 0 | | 0 | 1 | Bit 29 |
| todt | 0 | | 0 | 1 | Bit 30 |
| last | 0 | | 0 | 1 | Bit 31 |
| | RSS | RSS + 1 | RSS + 1 | RSS + 2 | |

Figure 10-72. Single-Beat Read Access to FPM DRAM




| | | | | | |
|---------|---|--|---------|---------|---------|
| cst1 | 0 | LALE pause (due to change in AMX) | 0 | 0 | Bit 0 |
| cst2 | 0 | | 0 | 0 | Bit 1 |
| cst3 | 0 | | 0 | 0 | Bit 2 |
| cst4 | 0 | | 0 | 1 | Bit 3 |
| bst1 | 1 | | 1 | 0 | Bit 4 |
| bst2 | 1 | | 1 | 0 | Bit 5 |
| bst3 | 1 | | 0 | 0 | Bit 6 |
| bst4 | 1 | | 0 | 1 | Bit 7 |
| g0i0 | | | | | Bit 8 |
| g0i1 | | | | | Bit 9 |
| g0h0 | | | | | Bit 10 |
| g0h1 | | | | | Bit 11 |
| g1t1 | 0 | | 0 | 0 | Bit 12 |
| g1t3 | 0 | | 0 | 0 | Bit 13 |
| g2t1 | | | | | Bit 14 |
| g2t3 | | | | | Bit 15 |
| g3t1 | | | | | Bit 16 |
| g3t3 | | | | | Bit 17 |
| g4t1 | | | | | Bit 18 |
| g4t3 | | | | | Bit 19 |
| g5t1 | | | | | Bit 20 |
| g5t3 | | | | | Bit 21 |
| redo[0] | | | | | Bit 22 |
| redo[1] | | | | | Bit 23 |
| loop | 0 | | 0 | 0 | Bit 24 |
| exen | 0 | | 0 | 0 | Bit 25 |
| amx0 | 1 | | 0 | 0 | Bit 26 |
| amx1 | 0 | | 0 | 0 | Bit 27 |
| na | 0 | | 0 | 0 | Bit 28 |
| uta | 0 | | 0 | 1 | Bit 29 |
| todt | 0 | | 0 | 1 | Bit 30 |
| last | 0 | | 0 | 1 | Bit 31 |
| | | WSS | WSS + 1 | WSS + 1 | WSS + 2 |

Figure 10-73. Single-Beat Write Access to FPM DRAM




| | | | | | | |
|---------|-----|--|---------|---------|---------|--------|
| cst1 | 0 | LALE pause (due to change in AMX) | 0 | 0 | 1 | Bit 0 |
| cst2 | 0 | | 0 | 0 | 1 | Bit 1 |
| cst3 | 0 | | 0 | 0 | 1 | Bit 2 |
| cst4 | 0 | | 0 | 0 | 1 | Bit 3 |
| bst1 | 1 | | 1 | 0 | 1 | Bit 4 |
| bst2 | 1 | | 1 | 0 | 1 | Bit 5 |
| bst3 | 1 | | 1 | 0 | 1 | Bit 6 |
| bst4 | 1 | | 0 | 0 | 1 | Bit 7 |
| g0i0 | | | | | | Bit 8 |
| g0i1 | | | | | | Bit 9 |
| g0h0 | | | | | | Bit 10 |
| g0h1 | | | | | | Bit 11 |
| g1t1 | 1 | | 1 | 1 | 1 | Bit 12 |
| g1t3 | 1 | | 1 | 1 | 1 | Bit 13 |
| g2t1 | | | | | | Bit 14 |
| g2t3 | | | | | | Bit 15 |
| g3t1 | | | | | Bit 16 | |
| g3t3 | | | | | Bit 17 | |
| g4t1 | | | | | Bit 18 | |
| g4t3 | | | | | Bit 19 | |
| g5t1 | | | | | Bit 20 | |
| g5t3 | | | | | Bit 21 | |
| redo[0] | | | | | Bit 22 | |
| redo[1] | | | | | Bit 23 | |
| loop | 0 | 1 | 1 | 0 | Bit 24 | |
| exen | 0 | 0 | 1 | 0 | Bit 25 | |
| amx0 | 1 | 0 | 0 | 0 | Bit 26 | |
| amx1 | 0 | 0 | 0 | 0 | Bit 27 | |
| na | 0 | 0 | 1 | 0 | Bit 28 | |
| uta | 0 | 0 | 1 | 0 | Bit 29 | |
| todt | 0 | 0 | 0 | 1 | Bit 30 | |
| last | 0 | 0 | 0 | 1 | Bit 31 | |
| | RBS | | RBS + 1 | RBS + 2 | RBS + 3 | |

Figure 10-74. Burst Read Access to FPM DRAM Using LOOP (Two Beats Shown)



| | | | | |
|---------|-----|---------|---------|--------|
| cst1 | 1 | 0 | 0 | Bit 0 |
| cst2 | 1 | 0 | 0 | Bit 1 |
| cst3 | 1 | 0 | 1 | Bit 2 |
| cst4 | 1 | 0 | 1 | Bit 3 |
| bst1 | 1 | 0 | 0 | Bit 4 |
| bst2 | 0 | 0 | 0 | Bit 5 |
| bst3 | 0 | 0 | 1 | Bit 6 |
| bst4 | 0 | 0 | 1 | Bit 7 |
| g0l0 | | | | Bit 8 |
| g0l1 | | | | Bit 9 |
| g0h0 | | | | Bit 10 |
| g0h1 | | | | Bit 11 |
| g1t1 | | | | Bit 12 |
| g1t3 | | | | Bit 13 |
| g2t1 | | | | Bit 14 |
| g2t3 | | | | Bit 15 |
| g3t1 | | | | Bit 16 |
| g3t3 | | | | Bit 17 |
| g4t1 | | | | Bit 18 |
| g4t3 | | | | Bit 19 |
| g5t1 | | | | Bit 20 |
| g5t3 | | | | Bit 21 |
| redo[0] | | | | Bit 22 |
| redo[1] | | | | Bit 23 |
| loop | 0 | 0 | 0 | Bit 24 |
| exen | 0 | 0 | 0 | Bit 25 |
| amx0 | 0 | 0 | 0 | Bit 26 |
| amx1 | 0 | 0 | 0 | Bit 27 |
| na | 0 | 0 | 0 | Bit 28 |
| uta | 0 | 0 | 0 | Bit 29 |
| todt | 0 | 0 | 1 | Bit 30 |
| last | 0 | 0 | 1 | Bit 31 |
| | PTS | PTS + 1 | PTS + 2 | |

Figure 10-75. Refresh Cycle (CBR) to FPM DRAM



| | | |
|------------|---|--------|
| cst1 | 1 | Bit 0 |
| cst2 | 1 | Bit 1 |
| cst3 | 1 | Bit 2 |
| cst4 | 1 | Bit 3 |
| bst1 | 1 | Bit 4 |
| bst2 | 1 | Bit 5 |
| bst3 | 1 | Bit 6 |
| bst4 | 1 | Bit 7 |
| g0l0 | | Bit 8 |
| g0l1 | | Bit 9 |
| g0h0 | | Bit 10 |
| g0h1 | | Bit 11 |
| g1t1 | | Bit 12 |
| g1t3 | | Bit 13 |
| g2t1 | | Bit 14 |
| g2t3 | | Bit 15 |
| g3t1 | | Bit 16 |
| g3t3 | | Bit 17 |
| g4t1 | | Bit 18 |
| g4t3 | | Bit 19 |
| g5t1 | | Bit 20 |
| g5t3 | | Bit 21 |
| redo[0] | | Bit 22 |
| redo[1] | | Bit 23 |
| loop | 0 | Bit 24 |
| exen | 0 | Bit 25 |
| amx0 | 0 | Bit 26 |
| amx1 | 0 | Bit 27 |
| na | 0 | Bit 28 |
| uta | 0 | Bit 29 |
| todt | 1 | Bit 30 |
| last | 1 | Bit 31 |
| EXS | | |

Figure 10-76. Exception Cycle

10.5.6 Interfacing to ZBT SRAM Using UPM

ZBT SRAMs have been designed to optimize the performance of table access in networking applications. This section describes how to interface to ZBT SRAMs. [Figure 10-77](#) shows the connections. The UPM is used to generate control signals. The same interfacing is used for pipelined and flow-through versions of ZBT SRAMs. However different UPM patterns must be generated for those cases. As ZBT SRAMs are

mostly used by performance-critical applications, it is assumed here that, typically, the maximum width of the local bus of 16 bits will be used.

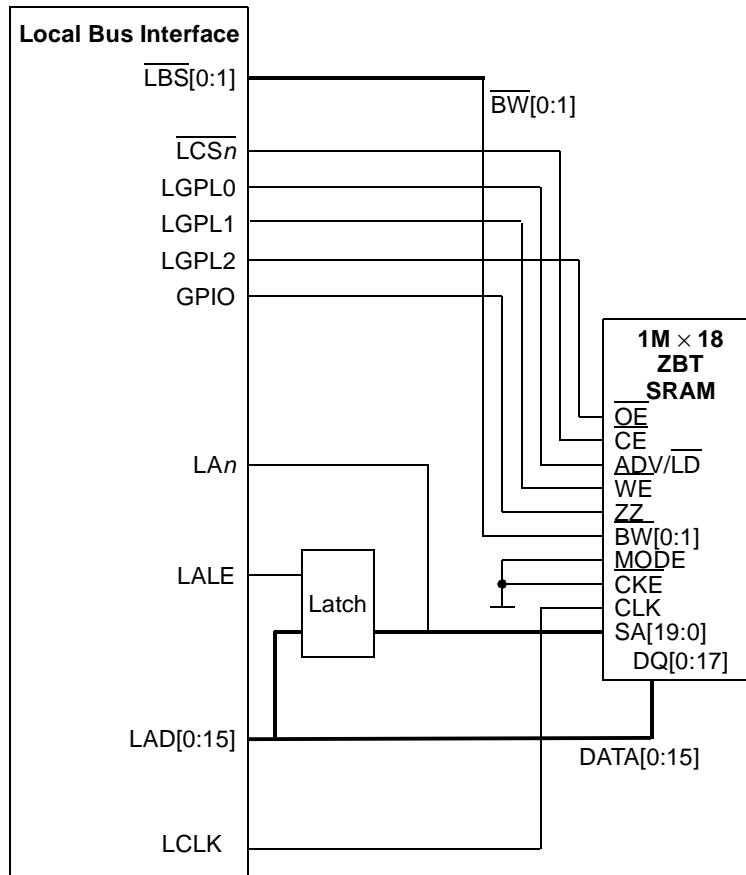


Figure 10-77. Interface to ZBT SRAM

ZBT SRAMs allow different configurations. For the local bus, the burst order should be set to linear burst order by tying the mode pin to GND. \overline{CKE} should also be tied to ground.

ZBT SRAMs perform four-beat bursts. Because the eLBC generates sixteen-beat transactions (for 16-bit ports) the UPM breaks down each burst into four consecutive four-beat bursts. The internal address generator of the eLBC generates the new $\{A21, A22\}$ for the second, third, and fourth burst. In other words, because linear burst is used on the SRAM, the device itself bursts with the burst addresses of $[0:1:2:3]$. The local bus always generates linear bursts and expects $[0:1:2:3:4:5:6:7:\dots:15]$. Therefore, four consecutive linear bursts of the ZBT SRAM with $\{A21, A22\} = \{0,0\}$ for the first burst, $\{A21, A22\} = \{0,1\}$ for the second burst, $\{A21, A22\} = \{1,0\}$ for the third burst, and $\{A21, A22\} = \{1,1\}$ for the fourth burst give the desired burst pattern.

The UPM also supports single beat accesses. Because the ZBT SRAM does not support this and always responds with a burst, the UPM pattern has to take care that data for the critical beat is provided (for write) or sampled (for read), and that the rest of the burst is ignored (by negating \overline{WE}). The UPM controller basically has to wait for the end of the SRAM burst to avoid bus contention with further bus activities.

Chapter 11 Sequencer

11.1 Sequencer Overview

The I/O sequencer switches transactions among its ports, using a buffer pool to minimize blocking. Figure 11-1 is a block diagram of the I/O sequencer (IOS).




Figure 11-1. I/O Sequencer Block Diagram

11.1.1 Sequencer Features

The I/O sequencer includes the following features:

- Switches transactions among its ports
- Contains 8 cache-line (32-byte) buffers to allow streaming of PCI transactions
- Performs address translation on outbound PCI transactions

Note that the number of CSB masters allowed to access to PCI outbound windows is restricted to no more than four non-CPU masters or no more than two non-CPU plus the CPU. If this number is exceeded, deadlock can occur on CSB arbitration.

11.2 Sequencer External Signal Description

The I/O sequencer has no external signals.

11.3 Sequencer Memory Map/Register Definition

Table 11-1 shows the I/O sequencer memory map.

Table 11-1. Sequencer Memory Map

| Offset | Register | Access | Reset | Section/Page |
|--|--|--------|-----------|-----------------------------|
| I/O Sequencer (IOS)—Block Base Address 0x0_8400 | | | | |
| 0x00 | POTAR0—PCI outbound translation address register 0 | R/W | All zeros | 11.4.1/11-3 |
| 0x08 | POBAR0—PCI outbound base address register 0 | R/W | All zeros | 11.4.2/11-3 |
| 0x10 | POCMR0—PCI outbound comparison mask register 0 | R/W | All zeros | 11.4.3/11-4 |
| 0x18 | POTAR1—PCI outbound translation address register 1 | R/W | All zeros | 11.4.1/11-3 |
| 0x20 | POBAR1—PCI outbound base address register 1 | R/W | All zeros | 11.4.2/11-3 |
| 0x28 | POCMR1—PCI outbound comparison mask register 1 | R/W | All zeros | 11.4.3/11-4 |
| 0x30 | POTAR2—PCI outbound translation address register 2 | R/W | All zeros | 11.4.1/11-3 |
| 0x38 | POBAR2—PCI outbound base address register 2 | R/W | All zeros | 11.4.2/11-3 |
| 0x40 | POCMR2—PCI outbound comparison mask register 2 | R/W | All zeros | 11.4.3/11-4 |
| 0x48 | POTAR3—PCI outbound translation address register 3 | R/W | All zeros | 11.4.1/11-3 |
| 0x50 | POBAR3—PCI outbound base address register 3 | R/W | All zeros | 11.4.2/11-3 |
| 0x58 | POCMR3—PCI outbound comparison mask register 3 | R/W | All zeros | 11.4.3/11-4 |
| 0x60 | POTAR4—PCI outbound translation address register 4 | R/W | All zeros | 11.4.1/11-3 |
| 0x68 | POBAR4—PCI outbound base address register 4 | R/W | All zeros | 11.4.2/11-3 |
| 0x70 | POCMR4—PCI outbound comparison mask register 4 | R/W | All zeros | 11.4.3/11-4 |
| 0x78 | POTAR5—PCI outbound translation address register 5 | R/W | All zeros | 11.4.1/11-3 |
| 0x80 | POBAR5—PCI outbound base address register 5 | R/W | All zeros | 11.4.2/11-3 |

Table 11-1. Sequencer Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--------|--|--------|-----------|--------------|
| 0x88 | POCMR5—PCI outbound comparison mask register 5 | R/W | All zeros | 11.4.3/11-4 |
| 0xF0 | PMCR—Power management control register | R/W | All zeros | 11.4.4/11-5 |
| 0xF8 | DTCR—Discard timer control register | R/W | All zeros | 11.4.5/11-6 |

11.4 Sequencer Register Descriptions

11.4.1 PCI Outbound Translation Address Registers (POTAR_n)

The PCI outbound translation address register defines the location of the outbound translation window in the PCI (translated) address space. [Figure 11-2](#) shows the POTAR_n register fields.


Figure 11-2. PCI Outbound Translation Address Registers (POTAR_n)

[Table 11-2](#) describes POTAR_n fields.

Table 11-2. POTAR_n Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–11 | — | Reserved |
| 12–31 | TA | Translation address. Contains the starting address of the outbound translated address. It also corresponds to the most-significant 20 bits of a 32-bit address. The translation address must be aligned based on the window's size. |

11.4.2 PCI Outbound Base Address Registers (POBAR_n)

The PCI outbound base address register (POBAR_n) defines the location of the outbound translation window in the local (source) memory space. [Figure 11-3](#) shows the POBAR_n register fields.



Figure 11-3. PCI Outbound Base Address Registers (POBAR_n)

Table 11-3 describes POBAR_n fields.

Table 11-3. POBAR_n Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–11 | — | Reserved |
| 12–31 | BA | Base address. This field contains the starting address of the outbound translated window. This field corresponds to the most-significant 20 bits of a 32-bit address. |

11.4.3 PCI Outbound Comparison Mask Registers (POCMR_n)

The PCI outbound comparison mask register (POCMR_n) defines the size and destination of the outbound translation window. It also defines some properties of the window in the PCI address space. See Section 11.5.1, “Transaction Forwarding,” for more information. Figure 11-4 shows the POCMR_n register fields.




Figure 11-4. PCI Outbound Comparison Mask Registers (POCMR_n)

Table 11-4 describes the bit settings of the POCMR_n register.

Table 11-4. POCMR_n Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0 | EN | Enable. Enables the address translation window. 0 Address translation is disabled for this window. 1 Address translation is enabled for this window. Local addresses that match the definition of the window will be recognized by the device and translated to the PCI memory space. |
| 1 | IO | I/O space. Determines whether the window is mapped to the PCI memory space or PCI I/O space. 0 Memory space 1 I/O space |
| 2–11 | — | Reserved, should be cleared. |

11.5.1.1 Transactions from the Coherency System Bus (CSB) Port

Transactions from the CSB port are forwarded as follows:

- If the address matches the 12-byte PCI controller software configuration memory space of the PCI controller, the transaction is forwarded to the PCI port. These address values are configuration options of the I/O sequencer. See [Table 13-3](#) for more information on PCI controller software configuration memory space.
- If the address matches the DMA register memory space, the transaction is forwarded to the DMA port.
- If the address hits any of the outbound translation windows, the transaction is forwarded to the PCI port, with the address translated. See [Section 11.5.2, “PCI Outbound Address Translation,”](#) for more information.

11.5.1.2 Transactions from the PCI Port

Transactions from the PCI port are forwarded as follows:

- If the address matches the DMA register memory space, the transaction is forwarded to the DMA port.
- All other transactions are forwarded to the CSB port.

11.5.1.3 Transactions from the DMA Port

Transactions from the DMA port are forwarded as follows:

- If the address hits any of the outbound translation windows, the transaction is forwarded to the PCI port, with the address translated. See [Section 11.5.2, “PCI Outbound Address Translation,”](#) for more information.
- All other transactions are forwarded to the CSB port.

11.5.2 PCI Outbound Address Translation

Outbound address translation is provided to allow the outbound transactions to access any address over the PCI memory or I/O space. Translation window base addresses are defined in the PCI outbound base address registers. See [Section 11.4.2, “PCI Outbound Base Address Registers \(POBARn\),”](#) for more information. Transactions to these address ranges are issued on the PCI bus with a translated address. The translation addresses are defined in the associated PCI outbound translation address registers (POTARs).

Figure 11-7 shows an example translation window for outbound memory accesses.

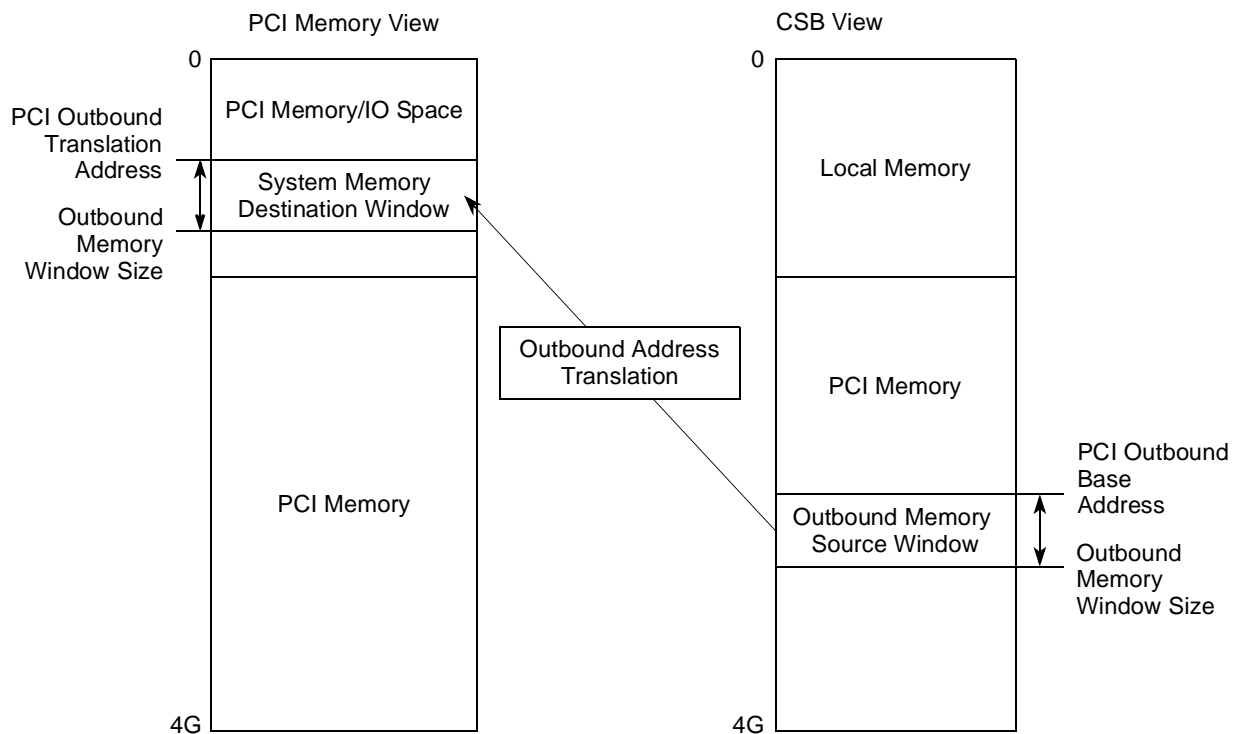


Figure 11-7. Outbound PCI Memory Address Translation

The six sets of outbound translation registers allow six simultaneous translation windows to the PCI port. Software can move and adjust the memory window translations and sizes during run-time. This allows software to access different PCI memory/IO spaces on-the-fly, but the PCI outbound translation source windows must not overlap. However, outbound translation destination windows can be overlapped.

11.5.3 Transaction Ordering

The following rules are applied to maintain proper ordering of transactions:

- The transactions arriving from each port are dispatched to the destination port in the order of arrival. The dispatch order of transactions arriving on different ports is not necessarily maintained.
- A read transaction that originates at the CSB port and reads from the PCI port pulls out of the IOS any posted writes that originated on the PCI port and were posted before the read data arrives from the PCI.
- The IOS can always accept a write from the PCI port without forcing the PCI port to first accept a read.

Chapter 12

DMA/Messaging Unit

The DMA/messaging unit supports communication between two processors on different buses, for example, a local processor and a processor on a PCI bus. This unit operates with generic messages and doorbell registers. [Figure 12-1](#) is a block diagram of the DMA/messaging unit.




Figure 12-1. DMA/Messaging Unit Block Diagram

This block also provides a DMA controller, which transfers blocks of data independent of the local processor or PCI hosts. The DMA module has four high-speed DMA channels, which share buffer space in the I/O sequencer (IOS) to facilitate the gathering and sending of data.

12.1 DMA Features

The DMA/messaging unit includes the following features:

- Message and doorbell registers for inter-processor communication
- DMA controller
 - Four DMA channels
 - Concurrent execution across multiple channels with programmable bandwidth control
 - Misaligned transfer capability
 - Data chaining and direct mode

- Interrupt on completed segment, chain, and error
- Optional external control signals (REQ/ACK/DONE) per channel

12.2 DMA External Signal Description

This section describes the DMA signals.

12.2.1 DMA Detailed Signal Descriptions

Table 12-1 contains the detailed descriptions of the DMA interface signals.

Table 12-1. DMA Interface Signals—Detailed Signal Descriptions

| Signal | I/O | Description | | |
|---|--|--|--|--|
| $\overline{\text{DREQ}}[0:1]$ | I | DMA request signals, one per channel. The DMA request signal indicates the start or continuation of a DMA transfer. The falling edge of $\overline{\text{DREQ}}_n$ causes $\text{DMAMR}_n[\text{CS}]$ to be set, thereby activating the DMA channel. | | |
| | | <table border="0"> <tr> <td style="vertical-align: top;">State Meaning</td> <td>Asserted—Assertion of $\overline{\text{DREQ}}_n$ starts or resumes a DMA transfer if $\text{DMAMR}_n[\text{EMSEN}]$ is 1. Negated—Negation of $\overline{\text{DREQ}}_n$ has no effect.</td> </tr> </table> | State Meaning | Asserted—Assertion of $\overline{\text{DREQ}}_n$ starts or resumes a DMA transfer if $\text{DMAMR}_n[\text{EMSEN}]$ is 1. Negated—Negation of $\overline{\text{DREQ}}_n$ has no effect. |
| | | State Meaning | Asserted—Assertion of $\overline{\text{DREQ}}_n$ starts or resumes a DMA transfer if $\text{DMAMR}_n[\text{EMSEN}]$ is 1. Negated—Negation of $\overline{\text{DREQ}}_n$ has no effect. | |
| <table border="0"> <tr> <td style="vertical-align: top;">Timing</td> <td>Assertion—Can be asserted asynchronously. Negation—Should remain asserted until $\overline{\text{DACK}}_n$ is asserted or the requested transaction to the peripheral occurs.</td> </tr> </table> | Timing | Assertion—Can be asserted asynchronously. Negation—Should remain asserted until $\overline{\text{DACK}}_n$ is asserted or the requested transaction to the peripheral occurs. | | |
| Timing | Assertion—Can be asserted asynchronously. Negation—Should remain asserted until $\overline{\text{DACK}}_n$ is asserted or the requested transaction to the peripheral occurs. | | | |
| $\overline{\text{DACK}}[0:1]$ | O | DMA acknowledge signals, one per channel. The DMA acknowledge signal reflects the value of $\text{DMAMR}_n[\text{CS}]$. | | |
| | | <table border="0"> <tr> <td style="vertical-align: top;">State Meaning</td> <td>Asserted—A DMA transfer is active. Negated—The DMA transfer is halted or complete.</td> </tr> </table> | State Meaning | Asserted—A DMA transfer is active. Negated—The DMA transfer is halted or complete. |
| | | State Meaning | Asserted—A DMA transfer is active. Negated—The DMA transfer is halted or complete. | |
| <table border="0"> <tr> <td style="vertical-align: top;">Timing</td> <td>Assertion—Asserted asynchronously when a DMA transfer is started or resumed in the internal control logic. Negation—Negated asynchronously when a DMA transfer is halted or completed in the internal control logic. Note that there may still be outstanding write transactions in the bus pipeline after the negation of $\overline{\text{DACK}}_n$.</td> </tr> </table> | Timing | Assertion—Asserted asynchronously when a DMA transfer is started or resumed in the internal control logic. Negation—Negated asynchronously when a DMA transfer is halted or completed in the internal control logic. Note that there may still be outstanding write transactions in the bus pipeline after the negation of $\overline{\text{DACK}}_n$. | | |
| Timing | Assertion—Asserted asynchronously when a DMA transfer is started or resumed in the internal control logic. Negation—Negated asynchronously when a DMA transfer is halted or completed in the internal control logic. Note that there may still be outstanding write transactions in the bus pipeline after the negation of $\overline{\text{DACK}}_n$. | | | |
| $\overline{\text{DDONE}}[0:1]$ | O | DMA done signals, one per channel. The DMA done signal indicates that the DMA transfer has completed. | | |
| | | <table border="0"> <tr> <td style="vertical-align: top;">State Meaning</td> <td>Asserted—A DMA transfer is complete. Negated—A DMA transfer is active or halted.</td> </tr> </table> | State Meaning | Asserted—A DMA transfer is complete. Negated—A DMA transfer is active or halted. |
| | | State Meaning | Asserted—A DMA transfer is complete. Negated—A DMA transfer is active or halted. | |
| <table border="0"> <tr> <td style="vertical-align: top;">Timing</td> <td>Assertion—Asserted asynchronously when a DMA transfer is completed in the internal control logic. Note that there may still be outstanding write transactions in the bus pipeline after the assertion of $\overline{\text{DDONE}}_n$. Negation—Negated asynchronously when a DMA transfer begins in the internal control logic.</td> </tr> </table> | Timing | Assertion—Asserted asynchronously when a DMA transfer is completed in the internal control logic. Note that there may still be outstanding write transactions in the bus pipeline after the assertion of $\overline{\text{DDONE}}_n$. Negation—Negated asynchronously when a DMA transfer begins in the internal control logic. | | |
| Timing | Assertion—Asserted asynchronously when a DMA transfer is completed in the internal control logic. Note that there may still be outstanding write transactions in the bus pipeline after the assertion of $\overline{\text{DDONE}}_n$. Negation—Negated asynchronously when a DMA transfer begins in the internal control logic. | | | |

12.3 DMA Memory Map/Register Definition

Table 12-2 lists the address and access of the memory map module.

Table 12-2. Module Memory Map

| Offset | Register | Access | Reset | Section/Page |
|--|--|--------|-----------|--------------------------------|
| DMA—Block Base Address 0x0_8100 | | | | |
| 0x0_8030 | OMISR—Outbound message interrupt status register | Mixed | All zeros | 12.4.1/12-4 |
| 0x0_8034 | OMIMR—Outbound message interrupt mask register | R/W | All zeros | 12.4.2/12-5 |
| 0x0_8050 | IMR0—Inbound message register 0 | R/W | All zeros | 12.4.3/12-6 |
| 0x0_8054 | IMR1—Inbound message register 1 | R/W | All zeros | 12.4.3/12-6 |
| 0x0_8058 | OMR0—Outbound message register 0 | R/W | All zeros | 12.4.4/12-6 |
| 0x0_805C | OMR1—Outbound message register 1 | R/W | All zeros | 12.4.4/12-6 |
| 0x0_8060 | ODR—Outbound doorbell register | R/W | All zeros | 12.4.5/12-6 |
| 0x0_8068 | IDR—Inbound doorbell register | R/W | All zeros | 12.4.5/12-6 |
| 0x0_8080 | IMISR—Inbound message interrupt status register | Mixed | All zeros | 12.4.6/12-8 |
| 0x0_8084 | IMIMR—Inbound message interrupt mask register | R/W | All zeros | 12.4.7/12-9 |
| 0x0_8100 | DMAMR0—DMA 0 mode register | R/W | All zeros | 12.4.8.1/12-10 |
| 0x0_8104 | DMASR0—DMA 0 status register | R/W | All zeros | 12.4.8.2/12-12 |
| 0x0_8108 | DMACDAR0—DMA 0 current descriptor address register | R/W | All zeros | 12.4.8.3/12-13 |
| 0x0_8110 | DMASAR0—DMA 0 source address register | R/W | All zeros | 12.4.8.4/12-14 |
| 0x0_8118 | DMADAR0—DMA 0 destination address register | R/W | All zeros | 12.4.8.5/12-14 |
| 0x0_8120 | DMABCR0—DMA 0 byte count register | R/W | All zeros | 12.4.8.6/12-15 |
| 0x0_8124 | DMANDAR0—DMA 0 next descriptor address register | R/W | All zeros | 12.4.8.7/12-15 |
| 0x0_8180 | DMAMR1—DMA 1 mode register | R/W | All zeros | 12.4.8.1/12-10 |
| 0x0_8184 | DMASR1—DMA 1 status register | R/W | All zeros | 12.4.8.2/12-12 |
| 0x0_8188 | DMACDAR1—DMA 1 current descriptor address register | R/W | All zeros | 12.4.8.3/12-13 |
| 0x0_8190 | DMASAR1—DMA 1 source address register | R/W | All zeros | 12.4.8.4/12-14 |
| 0x0_8198 | DMADAR1—DMA 1 destination address register | R/W | All zeros | 12.4.8.5/12-14 |
| 0x0_81A0 | DMABCR1—DMA 1 byte count register | R/W | All zeros | 12.4.8.6/12-15 |
| 0x0_81A4 | DMANDAR1—DMA 1 next descriptor address register | R/W | All zeros | 12.4.8.7/12-15 |
| 0x0_8200 | DMAMR2—DMA 2 mode register | R/W | All zeros | 12.4.8.1/12-10 |
| 0x0_8204 | DMASR2—DMA 2 status register | R/W | All zeros | 12.4.8.2/12-12 |
| 0x0_8208 | DMACDAR2—DMA 2 current descriptor address register | R/W | All zeros | 12.4.8.3/12-13 |
| 0x0_8210 | DMASAR2—DMA 2 source address register | R/W | All zeros | 12.4.8.4/12-14 |
| 0x0_8218 | DMADAR2—DMA 2 destination address register | R/W | All zeros | 12.4.8.5/12-14 |

12.4.3 Inbound Message Registers (IMR0–IMR1)

The inbound message registers can be read from the PCI bus and the CSB in both host and agent modes. They can be written only from the PCI bus. [Figure 12-4](#) shows the IMR0 and IMR1 fields.




Figure 12-4. Inbound Message Registers (IMR0, IMR1)

[Table 12-5](#) describes the IMR_n register.

Table 12-5. IMR0 and IMR1 Field Descriptions

| Bits | Name | Description |
|------|-------------------|---|
| 31–0 | IMSG _n | Inbound message <i>n</i> . Contains generic data to be passed between the local processor and external hosts. |

12.4.4 Outbound Message Registers (OMR0–OMR1)

The outbound message registers can be read from the PCI bus and the CSB in both host and agent modes. They can be written only from the CSB.




Figure 12-5. Outbound Message Registers (OMR0–OMR1)

[Table 12-6](#) describes the OMR_n registers.

Table 12-6. OMR0 and OMR1 Field Descriptions

| Bits | Name | Description |
|------|-------------------|--|
| 31–0 | OMSG _n | Outbound message <i>n</i> . Contains generic data to be passed between the local processor and external hosts. |

12.4.5 Doorbell Registers

The following sections describe the outbound and inbound doorbell registers.

12.4.5.1 Outbound Doorbell Register (ODR)

ODR is accessible from the PCI bus and the CSB in both host and agent modes. Figure 12-6 shows the ODR_n fields.

Offset: 0x060

Access: User Read/Write




Figure 12-6. Outbound Doorbell Register (ODR)

Table 12-7 describes the ODR registers.

Table 12-7. ODR Field Descriptions

| Bits | Name | Description |
|-------|------------------|---|
| 31–29 | — | Reserved |
| 28–0 | ODR _n | Outbound doorbell <i>n</i> . Write 1 from the CSB to set. Write 1 from the PCI bus to clear. Writing 0 has no effect. (Writing a bit in this register from the CSB causes an interrupt ($\overline{\text{PCI_INTA}}$) to be generated.) |

12.4.5.2 Inbound Doorbell Register (IDR)

IDR is accessible from the PCI bus and the CSB in both host and agent modes. Figure 12-7 shows the IDR fields.

Offset: 0x068

Access: User read/write




Figure 12-7. Inbound Doorbell Register (IDR)

Table 12-8 describes the IDR registers.

Table 12-8. IDR Field Descriptions

| Bits | Name | Descriptions |
|------|---------|---|
| 31 | IMC | Inbound machine check. Write 1 from the PCI bus to set. Write 1 from the CSB to clear. Writing 0 has no effect. Writing this bit from the PCI bus causes a machine check interrupt to be generated to the local processor. |
| 30–0 | IDR n | Inbound doorbell n . Write 1 from the PCI bus to set. Write 1 from the CSB to clear. Writing 0 has no effect. Writing a bit in this register from the PCI bus causes an interrupt to be generated to the local processor. |

12.4.6 Inbound Message Interrupt Status Register (IMISR)

The IMISR contains the interrupt status of the doorbell and message register events. Writing a 1 to IM1I clears the bit. The events are generated by the PCI masters.

Figure 12-8 shows the IMISR fields.




Figure 12-8. Inbound Message Interrupt Status Register (IMISR)

Table 12-9 describes the IMISR register.

Table 12-9. IMISR Field Descriptions


| Bits | Name | Descriptions |
|------|------|--|
| 31–5 | — | Reserved |
| 4 | MCI | Machine check interrupt. Indicates whether a machine check interrupt condition was generated by setting the IDR[31]. The interrupt is cleared by writing a 1 to IDR[IMC] from the CSB. 0 No machine check interrupt 1 There is a machine check interrupt |
| 3 | IDI | Inbound doorbell interrupt. Indicates whether an inbound doorbell interrupt occurred. 0 No inbound doorbell interrupt 1 There is an inbound doorbell interrupt |
| 2 | — | Reserved |

Table 12-9. IMISR Field Descriptions (continued)

| Bits | Name | Descriptions |
|------|------|--|
| 1 | IM1I | Inbound message 1 interrupt. Indicates whether an inbound message 1 interrupt occurred. Write 1 to this position to clear this bit. 0 No inbound message 1 interrupt. 1 There is an inbound message 1 interrupt. |
| 0 | IM0I | Inbound message 0 interrupt. Indicates whether an inbound message 0 interrupt occurred. Write 1 to this position to clear this bit. 0 No inbound message 0 interrupt. 1 There is an inbound message 0 interrupt. |

12.4.7 Inbound Message Interrupt Mask Register (IMIMR)

This register contains the interrupt mask of the doorbell and message register events generated by the PCI master. [Figure 12-9](#) shows the IMIMR fields.


Figure 12-9. Inbound Message Interrupt Mask Register (IMIMR)

[Table 12-10](#) describes the IMISR register.

Table 12-10. IMIMR Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 31–5 | — | Reserved |
| 4 | MCIM | Machine check interrupt mask. 0 Machine check interrupt from the IDR is allowed 1 Machine check interrupt is masked. IMISR[MC1] is cleared |
| 3 | IDIM | Inbound doorbell interrupt mask. 0 Inbound doorbell interrupt is allowed 1 Inbound doorbell interrupt is masked. IMISR[IDI] is cleared. |
| 2 | — | Reserved |
| 1 | IM1IM | Inbound message 1 interrupt mask. 0 Inbound message 1 interrupt is allowed 1 Inbound message 1 interrupt is masked. IMISR[IM1] is cleared |
| 0 | IM0IM | Inbound message 0 interrupt mask. 0 Inbound message 0 interrupt is allowed 1 Inbound message 0 interrupt is masked. IMISR[IM0] is cleared |

Table 12-11. DMAMR_n Field Descriptions (continued)

| Bits | Name | Description |
|-------|-------|---|
| 20 | DMSEN | Direct mode snoop enable. This bit controls snooping of direct mode DMA transactions. 0 Snooping is disabled 1 Snooping is enabled |
| 19 | IRQS | Interrupt steer. This bit determines the destination of the DMA interrupts. 0 All DMA interrupts are routed to the on-chip interrupt controller 1 All DMA interrupts are routed to the PCI bus through <code>PCI_INTA</code> |
| 18 | EMSEN | External master start enable. This bit is cleared when the DMA transfer has completed, so it must be set again for each transfer. 0 The channel is started by software setting the CS bit 1 The channel is started by hardware asserting the DREQ pin |
| 17–16 | DAHTS | Destination address hold transfer size. This field indicates the transfer size used for each transaction when DAHE is 1. The byte count register must be in multiples of the size, and the destination address register must be aligned based on the size. 00 1 byte 01 2 bytes 10 4 bytes 11 8 bytes |
| 15–14 | SAHTS | Source address hold transfer size. This field indicates the transfer size used for each transaction when SAHE is 1. The byte count register must be in multiples of the size, and the source address register must be aligned based on the size. 00 1 byte 01 2 bytes 10 4 bytes 11 8 bytes |
| 13 | DAHE | Destination address hold enable. This bit allows the DMA controller to hold the destination address constant for every transfer. The size used for transfer is indicated by DAHTS. Note that hardware supports only aligned transfers for this feature. 0 Do not hold the destination address constant 1 Hold the destination address constant Note: The DMA does not support address hold when the external trigger mode is selected (<code>EMSEN = 1</code>). Note: The DMA does not support address hold for both the source and the destination at the same transfer. |
| 12 | SAHE | Source address hold enable. This bit allows the DMA controller to hold the source address constant for every transfer. The size used for transfer is indicated by SAHTS. Note that hardware supports only aligned transfers for this feature. 0 Do not hold the source address constant 1 Hold the source address constant Note: The DMA does not support address hold when the external trigger mode is selected (<code>EMSEN = 1</code>). Note: The DMA does not support address hold for both the source and the destination at the same transfer. |
| 11–10 | PRC | PCI read command. This field indicates the type of PCI read command to use. 00 Reserved 01 PCI read line 10 PCI read multiple 11 Reserved |
| 9–8 | — | Reserved |

Table 12-12. DMASR_n Field Descriptions (continued)

| Bits | Name | Description |
|------|-------|--|
| 6–3 | — | Reserved |
| 2 | CB | Channel busy. This bit indicates whether the channel is busy. It is cleared as a result of any of the following conditions: an error or completion of the DMA transfer. 0 No DMA transfer is currently in progress 1 A DMA transfer is currently in progress |
| 1 | EOSI | End-of-segment interrupt. After transferring a segment of data, if the DMACDAR _n [EOSIE] bit in the current descriptor address register is set, this bit is set and an interrupt is generated. |
| 0 | EOCDI | End-of-chain/direct interrupt. When the last DMA transfer is finished, either in chaining or direct mode, if DMAMR[EOTIE] is set, this bit is set and an interrupt is generated. |

12.4.8.3 DMA Current Descriptor Address Register (DMACDAR_n)

DMACDAR_n contains the address of the current segment descriptor being transferred. In chaining mode, software must initialize this register to point to the first descriptor in the chain. After processing the first descriptor, the DMA controller moves the contents of the next descriptor address register into DMACDAR, loads the following descriptor into DMANDAR, and executes the current transfer.

Figure 12-12 shows the DMACDAR_n fields.



Figure 12-12. DMA Current Descriptor Address Register (DMACDAR_n)

Table 12-13 describes the DMACDAR_n register.

Table 12-13. DMACDAR_n Field Descriptions

| Bits | Name | Description |
|------|-------|---|
| 31–5 | CDA | Current descriptor address. This field contains the current descriptor address of the segment descriptor in memory. It must be aligned on an 8-word boundary. |
| 4 | SNEN | Snoop enable. 0 Snooping is disabled on DMA transactions of the current segment. 1 Snooping is enabled on DMA transactions of the current segment. |
| 3 | EOSIE | End-of-segment interrupt enable 0 No end-of-segment interrupt is generated. 1 An interrupt is generated when the current DMA transfer for the current descriptor is finished. |
| 2–0 | — | Reserved |

12.4.8.4 DMA Source Address Register (DMASAR_n)

DMASAR_n indicates the address from which the DMA controller will be reading data. The software must ensure that this is a valid memory address. [Figure 12-13](#) shows the DMASAR_n.



Figure 12-13. DMA Source Address Register (DMASAR_n)

[Table 12-14](#) describes the DMASAR_n register.

Table 12-14. DMASAR_n Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 31–0 | SA | Source address of DMA transfer. The content of this field is updated after each DMA read operation. |

12.4.8.5 DMA Destination Address Register (DMADAR_n)

DMADAR_n indicates the address to which the DMA controller will be writing data. The software must ensure that this is a valid memory address. [Figure 12-14](#) shows the DMADAR_n fields.




Figure 12-14. DMA Destination Address Register (DMADAR_n)

[Table 12-15](#) describes the DMADAR_n register.

Table 12-15. DMASAR_n Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 31–0 | DA | Destination address of DMA transfer. Updated after each DMA write operation. |

Table 12-17. DMANDAR_n Field Descriptions (continued)

| Bits | Name | Descriptions |
|------|------|--|
| 2–1 | — | Reserved |
| 0 | EOTD | End-of-transfer descriptor. 0 This descriptor contains a link to another descriptor. 1 This descriptor is the last to be executed. |

12.4.8.8 DMA General Status Register (DMAGSR)

DMAGSR provides faster access to the status bits by combining the status bits of all of the DMA channels into one register. Each byte of this register provides the value of bits 7–0 of a channel’s DMA status register. These bits are cleared by writing to the individual DMA status registers. [Figure 12-17](#) shows the DMAGSR fields.




Figure 12-17. DMA General Status Register (DMAGSR)

12.5 Functional Description

12.5.1 Message Unit

An embedded processor is often part of a larger system containing many processors and distributed memory. These processors tend to work on tasks independent of the host and other peripheral processors in the system. Because of the independent nature of the tasks, it is necessary to provide a communication mechanism between the peripheral processors and the rest of the system. One such method is the use of messages. This block provides a messaging unit to further facilitate communications between host and peripheral. The message unit uses generic messages and doorbell registers.

12.5.1.1 Messaging Registers (IMR0–IMR1, OMR0–OMR1)

There are two 32-bit inbound message registers (IMR0–IMR1) and two 32-bit outbound message registers (OMR0–OMR1). IMR0 and IMR1 allow a remote host or PCI master to write a 32-bit value that, in turn, causes an interrupt request to the on-chip interrupt controller that drives an interrupt line to the local processor. OMR0 and OMR1 allow the local processor to write an outbound message which, in turn, causes the outbound interrupt signal $\overline{\text{PCI_INTA}}$ to assert.

The interrupt to the local processor is cleared by writing 1 to the appropriate IMISR bit. The interrupt to PCI ($\overline{\text{PCI_INTA}}$) is cleared by writing 1 to the appropriate OMISR bit.

12.5.1.2 Doorbell Registers (IDR and ODR)

This block contains the inbound doorbell register (IDR) and the outbound doorbell register (ODR). The inbound doorbell allows a remote processor to set a bit in the register from the PCI bus. This, in turn, generates an interrupt request to the on-chip interrupt controller that drives an interrupt line to the local processor. The local processor can write to the ODR, which causes the outbound interrupt signal `PCI_INTA` to assert, thus interrupting the remote processor on the PCI bus.

The interrupt to the local processor is cleared by writing 1 to the appropriate IDR bit. The interrupt to PCI (`PCI_INTA`) is cleared by writing 1 to the appropriate ODR bit.

12.5.2 DMA Controller

The DMA controller transfers blocks of data independent of the local processor or PCI hosts. Data movement occurs on the PCI bus and/or CSB. The DMA module has four high-speed DMA channels, which share buffer space in the IOS to facilitate the gathering and sending of data. Both the local processor and PCI masters can initiate a DMA transfer.

Features of the DMA controller include the following:

- Four channels
- Concurrent execution across multiple channels with programmable bandwidth control
- All channels are accessible by local processor and remote PCI masters
- Unaligned transfer capability
- Data chaining and direct mode
- Interrupt on completed segment, chain, and error

Figure 12-18 shows a diagram of the DMA controller in the integrated device.




Figure 12-18. DMA Controller Block Diagram

12.5.3 DMA Operation

The DMA controller operates in the following two modes:

- Direct mode, in direct mode, the DMA controller does not read a chain of descriptors from memory but instead uses the current parameters in the DMA registers to start a DMA transfer. The DMA

transfer finishes after all the bytes specified in the byte count register have been transferred. See [Section 12.6.1, “Initialization Steps in Direct Mode,”](#) for more details on initialization steps.

- Chaining mode, in chaining mode, the DMA controller loads descriptors from memory prior to a DMA transfer. The DMA controller begins the transfer according to the descriptor information loaded for each segment. Once the current segment is finished, the DMA controller reads the next descriptor from memory and begins another DMA transfer. The process is finished if the current descriptor is the last one in the chain. See [Section 12.6.2, “Initialization Steps in Chaining Mode,”](#) for more details on initialization steps.

In both modes, setting the start bit in the DMA mode register begins the DMA transfer.

The DMA controller supports misaligned transfers for both the source and destination addresses. It gathers data beginning at the source address and aligns the data accordingly before sending it to the destination address. The DMA controller assumes that the source and destination addresses are valid PCI or CSB memory addresses.

Accesses to CSB memory depend on the alignment of the source and destination addresses and the size of the transfer. The DMA controller transfers a full cache line whenever possible. On misaligned addresses, full cache line transfers (32 byte bursting) occur if the source and destination offsets end in the same byte offset. For example, if the destination address is 0x4000_1050 and the source address is 0x9000_2050, the transfer bursts because both end in 0x50. However, if the destination address is 0x4000_1050 and the source is 0x9000_2000, the transfer does not burst because the last byte offset is not the same. On misaligned destination addresses, subtransfers of less than a cache line occur on the initial and final beats of the transfer while full cache lines occur on intermediate beats.

Configuring a DMA channel for address hold mode $DMAMR_n$ precludes cache line transfers.

PCI memory read operations depend on the PRC (PCI read command) field in the mode register, the alignment of the source address, and the size of the transfer. The DMA controller attempts to read a full cache line whenever possible. Writing to PCI memory depends on the alignment of the destination address and the size of the transfer.

12.5.3.1 External Control

The DMA transfer of any channel, in either direct mode or chaining mode, can be controlled by the DMA request input signals. External control is enabled by setting the external master start enable (EMSEN) bit instead of the channel start (CS) bit in the DMA mode register ($DMAMR_n$).

When using external control, the following restrictions apply:

- Both the source and destination addresses must be aligned to 32-byte boundaries.
- In chaining mode all byte count values except the last must be multiples of 32 bytes.

A falling edge on \overline{DREQ}_n sets $DMAMR_n[CS]$ to start the transfer and asserts the corresponding \overline{DACK}_n output signal. The number of cache lines specified by the DMA request count (DRCNT) bit in the DMA mode register ($DMAMR_n$)—or the remaining byte count, if smaller— is transferred, and the CS bit and the \overline{DACK}_n output signal are then cleared and negated to halt the transfer until the next \overline{DREQ}_n assertion.

When using the \overline{DACK}_n handshake signal, \overline{DREQ}_n should remain asserted until \overline{DACK}_n is asserted, at which point \overline{DREQ}_n may be negated. \overline{DREQ}_n may be asserted again to resume the transfer or start a new transfer once \overline{DACK}_n has been negated.

If the \overline{DACK}_n handshake signal is not used, \overline{DREQ}_n should remain asserted until the first transaction of the DMA transfer appears on the external interface, at which point \overline{DREQ}_n may be negated. \overline{DREQ}_n may be asserted again to resume the transfer or start a new transfer as early as one clock cycle after its negation, even if the current transfer is still active. The DMA controller is able to record a new assertion of \overline{DREQ}_n while a transfer is in progress, with the effect of setting $DMAMR_n[CS]$ again once the transfer has been halted.

Once a transfer has been completed, $DMAMR_n[EMSEN]$ is cleared by the DMA controller, and $DMAMR_n[CS]$ will not be set again until $DMAMR_n[EMSEN]$ has been set by the user. The assertion of \overline{DREQ}_n and the setting of $DMAMR_n[EMSEN]$ may occur in either order; whichever occurs later will trigger the DMA transfer.

The \overline{DDONE}_n output signal is asserted when the DMA transfer has completed, that is, all bytes specified in the byte count register or the descriptors have been transferred. This signal could be used as an indication that the number of cache lines transferred might be smaller than that specified by the $DRCNT$ field.

NOTE

The \overline{DACK}_n and \overline{DDONE}_n output signals are intended as handshake signals for \overline{DREQ}_n . They are asserted and negated according to the DMA controller internal logic. These signals are not synchronized to the transactions appearing on the external pins of the device. Specifically, the negation of \overline{DACK}_n or the assertion of \overline{DDONE}_n does not mean that all transactions of the DMA transfer have been completed as seen on the external pins.

See [Section 12.6.3, “Initialization Steps in Direct Mode with External Control,”](#) and [Section 12.6.4, “Initialization Steps in Chaining Mode with External Control,”](#) for more details on initialization steps.

12.5.3.2 DMA Coherency

The four DMA channels use up to four cache lines (128 bytes) of buffer space in the IOS in addition to 16 bytes of local buffer space. Because no address snooping occurs in these internal queues, data posted in these queues is not visible to the rest of the system while a DMA transfer is in progress. It is the responsibility of application software to ensure the coherency of the region being transferred during the DMA process.

Snooping of the CPU or processor data cache is selectable during DMA transactions. A snoop bit is provided in the DMA current descriptor address register ($DMACDAR_n$) and the DMA next descriptor address register ($DMANDAR_n$) that allows software to control when the cache is snooped on a per segment basis.

12.5.3.3 Halt and Error Conditions

DMA transfers are halted either by clearing the CS (channel start) bit in the DMA mode register (DMAMR n) or when encountering an error condition. In either case, the application software can do one of the following:

- Continue the DMA transfer
- Reconfigure the DMA for a new transfer
- Leave the channel in the halted state

When a DMA channel is halted, its programming model is completely accessible. If the DMA is halted due to an error condition, the TE (transfer error) bit in the DMA status register (DMASR n) must be cleared before the transfer can be resumed or a new transfer initiated. Note that the TE bit is not cleared automatically by hardware.

12.5.4 DMA Segment Descriptors

DMA segment descriptors contain the source and destination addresses of the data segment, the segment byte count, and a link to the next descriptor. Segment descriptors are built on cache-line (32-byte) boundaries in either CSB or PCI memory and are linked together into chains using the next-descriptor-address field.

Table 12-18. DMA Segment Descriptor Fields

| Descriptor Field | Description |
|-------------------------|---|
| Source address | Contains the source address of the DMA transfer. After the DMA controller reads the descriptor from memory, this field will be loaded into the DMA source address register (DMASAR n). |
| Destination address | Contains the destination address of the DMA transfer. After the DMA controller reads the descriptor from memory, this field will be loaded into the DMA destination address register (DMADAR n). |
| Next descriptor address | Points to the next descriptor in memory. After the DMA controller reads the descriptor from memory, this field will be loaded into the DMA next descriptor address register (DMANDAR n). |
| Byte count | Contains the number of bytes to transfer. After the DMA controller reads the descriptor from memory, this field will be loaded into the DMA byte count register (DMABCR n). |

Application software initializes the current DMA current descriptor address register (DMACDAR n) to point to the first descriptor in the chain. For each descriptor in the chain, the DMA controller starts a new DMA transfer with the control parameters specified by the descriptor. The DMA controller traverses the descriptor chain until reaching the last descriptor (with its EOTD bit set).

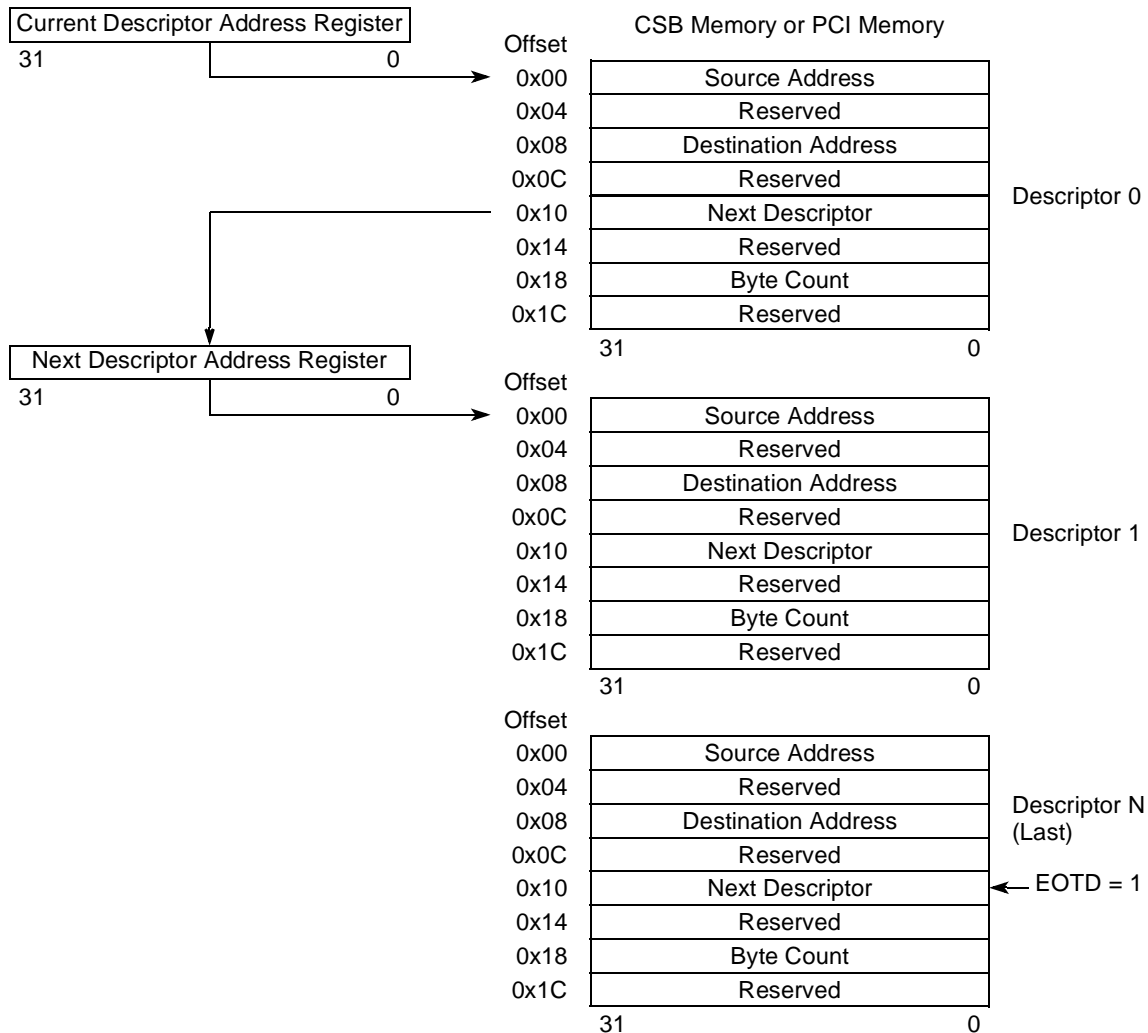


Figure 12-19. DMA Chain of Segment Descriptors

12.5.4.1 Descriptor in Big-Endian Mode

In big-endian mode, the descriptor in CSB memory should be programmed such that data appears in ascending significant-byte order. If segment descriptors are written to memory located in the CSB, they should be treated like they are translated from big-endian to little-endian mode.

Example: Big-endian mode descriptor's data structure. Note that the descriptor structure must be aligned on an 8-word boundary.

```
struct {
    double a;          /* 0x1122334455667788 double word*/
    double b;          /* 0x55667788aabbccdd double word*/
    double c;          /* 0x8765432101234567 double word */
    double d;          /* 0x0123456789abcdef double word */
} Descriptor;
Results: Source Address = 0x44332211 <MSB..LSB>
        Destination Address = 0x88776655 <MSB..LSB>
        Next Descriptor Address = 0x21436587 <MSB..LSB>
        Byte Count = 0x67452301 <MSB..LSB>
```

12.5.4.2 Descriptor in Little-Endian Mode

In little-endian mode, each segment descriptor should be programmed in descending significant-byte order.

Example: Little-endian mode descriptor's data structure. Note that the descriptor structure must be aligned on an 8-word boundary.

```
struct {
    double a;          /* 0x8877665544332211 double word*/
    double b;          /* 0x1122334488776655 double word*/
    double c;          /* 0x7654321012345678 double word */
    double d;          /* 0x0123456776543210 double word */
} Descriptor;
Results: Source Address = 0x44332211 <MSB..LSB>
        Destination Address = 0x88776655 <MSB..LSB>
        Next Descriptor Address = 0x12345678 <MSB..LSB>
        Byte Count = 0x76543210 <MSB..LSB>
```

12.6 Initialization/Application Information

12.6.1 Initialization Steps in Direct Mode

The initialization steps of a DMA transfer in direct mode are described as follows:

1. Poll the CB (channel busy) bit in the DMA status register (DMASR n) to make sure the DMA channel is idle.
2. Initialize the DMASAR n , DMADAR n , and the DMABCR n .
3. Initialize DMAMR n [CTM]) to indicate direct mode. Other control parameters in the mode register can also be initialized here if necessary.
4. First clear then set the DMAMR n [CS] to start the DMA transfer.

12.6.2 Initialization Steps in Chaining Mode

The initialization steps of a DMA transfer in chaining mode are described as follows:

1. Build a chain of descriptor segments in memory. Refer to [Section 12.5.4, “DMA Segment Descriptors.”](#)

2. Poll the $DMASR_n[CB]$ to make sure the DMA channel is idle.
3. Initialize the $DMACDAR_n$ to point to the first descriptor in the chain.
4. Initialize the $DMAMR_n[CTM]$ to indicate chaining mode. Other control parameters in the mode register can also be initialized here if necessary.
5. First clear then set the $DMAMR_n[CS]$ to start the DMA transfer.

12.6.3 Initialization Steps in Direct Mode with External Control

The initialization steps of a DMA transfer in direct mode with external control are described as follows:

1. Poll the CB (channel busy) bit in the DMA status register ($DMASR_n$) to make sure the DMA channel is idle.
2. Initialize the DMA source address register ($DMASAR_n$), the DMA destination address register ($DMADAR_n$), and the DMA byte count register ($DMABCR_n$).
3. Set $DMAMR_n[CTM]$ to indicate direct mode, program the DRCNT field, and set the EMSEN bit. Other control parameters in the mode register can also be initialized here if necessary.

12.6.4 Initialization Steps in Chaining Mode with External Control

The initialization steps of a DMA transfer in chaining mode with external control are described as follows:

1. Build a chain of descriptor segments in memory. Refer to [Section 12.5.4, “DMA Segment Descriptors,”](#) for more information.
2. Poll the CB (channel busy) bit in the DMA status register ($DMASR_n$) to make sure the DMA channel is idle.
3. Initialize the DMA current descriptor address register ($DMACDAR_n$) to point to the first descriptor in the chain.
4. Clear the $DMAMR_n[CTM]$ to indicate chaining mode, program the DRCNT field, and set the EMSEN bit. Other control parameters in the mode register can also be initialized here if necessary.



Chapter 13

PCI Bus Interface

The PCI interface is compatible with the *PCI Local Bus Specification*, Rev. 2.3. It is beyond the scope of this manual to document the intricacies of PCI. This chapter describes the PCI controller and provides a basic description of the PCI bus operations. The specific emphasis is directed at how this device implements the PCI specification. Designers of systems incorporating PCI devices should refer to the respective specifications for a thorough description of the PCI buses.

NOTE

Much of the available PCI literature refers to a 16-bit quantity as a WORD and a 32-bit quantity as a DWORD. Because this is inconsistent with the terminology in this manual, the terms 'word' and 'double word' are not used in this chapter. Instead, the number of bits or bytes indicates the exact quantity.

13.1 PCI Introduction

The PCI controller acts as a bridge between the PCI interface and the CSB. The I/O sequencer buffers the data. [Figure 13-1](#) is a high-level block diagram of the PCI controller.




Figure 13-1. PCI Controller Block Diagram

The PCI controller connects the processor and memory system to the I/O components through the PCI system bus. This interface acts as both initiator (master) and target (slave) device. The PCI controller uses a 32-bit multiplexed, address/data bus that can run at frequencies up to 66-MHz. The interface provides address and data parity with error checking and reporting. The interface provides for three physical address spaces—64-bit address memory, 32-bit address I/O, and PCI configuration space.

Note that PCI supports up to three external masters.

The PCI interface can function as either a PCI host bridge referred to as host mode or a peripheral device on the PCI bus referred to as agent mode. See [Section 13.4.4.4, “Host Mode Configuration Access,”](#) for more information. Note that the PCI controller can be configured from the PCI bus while in agent mode.

An address translation mechanism is provided to map PCI memory windows between the PCI bus and the internal bus.

The PCI interface does not flush pending outbound writes as a result of an inbound read command. Systems must not rely on inbound reads to ensure all pending outbound writes have completed. For example, consider the case where a core writes data to a PCI device and then updates a flag in the local DDR memory indicating the write to PCI has completed. An external PCI master may misread the flag ahead of the actual write transaction's completion on the PCI bus.

13.1.1 PCI Features

The PCI controller includes the following features:

- PCI specification revision 2.3 compliant
- 32-bit PCI interface support
- Host and agent mode support
- PCI bus power management unit
- Supports accesses to all PCI address spaces
- 64-bit dual-address cycle (DAC) support (as a target only)
- Internal configuration registers accessible from PCI
- On-chip arbitration supporting three masters on PCI
- Arbiter supports two-level priority request/grant signal pairs
- Supports PCI-to-memory and memory-to-PCI streaming
- Memory prefetching of PCI read accesses and support for delayed read transactions
- Supports posting of processor-to-PCI and PCI-to-memory writes
- Supports selectable snooping for inbound transactions
- Address translation units for address mapping between host and peripheral
- Supports parity
- PCI 3.3-V compatible

13.1.2 PCI Modes of Operation

PCI controller modes of operation are determined at reset by the reset configuration word high (RCWH) as described in [Section 4.3.2, “Reset Configuration Words.”](#) [Table 13-1](#) summarizes these modes.

13.1.2.1 Host/Agent Mode Configuration

The PCI controller can function as either a PCI host bridge (referred to as host mode) or a peripheral device on the PCI bus (referred to as agent mode). Note that host/agent mode selection is determined at power-up as summarized in [Section 4.3.2.2.1, “PCI Host/Agent Configuration.”](#)

When the device powers up in host mode, all inbound configuration accesses are ignored (and thus master aborted). When the device powers up in agent mode, it acknowledges inbound configuration accesses. Note that in PCI agent mode, the PCI controller ignores all PCI memory accesses except those to the

memory-mapped registers until inbound address translation is enabled. In agent mode, configuration cycles are acknowledged if CFG_LOCK is 0 (see [Section 13.3.3.24, “PCI Function Configuration Register”](#)), either from reset configuration or after being cleared by software.

13.1.2.2 PCI Arbiter Configuration

The interface can be configured to use an on-chip or off-chip PCI arbiter. Arbitration for PCI is determined by the value in RCWH[PCIARB]. See [Section 4.3.2.2, “Reset Configuration Word High Register \(RCWHR\),”](#) for more information.

13.2 PCI External Signal Description

Table 13-2 shows the properties of the PCI signals.

Table 13-2. Signal Properties

| Name | Function | Reset State | Pull Up |
|---------------|------------------------------------|-------------------------|--------------------|
| CPCI_HS_ENUM | CompactPCI hot swap enumerator | High impedance | Required |
| CPCI_HS_ES | CompactPCI hot swap ejector switch | — | — |
| CPCI_HS_LED | CompactPCI hot swap LED | Asserted | — |
| PCI_AD[31:0] | PCIn address / data | High impedance | — |
| PCI_C/BE[3:0] | PCI bus command / byte enable | High impedance | — |
| PCI_DEVSEL | PCI device select | High impedance | Required |
| PCI_FRAME | PCI cycle frame | High impedance | Required |
| PCI_REQ[0:2] | PCI arbiter requests | Configuration-dependent | Required on inputs |
| PCI_GNT[0:2] | PCI arbiter grants | Configuration-dependent | — |
| PCI_IDSEL | PCI initialization device select | — | — |
| PCI_INTA | PCI interrupt A | High impedance | Required |
| PCI_IRDY | PCI initiator ready | High impedance | Required |
| PCI_PAR | PCI parity | High impedance | — |
| PCI_PERR | PCI parity error | High impedance | Required |
| PCI_RESET_OUT | PCI reset output | Asserted | |
| PCI_SERR | PCI system error | High impedance | Required |
| PCI_STOP | PCI stop | High impedance | Required |
| PCI_TRDY | PCI target ready | High impedance | Required |
| PCI_PME | PCI PME assertion request | High impedance | Required |

Figure 13-2 shows the external PCI signals.




Figure 13-2. PCI Interface External Signals

Table 13-3 contains detailed descriptions of the external PCI interface signals.

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|--------------|-----|---|
| CPCI_HS_ENUM | O | CompactPCI hot swap enumerator. Used for the hot swap interface to connect to the host as the enumeration request in a compact PCI system. This signal is used for agent mode only. |
| | | State Meaning Asserted—This card was inserted and needs to be configured, or this card is about to be extracted and needs to be removed from the system resources list. Negated—No action is needed. |
| | | Timing Assertion/Negation—No timing is specified |
| CPCI_HS_ES | I | CompactPCI hot swap ejector switch. Used for agent mode only. In a compact PCI system this input signal is used for the hot swap interface to connect to the ejector switch logic. |
| | | State Meaning Asserted—The switch is open. Negated—The switch is closed. |
| | | Timing Assertion/Negation—No timing is specified |
| CPCI_HS_LED | O | CompactPCI hot swap LED. Used for the hot swap interface to connect to the hot swap LED in a CompactPCI system. This signal is used for agent mode only. |
| | | State Meaning Asserted—Output is driving logic 1 to illuminate the hot swap LED. Negated—Output is driving logic 0 to turn off the hot swap LED. |
| | | Timing Assertion/Negation—No timing is specified |
| M66EN | I | 66-MHz enable. Determines the AC timing of the PCI interface. |
| | | State Meaning Asserted—The PCI interface signals use the 66-MHz PCI AC timing parameters. Negated—The PCI interface signals use the 33-MHz PCI AC timing parameters. |
| | | Timing Assertion/Negation—Constant |

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|---------------|-----|---|---|
| PCI_AD[31:0] | I/O | PCI address/data bus. During an address phase, these signals contain a physical address. During a data phase, these signals contain the data bytes. | |
| | O | Outputs for the bi-directional PCI address/data bus. | |
| | | State Meaning | Asserted/Negated—Represents the physical address during the address phase of a PCI transaction. During the data phase(s) of a PCI transaction, the PCI address/data bus contain the data being written. PCI_AD[7:0] define the LSB and, PCI_AD[31:24] define the MSB. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional PCI address/data bus. | |
| | | State Meaning | Asserted/Negated—Represents the address to be decoded as a check for device select during the address phase of a PCI transaction or the data being received during the data phase(s) of a PCI transaction. PCI_AD[7:0] define the LSB and, PCI_AD[31:24] define the MSB. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |
| PCI_C/BE[3:0] | I/O | PCI bus command/byte enable. | |
| | O | Outputs for the bi-directional command/byte enable. | |
| | | State Meaning | Asserted/Negated—During the address phase, PCI_CBE[3:0], define the bus command. Byte enables determine which byte lanes carry meaningful data for PCI bus data phases. The PCI_CBE[0] signal applies to the LSB. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional command/byte enable. | |
| | | State Meaning | Asserted/Negated—During the address phase, PCI_CBE[3:0], indicate the command that another master is sending. During the PCI bus data phase, PCI_CBE[3:0], indicate which byte lanes are valid. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |
| PCI_DEVSEL | I/O | PCI device select. | |
| | O | Outputs for the bi-directional device select. | |
| | | State Meaning | Asserted—The PCI controller has decoded the address and is the target of the current access. Negated—The PCI controller has decoded the address and is not the target of the current access. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional device select. | |
| | | State Meaning | Asserted—Some PCI agents (other than this PCI controller) have decoded its address as the target of the current access. Negated—No PCI agent has been selected. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|-----------------------------------|-----|---|---|
| $\overline{\text{PCI_FRAME}}$ | I/O | PCI cycle frame. Used by the current PCI master to indicate the beginning and duration of an access. | |
| | O | Outputs for the bi-directional frame. | |
| | | State Meaning | Asserted—The PCI controller acting as a PCI master which is initiating a bus transaction. While $\overline{\text{PCI_FRAME}}$ is asserted, data transfers may continue. Negated—If $\overline{\text{PCI_IRDY}}$ is asserted, indicates that the PCI transaction is in the final data phase; if $\overline{\text{PCI_IRDY}}$ is negated, indicates that the PCI bus is idle. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional frame. | |
| | | State Meaning | Asserted—Another PCI master is initiating a bus transaction. Negated—The transaction is in the final data phase or that the bus is idle. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |
| $\overline{\text{PCI_GNT0}}$ | I/O | PCI arbiter grants. Output signal on this PCI controller when the arbiter is enabled. Input signal when the arbiter is disabled. Note: $\overline{\text{PCI_GNT}}[0]$ is a point-to-point signal. Every master has its own bus grant signal. | |
| | O | Outputs for the bi-directional arbiter grants. | |
| | | State Meaning | Asserted—The PCI controller granted control of the PCI bus to agent 0. Negated—The PCI controller did not grant control of the PCI bus to agent 0. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional arbiter grants. | |
| | | State Meaning | Asserted—The PCI controller has been granted control of the PCI bus by an external arbiter. Negated—The PCI controller has not been granted control of the PCI bus by an external arbiter. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |
| $\overline{\text{PCI_GNT}}[1:2]$ | O | PCI arbiter grants. Output signals on this PCI controller when the arbiter is enabled. Note that $\overline{\text{PCI_GNT}}n$ is a point-to-point signal. Every master has its own bus grant signal. | |
| | | State Meaning | Asserted—The PCI controller granted control of the PCI bus to agent n . Negated—The PCI controller did not grant control of the PCI bus to agent n . |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| $\overline{\text{PCI_IDSEL}}$ | I | PCI initialization device select. Used as a chip select during a PCI configuration cycle in agent mode. This signal should be tied low in host mode. | |
| | | State Meaning | Asserted—The PCI controller is being selected as a target of a configuration read or write transactions. Negated—The PCI controller is not being selected as a target of configuration read or write transactions. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|-------------------------------|---|---|
| $\overline{\text{PCI_INTA}}$ | O | PCI interrupt A. |
| | State Meaning | Asserted—The PCI controller signals an interrupt to the PCI host. Negated—The PCI controller is not currently signalling an interrupt. |
| | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| $\overline{\text{PCI_IRDY}}$ | I/O | PCI initiator ready. This signal is driven by the PCI controller when it is the initiator of a PCI transfer. |
| | O | Outputs for the bi-directional initiator ready. |
| | State Meaning | Asserted—The PCI controller, acting as a PCI master, can complete the current data phase of a PCI transaction. During a write, this PCI controller asserts $\overline{\text{PCI_IRDY}}$ to indicate that valid data is present on $\text{PCI_AD}[31:0]$. During a read, this PCI controller asserts $\overline{\text{PCI_IRDY}}$ to indicate that it is prepared to accept data. Negated—The PCI target needs to wait before this PCI controller, acting as a PCI master, can complete the current data phase. During a write, this PCI controller negates $\overline{\text{PCI_IRDY}}$ to insert a wait cycle when it cannot provide valid data to the target. During a read, this PCI controller negates $\overline{\text{PCI_IRDY}}$ to insert a wait cycle when it cannot accept data from the target. |
| | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional initiator ready. |
| | State Meaning | Asserted—Another PCI master can complete the current data phase of a transaction. Negated—If $\overline{\text{PCI_FRAME}}$ is asserted, indicates a wait cycle from another master. If $\overline{\text{PCI_FRAME}}$ is negated, indicates that the PCI bus is idle. |
| PCI_PAR | I/O | PCI parity. |
| | O | Outputs for the bi-directional parity. |
| | State Meaning | Asserted—Odd parity across $\text{PCI_AD}[31:0]$ and $\text{PCI_CBE}[3:0]$ during address and data phases. Negated—Even parity across $\text{PCI_AD}[31:0]$ and $\text{PCI_AD}[31:0]$ during address and data phases. |
| | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional parity. |
| | State Meaning | Asserted—Odd parity driven by another PCI master or the PCI target during address and data phases. Negated—Even parity driven by another PCI master or the PCI target during address and data phases. |
| Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|-------------------------------------|-----|---|--|
| $\overline{\text{PCI_PERR}}$ | I/O | PCI parity error | |
| | O | Outputs for the bi-directional parity error. | |
| | | State Meaning | Asserted—The PCI controller, acting as a PCI agent, detected a data parity error. (driven by the PCI initiator on reads; driven by the PCI target on writes.) Negated—No error. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional parity error. | |
| | | State Meaning | Asserted—Another PCI agent detects a data parity error while this PCI controller is sourcing data (this PCI controller was acting as the PCI initiator during a write, or is acting as the PCI target during a read). Negated—No error. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |
| $\overline{\text{PCI_PME}}$ | I/O | PCI PME signal | |
| | O | Outputs for the bi-directional $\overline{\text{PCI_PME}}$ signal. This is an open-drain signal. | |
| | | State Meaning | Asserted—Indicates that a power management event has occurred Negated—Indicates that no power management event has occurred |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| $\overline{\text{PCI_REQ0}}$ | I/O | PCI bus request. Input signal on this PCI controller when the arbiter is enabled. Output signal when the arbiter is disabled. Note that $\text{PCI_REQ}n$ is a point-to-point signal. Every master has its own bus request signal. | |
| | O | Outputs for the bi-directional bus request. | |
| | | State Meaning | Asserted—The PCI controller is requesting control of the PCI bus to perform a transaction. Negated—The PCI controller does not require use of the PCI bus. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Input for the bi-directional bus request. | |
| | | State Meaning | Asserted—Agent 0 is requesting control of the PCI bus to perform a transaction. Negated—Agent 0 does not require use of the PCI bus. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| $\overline{\text{PCI_REQ}}[1:2]$ | I | PCI bus request. Input signals on this PCI controller when the arbiter is enabled. Note that $\text{PCI_REQ}[n]$ is a point-to-point signal. Every master has its own bus request signal. Following is the state meaning for the $\text{PCI_REQ}[n]$ input. | |
| | | State Meaning | Asserted—An agent n is requesting control of the PCI bus to perform a transaction. Negated—An agent n does not require use of the PCI bus. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| $\overline{\text{PCI_RESET_OUT}}$ | O | PCI reset. This signal is used only in host mode. It should be left unconnected in agent mode. | |
| | | State Meaning | Asserted—Devices on the PCI bus are in reset. Negated—Devices on the PCI bus operate normally. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |

Table 13-3. PCI Interface Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|-------------------------------|-----|---|---|
| $\overline{\text{PCI_SERR}}$ | I/O | PCI system error | |
| | O | Outputs for the bi-directional system error. | |
| | | State Meaning | Asserted—An address parity error, a target-abort (when this PCI controller is acting as the initiator), or some other system error (where the result is a catastrophic error) was detected. Negated—No error. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional system error. | |
| | | State Meaning | Asserted—A device (other than this PCI controller) has detected a catastrophic error. Negated—No error. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |
| $\overline{\text{PCI_STOP}}$ | I/O | PCI stop. | |
| | O | Outputs for the bi-directional stop. | |
| | | State Meaning | Asserted—The PCI controller, acting as a PCI target, is requesting that the initiator stop the current transaction. Negated—The current transaction can continue. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional stop. | |
| | | State Meaning | Asserted—A target is requesting that this PCI controller, as the initiator, stop the current transaction. Negated—The current transaction can continue. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |
| $\overline{\text{PCI_TRDY}}$ | I/O | PCI target ready. | |
| | O | Outputs for the bi-directional target ready. | |
| | | State Meaning | Asserted—The PCI controller, acting as a PCI target, can complete the current data phase of a PCI transaction. During a read, this PCI controller asserts $\overline{\text{PCI_TRDY}}$ to indicate that valid data is present on PCI_AD[31:0]. During a write, this PCI controller asserts $\overline{\text{PCI_TRDY}}$ to indicate that it is prepared to accept data. Negated—The PCI initiator needs to wait before this PCI controller, acting as a PCI target, can complete the current data phase. During a read, this PCI controller negates $\overline{\text{PCI_TRDY}}$ to insert a wait cycle when it cannot provide valid data to the initiator. During a write, this PCI controller negates $\overline{\text{PCI_TRDY}}$ to insert a wait cycle when it cannot accept data from the initiator. |
| | | Timing | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> |
| | I | Inputs for the bi-directional target ready. | |
| | | State Meaning | Asserted—Another PCI target is able to complete the current data phase of a transaction. Negated—A wait cycle from another target. |
| Timing | | Assertion/Negation—As specified by <i>PCI Local Bus Specification Rev 2.3</i> | |

13.3 PCI Memory Map/Register Definitions

The PCI controller has the following types of registers:

- The PCI configuration access registers. Used for generating PCI configuration accesses from the CSB. These registers, listed in [Table 13-4](#), are memory-mapped on the CSB and accessed through the IMMR window.
- The PCI memory-mapped registers. Used to manage error functions, general control and status, and address translation control for the inbound path. These registers are shown in [Table 13-5](#). They can be accessed by PCI masters via the PCI controller to the CSB through the PIMMR inbound window. Note that [Table 13-5](#) does not list outbound address translation registers; these are contained in the I/O sequencer (IOS) memory-mapped registers. See [Chapter 12](#), “DMA/Messaging Unit,” for more information.
- The PCI configuration space registers. Defined by the PCI specification. These registers are accessed by PCI masters using configuration accesses and are described in [Section 13.3.3](#), “PCI Configuration Space Registers.”

Table 13-4. PCI Configuration Access Registers

| Offset | Register | Access | Reset | Section/Page |
|---|---|--------|-------------|---------------------------------|
| PCI Configuration Access Registers—Block Base Address 0x0_8300 | | | | |
| 0x00 | PCI_CONFIG_ADDRESS | W | All zeros | 13.3.1.1/13-13 |
| 0x04 | PCI_CONFIG_DATA | R/W | All zeros | 13.3.1.2/13-14 |
| 0x08 | PCI_INT_ACK | R | N/A | 13.3.1.3/13-15 |
| 0x80 | PCIPMR0—PCI power management register 0 | R | 0x7E4B_0001 | 13.3.3.27/13-41 |
| 0x84 | PCIPMR1—PCI power management register 1 | R/W | 0x00n0_0000 | 13.3.3.28/13-42 |
| 0x88–0xFF | Reserved | — | — | — |

Table 13-5. PCI Memory-Mapped Registers

| Offset | Register | Access | Reset | Section/Page |
|---|---|--------|-----------|--------------------------------|
| PCI Controller—Block Base Address 0x0_8500 | | | | |
| PCI Error Management Registers | | | | |
| 0x00 | PCI error status register (PCI_ESR) | w1c | All zeros | 13.3.2.1/13-15 |
| 0x04 | PCI error capture disable register (PCI_ECDR) | R/W | All zeros | 13.3.2.2/13-16 |
| 0x08 | PCI error enable register (PCI_EER) | R/W | All zeros | 13.3.2.3/13-17 |
| 0x0C | PCI error attributes capture register (PCI_EATCR) | R/W | All zeros | 13.3.2.4/13-18 |
| 0x10 | PCI error address capture register (PCI_EACR) | R | All zeros | 13.3.2.5/13-19 |
| 0x14 | PCI error extended address capture register (PCI_EEACR) | R | All zeros | 13.3.2.6/13-20 |
| 0x18 | PCI error data capture register (PCI_EDCR) | R/W | All zeros | 13.3.2.7/13-20 |

Table 13-5. PCI Memory-Mapped Registers (continued)

| Offset | Register | Access | Reset | Section/Page |
|---|--|--------|-----------|---------------------------------|
| PCI Control and Status Registers | | | | |
| 0x20 | PCI general control register (PCI_GCR) | R/W | All zeros | 13.3.2.8/13-20 |
| 0x24 | PCI error control register (PCI_ECR) | R/W | All zeros | 13.3.2.9/13-21 |
| 0x28 | PCI general status register (PCI_GSR) | R | All zeros | 13.3.2.10/13-22 |
| PCI Inbound ATU Registers | | | | |
| 0x38 | PCI inbound translation address register 2 (PITAR2) | R/W | All zeros | 13.3.2.11/13-22 |
| 0x3C | Reserved | — | — | — |
| 0x40 | PCI inbound base address register 2 (PIBAR2) | R/W | All zeros | 13.3.2.12/13-23 |
| 0x44 | PCI inbound extended base address register 2 (PIEBAR2) | R/W | All zeros | 13.3.2.13/13-24 |
| 0x48 | PCI inbound window attributes register 2 (PIWAR2) | R/W | All zeros | 13.3.2.14/13-24 |
| 0x50 | PCI inbound translation address register 1 (PITAR1) | R/W | All zeros | 13.3.2.11/13-22 |
| 0x54 | Reserved | — | — | — |
| 0x58 | PCI inbound base address register 1 (PIBAR1) | R/W | All zeros | 13.3.2.12/13-23 |
| 0x5C | PCI inbound extended base address register 1 (PIEBAR1) | R/W | All zeros | 13.3.2.13/13-24 |
| 0x60 | PCI inbound window attributes register 1 (PIWAR1) | R/W | All zeros | 13.3.2.14/13-24 |
| 0x68 | PCI inbound translation address register 0 (PITAR0) | R/W | All zeros | 13.3.2.11/13-22 |
| 0x6C | Reserved | — | — | — |
| 0x70 | PCI inbound base address register 0 (PIBAR0) | R/W | All zeros | 13.3.2.12/13-23 |
| 0x78 | PCI inbound window attributes register 0 (PIWAR0) | R/W | All zeros | 13.3.2.13/13-24 |
| 0x7C–0xFF | Reserved | — | — | — |

13.3.1 PCI Configuration Access Registers

This section describes the registers used to allow a local bus master to access the PCI configuration space, and generate special cycle or interrupt acknowledge transactions on the PCI bus. A special case provides access to the PCI controller’s internal PCI configuration registers. The PCI registers, `PCI_CONFIG_ADDRESS`, `PCI_CONFIG_DATA`, and `PCI_INT_ACK`, are little endian registers.

13.3.1.1 PCI_CONFIG_ADDRESS

Figure 13-3 shows the PCI_CONFIG_ADDRESS register fields.




Figure 13-3. PCI_CONFIG_ADDRESS Register

The PCI_CONFIG_ADDRESS register holds the address for an access to the PCI configuration space from the local bus. This register must be programmed before accessing PCI_CONFIG_DATA to perform the transaction. Only 32-bit accesses are permitted.

If EN=1, BN=0, and DN=0, the access is to the internal PCI configuration registers, so no transaction is generated on the PCI bus.

If EN=1, BN=0, DN=31, FN=7, and RN=0, writing to PCI_CONFIG_DATA generates a special cycle transaction and reading from PCI_CONFIG_DATA generates an interrupt acknowledge transaction.

Table 13-6 shows the bit settings of the PCI_CONFIG_ADDRESS register.

Table 13-6. PCI_CONFIG_ADDRESS Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31 | EN | Enable configuration transaction. Determines the type of transaction to be generated. 0 No configuration transaction will be generated by accessing the CONFIG_DATA register. Such an access will be passed through to the PCI bus as an I/O transaction. Since this is generally not desirable, the user should not access CONFIG_DATA when the EN bit is 0. 1 A configuration transaction will be generated by accessing the CONFIG_DATA register if BN and DN are not both zero. |
| 30–24 | — | Reserved |
| 23–16 | BN | Bus number. Specifies the bus segment to which a configuration transaction is directed. If this field is 0, a Type 0 configuration transaction is generated. Otherwise, a Type 1 configuration transaction is generated. |

Table 13-6. PCI_CONFIG_ADDRESS Field Descriptions (continued)

| Bits | Name | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|--------|---------------------------------------|--|--------------------------------|--------------------------------|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|----|-------|---------------------------------------|-------|-----------------|--------|----------|
| 15–11 | DN | Device number. Specifies the device to which a configuration transaction is directed. For a Type 0 configuration transaction, this field is decoded to individual PCI1_IDSEL signals for the address phase according to the following values. For a Type 1 configuration transaction, this field is used directly for the address phase. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | <table border="0"> <thead> <tr> <th>Value</th> <th>AD Signal that is Driving High</th> </tr> </thead> <tbody> <tr><td>01010</td><td>31</td></tr> <tr><td>01011</td><td>11</td></tr> <tr><td>01100</td><td>12</td></tr> <tr><td>01101</td><td>13</td></tr> <tr><td>01110</td><td>14</td></tr> <tr><td>01111</td><td>15</td></tr> <tr><td>10000</td><td>16</td></tr> <tr><td>10001</td><td>17</td></tr> <tr><td>10010</td><td>18</td></tr> <tr><td>10011</td><td>19</td></tr> <tr><td>10100</td><td>20</td></tr> <tr><td>10101</td><td>21</td></tr> <tr><td>10110</td><td>22</td></tr> <tr><td>10111</td><td>23</td></tr> <tr><td>11000</td><td>24</td></tr> <tr><td>11001</td><td>25</td></tr> <tr><td>11010</td><td>26</td></tr> <tr><td>11011</td><td>27</td></tr> <tr><td>11100</td><td>28</td></tr> <tr><td>11101</td><td>29</td></tr> <tr><td>11110</td><td>30</td></tr> <tr><td>11111</td><td>Special cycle / interrupt acknowledge</td></tr> <tr><td>00000</td><td>Internal access</td></tr> <tr><td>Others</td><td>Reserved</td></tr> </tbody> </table> | Value | AD Signal that is Driving High | 01010 | 31 | 01011 | 11 | 01100 | 12 | 01101 | 13 | 01110 | 14 | 01111 | 15 | 10000 | 16 | 10001 | 17 | 10010 | 18 | 10011 | 19 | 10100 | 20 | 10101 | 21 | 10110 | 22 | 10111 | 23 | 11000 | 24 | 11001 | 25 | 11010 | 26 | 11011 | 27 | 11100 | 28 | 11101 | 29 | 11110 | 30 | 11111 | Special cycle / interrupt acknowledge | 00000 | Internal access | Others | Reserved |
| | | Value | AD Signal that is Driving High | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 01010 | 31 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 01011 | 11 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 01100 | 12 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 01101 | 13 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 01110 | 14 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 01111 | 15 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10000 | 16 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10001 | 17 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10010 | 18 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10011 | 19 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10100 | 20 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10101 | 21 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10110 | 22 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10111 | 23 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 11000 | 24 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 11001 | 25 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 11010 | 26 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 11011 | 27 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 11100 | 28 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 11101 | 29 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 11110 | 30 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 11111 | Special cycle / interrupt acknowledge | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 00000 | Internal access | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Others | Reserved | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 10–8 | FN | Function number. Specifies the function to which the configuration transaction is directed on a multi-function device. It is used directly in the address phase of the configuration transaction. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 7–2 | RN | Register number. Specifies the register being accessed in the PCI configuration space. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 1–0 | — | Reserved | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

13.3.1.2 PCI_CONFIG_DATA

An access to PCI_CONFIG_DATA usually generates a PCI configuration transaction if PCI_CONFIG_ADDRESS[EN] is set. There are some exceptions contained in the description of PCI_CONFIG_ADDRESS[EN].

This register may be accessed with an 8-, 16-, or 32-bit access, depending on the width of the register targeted by the configuration transaction.

Figure 13-4 shows the PCI_CONFIG_DATA register fields.




Figure 13-4. PCI_CONFIG_DATA

Table 13-7 shows the bit settings of the PCI_CONFIG_DATA register.

Table 13-7. PCI_CONFIG_DATA Field Descriptions

| Bits | Name | Description |
|------|----------|--|
| 31-0 | CFG_DATA | Configuration data. This field contains the data transferred on a PCI configuration transaction. |

13.3.1.3 PCI Interrupt Acknowledge Register (PCI_INT_ACK)

Reading this register generates an interrupt acknowledge transaction on the PCI bus. The value that is read is undefined.

13.3.2 PCI Memory-Mapped Control and Status Registers

This section describes the control and status registers.

13.3.2.1 PCI Error Status Register (PCI_ESR)

The PCI error status register (PCI_ESR) contains status bits for various types of error conditions captured by the PCI controller. Each status bit is set when the corresponding error condition is captured. PCI_ESR is a write-1-to-clear type register. A bit is cleared whenever the register is written and the data in the corresponding bit location is a 1. Figure 13-5 shows the PCI_ESR fields.




Figure 13-5. PCI Error Status Register (PCI_ESR)

Table 13-10. PCI_EER Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 26 | TABT | Target abort. Generate an interrupt when the corresponding bit of the PCI_ESR is 1. |
| 27–31 | — | Reserved |

13.3.2.4 PCI Error Attributes Capture Register (PCI_EATCR)

PCI_EATCR contains fields for storing information associated with the first PCI error captured. Figure 13-8 shows the PCI_EATCR fields.




Figure 13-8. PCI Error Attributes Capture Register (PCI_EATCR)

Table 13-11 describes the bit settings of the PCI_EATCR register.

Table 13-11. PCI_EATCR Field Descriptions


| Bits | Name | Description |
|------|---------|--|
| 0 | — | Reserved |
| 1–3 | ERRTYPE | First error type. This field is encoded to indicate the type of the first PCI error captured. 000 Address parity error 001 Write data parity error 010 Read data parity error 011 Master abort 100 Target abort 101 System error indication received 110 Parity error indication received on a read 111 Parity error indication received on a write |
| 4–7 | BN | Beat number. This field provides the data beat number on which the error occurred for data parity errors. The value of this field is undefined for other error types. The beat values are described as follows: 0000 1st beat 0001 2nd beat 0010 3rd beat 0011 4th beat 0100 5th beat 0101 6th beat 0110 7th beat 0111 8th beat 1000 9th beat or beyond (transaction larger than one cache line) Others Reserved |

Table 13-11. PCI_EATCR Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 8–9 | — | Reserved |
| 10–11 | TS | Transaction size. Indicates the size of the transaction in units of doublewords (8 bytes). If the transaction crossed a cache line (32-byte) boundary, this field indicates the number of actual double words in the cache line on which the error occurred. This field is valid only if the PCI controller was the master of the transaction. 00 4 double words 01 1 double word 10 2 double words 11 3 double words |
| 12–15 | ES | Error source. This field indicates the source of the PCI transaction. 0000 External master 0101 DMA Others reserved |
| 16–19 | CMD | PCI command. Contains the PCI command PCI_CBE[3:0] of the transaction. |
| 20–23 | — | Reserved |
| 24–27 | BE | PCI byte enables. Contains the PCI byte enables PCI_CBE[3:0] for the data word. |
| 28–29 | — | Reserved |
| 30 | PB | Parity bit. Contains the PCI parity bit for the captured data word. |
| 31 | VI | Error information valid. This bit indicates that the error information captured in this register, PCI_EACR, PCI_EEACR, and PCI_EDCR is valid. 0 No valid error information 1 Error information is valid |

13.3.2.5 PCI Error Address Capture Register (PCI_EACR)

PCI_EACR contains fields for storing the low portion of the address associated with the first PCI error captured. [Figure 13-9](#) shows the PCI_EACR fields.


Figure 13-9. PCI Error Address Capture Register (PCI_EACR)

[Table 13-12](#) describes the bit settings of the PCI_EACR register.

Table 13-12. PCI_EACR Field Description

| Bits | Name | Description |
|------|--------|---|
| 0–31 | PCI_EA | PCI error address. Contains the low portion of the address associated with the first detected error. Read only. |

13.3.2.6 PCI Error Extended Address Capture Register (PCI_EEACR)

PCI_EEACR contains fields for storing the high portion of the address associated with the first PCI error captured. [Figure 13-10](#) shows the PCI_EEACR fields.




Figure 13-10. PCI Error Extended Address Capture Register (PCI_EEACR)

[Table 13-13](#) describes the bit settings of the PCI_EEACR register.

Table 13-13. PCI_EEACR Field Description

| Bits | Name | Description |
|------|---------|--|
| 0–31 | PCI_EEA | PCI error extended address. Contains the high portion of the address associated with the first detected error. |

13.3.2.7 PCI Error Data Low Capture Register (PCI_EDLCR)

PCI_EDLCR contains fields for storing the data associated with the first PCI error captured. [Figure 13-11](#) shows the PCI_EDLCR fields.



Figure 13-11. PCI Error Data Low Capture Register (PCI_EDLCR)

[Table 13-14](#) describes the bit settings of the PCI_EDLCR register.

Table 13-14. PCI_EDLCR Field Description

| Bits | Name | Description |
|------|---------|---|
| 0–31 | PCI_EDR | PCI error data. Contains the data associated with the first detected error. |

13.3.2.8 PCI General Control Register (PCI_GCR)

PCI_GCR contains fields for controlling the behavior of the internal arbiter, the state of the bus signals, and the PCI reset signal for host mode. [Figure 13-12](#) shows the PCI_GCR fields.

outbound window, or where an outbound translation window points back into an inbound window, are not allowed.




Figure 13-15. PCI Inbound Translation Address Registers (PITAR n)

Table 13-18 shows the bit settings of PITAR n .

Table 13-18. PITAR n Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–11 | — | Reserved |
| 12–31 | TA | Translation address. Contains the starting address of the inbound translated address. TA corresponds to the 20 highest-order bits of a 32-bit local address. The specified address must be aligned to the window size, as defined by PIWAR n [IWS]. |

13.3.2.12 PCI Inbound Base Address Registers (PIBAR n)

PIBAR n contains fields for defining the starting point of the inbound windows in the PCI memory space. A write to a PIBAR n register also causes a change in the base address bits in the corresponding GPL base address register in the PCI configuration space. Figure 13-16 shows the PIBAR x fields.



Figure 13-16. PCI Inbound Base Address Registers (PIBAR n)

Table 13-19 shows the bit settings of PIBAR n .

Table 13-19. PIBAR n Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–31 | BA | Base address. Contains the starting address in the PCI memory space of the inbound window. This field corresponds to bits 43–12 of a 64-bit address. In PIBAR0, the upper 12 bits are reserved because only a 32-bit address is supported. The specified address must be aligned to the window size, as defined by PIWAR n [IWS]. |

13.3.2.13 PCI Inbound Extended Base Address Registers (PIEBAR_n)

PIEBAR_n contains fields for defining the high portion of the starting point of the inbound windows in the PCI memory space. Figure 13-17 shows the PIEBAR_n fields.




Figure 13-17. PCI Inbound Extended Base Address Registers (PIEBAR_n)

Table 13-20 shows the bit settings of PIEBAR_n.

Table 13-20. PIEBAR_n Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–11 | — | Reserved |
| 12–31 | EBA | Extended base address. Contains the high portion of the starting address in the PCI memory space of the inbound base address. This 20-bit field corresponds to bits 63–44 of a 64-bit address. |

13.3.2.14 PCI Inbound Window Attribute Registers (PIWAR_n)

PIWAR_n contains fields for defining the size of an inbound translation window. It also defines some properties of the window. Figure 13-18 shows the PIWAR_n fields. See Section 4.3.1.1, “Reset Configuration Word Source,” for more information on reset configuration.




Figure 13-18. PCI Inbound Window Attribute Registers (PIWAR_n)

Table 13-21 shows the bit settings of PIWAR_n.

Table 13-21. PIWAR_n Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0 | EN | Enable. Used to enable the address translation window. 0 Address translation is disabled for this window. 1 Address translation is enabled for this window. PCI addresses that match the definition of the window will be recognized by the PCI controller and translated to the local memory space. |
| 1 | — | Reserved |

Table 13-21. PIWAR_n Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 2 | PF | Prefetchable. Defines whether the transactions that are translated through this window are prefetchable on the local bus. Streaming the transactions requires the memory space to be prefetchable. 0 Not prefetchable 1 Prefetchable |
| 3–11 | — | Reserved |
| 12–15 | RTT | Read transaction type. Determines the type of transaction performed on the local bus when the PCI transaction is a read. The RTT values are described as follows: 0100 Read without snoop on system bus 0101 Read with snoop on system bus Others reserved |
| 16–19 | WTT | Write transaction type. Determines the type of transaction performed on the local bus when the PCI transaction is a write. The WTT values are described as follows: 0100 Write without snoop of local processor 0101 Write with snoop of local processor Others reserved |
| 20–25 | — | Reserved |
| 26–31 | IWS | Inbound window size. Indicates the size of the inbound translation window. Inbound translation window size N which is the encoded $2^{(N+1)}$ bytes window size. The smallest window is 4 Kbytes (N = 11) 000000–001010 Reserved 001011 4-Kbyte window size 001100 8-Kbyte window size ... 011110 2-Gbyte window size 011111–111111 Reserved |

13.3.3 PCI Configuration Space Registers

This section describes the PCI configuration space registers. These registers are shown with descending bit numbering to correspond to the PCI standard.

NOTE

The registers described in this section use little-endian byte ordering. Software running on the local processor in big-endian mode must byte-swap the data. No byte swapping occurs when the registers are accessed from the PCI bus.

Table 13-22 shows the PCI configuration registers that are mapped in PCI configuration space. Some fields are common to registers in both spaces to ensure consistency. These fields are discussed in the register definitions.

Table 13-22. PCI Configuration Space Registers

| Address | Use | Access |
|---------|---|----------------|
| 00 | Vendor ID configuration register | R |
| 02 | Device ID configuration register | R |
| 04 | PCI command configuration register | R/W |
| 06 | PCI status configuration register | Read/bit-reset |
| 08 | Revision ID configuration register | R |
| 09 | Standard programming interface | R |
| 0A | Subclass code configuration register | R |
| 0B | Base class code configuration register | R |
| 0C | Cache line size configuration register | R/W |
| 0D | Latency timer configuration register | R/W |
| 0E | Header type configuration register | R |
| 0F | BIST control configuration register | R |
| 10 | PIMMR base address register | R/W |
| 14 | GPL base address register 0 | R/W |
| 18 | GPL base address register 1 | R/W |
| 1C | GPL extended base address register 1 | R/W |
| 20 | GPL base address register 2 | R/W |
| 24 | GPL extended base address register 2 | R/W |
| 2C | Subsystem vendor ID configuration register | R |
| 2E | Subsystem device ID configuration register | R |
| 34 | Capabilities pointer configuration register | R |
| 3C | Interrupt line configuration register | R/W |
| 3D | Interrupt pin configuration register | R |
| 3E | Minimum grant configuration register | R |
| 3F | Maximum latency configuration register | R |
| 44 | PCI function configuration register | R/W |
| 46 | PCI arbiter control register (PCIACR) | R/W |
| 48 | Hot swap register block | R/W |

13.3.3.1 Vendor ID Configuration Register

Figure 13-19 shows the vendor ID fields. This is a read-only register.




Figure 13-19. Vendor ID Configuration Register

Table 13-23 shows the bit settings of the vendor ID register.

Table 13-23. Vendor ID Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 15–0 | VID | Vendor ID. The read-only value 0x1957 specifies Freescale Semiconductor as the manufacturer of the device. |

13.3.3.2 Device ID Configuration Register

Figure 13-20 shows the device ID fields. This is a read only register.




Figure 13-20. Device ID Configuration Register

Table 13-24 shows the bit settings of the device ID register.

Table 13-24. Device ID Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 15–0 | DID | Device ID. This field identifies the device. 00B4 MPC8315E 00B5 MPC8315 00B6 MPC8314E 00B7 MPC8314 |

13.3.3.3 PCI Command Configuration Register

Figure 13-21 shows the PCI command fields.

Offset 04

Access: Mixed




Figure 13-21. PCI Command Configuration Register

Table 13-25 shows the bit settings of the PCI command register.

Table 13-25. PCI Command Configuration Register Field Descriptions

| Bits | Name | Description |
|-------|--------|---|
| 15–11 | — | Reserved |
| 10 | INTD | Interrupt Disable. Setting this bit masks the $\overline{\text{PCI_INTA}}$ output. 0 $\overline{\text{PCI_INTA}}$ provides the device interrupt status. 1 $\overline{\text{PCI_INTA}}$ is always negated. |
| 9 | FB-B | Fast back-to-back. Hard-wired to 0. |
| 8 | SERREN | SERR enable. This bit is an enable bit for the SERR driver. Address parity errors are reported only if this bit and bit 6 are 1. 0 $\overline{\text{PCI_SERR}}$ is never asserted. 1 $\overline{\text{PCI_SERR}}$ may be asserted to indicate error conditions. |
| 7 | — | Reserved |
| 6 | PERRR | Parity error response. Controls the PCI controller's response to a parity error. 0 Parity errors are ignored and normal operation continues. 1 Standard parity error treatment. |
| 5 | — | Reserved |
| 4 | MWI | Memory-write-and-invalidate. Hard-wired to 0. |
| 3 | SC | Special cycles. Hard-wired to 0. |
| 2 | BMST | Bus master. Controls the PCI controller's ability to be a master on the PCI bus. At reset, this bit is cleared in Agent Mode and set in Host Mode. 0 The PCI controller does not generate PCI accesses. 1 The PCI controller behaves as a bus master. |
| 1 | MEM | Memory space. Controls the response to memory space accesses. 0 The PCI controller does not respond to Memory Space accesses. 1 The PCI controller as a target responds to Memory Space accesses. |
| 0 | I/O | I/O space. Hard-wired to 0. |

13.3.3.4 PCI Status Configuration Register

This register is used to record status information for PCI bus-related events. Some of the bits are hard-wired to indicate the capabilities of the PCI controller. Other bits can be cleared by writing 1 to the bit location. Figure 13-22 shows the PCI status fields.

| | | | | | | | | | | | | | | | |
|-----------|-------|-------|-----|-----|-----|----------|---|-----|-------|---|---------------|----|------|---|---|
| Offset 06 | | | | | | | | | | | Access: Mixed | | | | |
| | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 0 |
| R | DPERR | SSERR | RMA | RTA | STA | DEVSEL_T | | DPD | FB-BC | — | 66M | CL | INTS | — | |
| W | w1c | w1c | w1c | w1c | w1c | | | w1c | | | | | w1c | | |
| Reset | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 0 | 0 | 0 |

Figure 13-22. PCI Status Configuration Register

Table 13-26 shows the bit settings of the PCI status register.

Table 13-26. PCI Status Configuration Register Field Descriptions

| Bits | Name | Description |
|------|----------|---|
| 15 | DPERR | Detected parity error. Set whenever the PCI controller detects a parity error on the PCI bus, even if parity error handling is disabled (as controlled by bit 6 in the PCI Command register). |
| 14 | SSERR | Signaled system error. Set whenever $\overline{\text{PCI_SERR}}$ is asserted. |
| 13 | RMA | Received master abort. Set whenever the PCI controller, acting as the PCI master on the PCI bus, terminates a transaction (except for a special-cycle) using master-abort. |
| 12 | RTA | Received target abort. Set whenever a transaction initiated by this PCI controller on the PCI bus is terminated by a target-abort. |
| 11 | STA | Signaled target abort. Set whenever the PCI controller, acting as the PCI target on the PCI bus, issues a target-abort to a PCI master. |
| 10–9 | DEVSEL_T | DEVSEL timing. Hard-wired to 00. |
| 8 | DPD | Master data parity error. Set when a data parity error is detected on the PCI bus, if the PCI controller is the master that initiated the transaction and bit 6 in the PCI command register is set. |
| 7 | FB-BC | Fast back-to-back capable. Hard-wired to 1. |
| 6 | — | Reserved |
| 5 | 66M | 66-MHz capable. Hard-wired to 1. |
| 4 | CL | Capabilities list. Hard-wired to 1. |
| 3 | INTS | Interrupt status. Contains the status of the device interrupt. The value of this bit is not affected by the INTD bit of the PCI command configuration register. |
| 2–0 | — | Reserved |

13.3.3.5 Revision ID Configuration Register

Figure 13-23 shows the revision ID fields.




Figure 13-23. Revision ID Configuration Register

Table 13-27 shows the bit settings of the revision ID register.

Table 13-27. Revision ID Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 7-0 | RID | Revision ID. Specifies a revision code of the PCI controller. 8'h10 Revision 010.8'h21 Revision 021. |

13.3.3.6 Standard Programming Interface Configuration Register

Figure 13-24 shows the standard programming interface fields. This is the lower byte of the class code.




Figure 13-24. Standard Programming Interface Configuration Register

Table 13-28 shows the bit settings of the standard programming interface register.

Table 13-28. Standard Programming Interface Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7-0 | PI | Programming interface. This field is hard-wired to 0x00. |

13.3.3.7 Subclass Code Configuration Register

Figure 13-25 shows the subclass code fields. This is the middle byte of the class code.




Figure 13-25. Subclass Code Configuration Register

Table 13-29 shows the bit settings of the subclass code register.

Table 13-29. Subclass Code Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7–0 | SC | Sub-class code. This field is hard-wired to 0x20, indicating a Power PC processor. |

13.3.3.8 Base Class Code Configuration Register

Figure 13-26 shows the base class code fields. This is the upper byte of the class code.




Figure 13-26. Base Class Code Configuration Register

Table 13-30 shows the bit settings of the class code register.

Table 13-30. Class Code Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7–0 | BCC | Base class code. This field is hard-wired to 0x0B, indicating a processor. |

13.3.3.9 Cache Line Size Configuration Register

Figure 13-27 shows the cache line size fields.




Figure 13-27. Cache Line Size Configuration Register

Table 13-31 shows the bit settings of the cache line size register.

Table 13-31. Cache Line Size Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7-0 | CLS | Cache line size. Cache-line in terms of 32-bit words. Although the register is writable, only the value 0x08 is legal. |

13.3.3.10 Latency Timer Configuration Register

Figure 13-28 shows the latency timer fields.




Figure 13-28. Latency Timer Configuration Register

Table 13-32 shows the bit settings of the latency timer register.

Table 13-32. Latency Timer Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7-3 | LT | Latency timer. Specifies a granularity of 8 PCI clocks, the length of time that the PCI controller, when mastering a transaction, may hold the bus as the result of a bus grant. Refer to the PCI 2.3 specification for the rules by which the PCI controller completes transactions when the timer has expired. |
| 2-0 | — | Reserved |

13.3.3.11 Header Type Configuration Register

Figure 13-29 shows the read-only header type register, which is hard-wired to 0x00.




Figure 13-29. Header Type Configuration Register

13.3.3.12 BIST Control Configuration Register

Figure 13-30 shows the read-only BIST control register, which is hard-wired to 0x00.




Figure 13-30. BIST Control Configuration Register

13.3.3.13 PIMMR Base Address Configuration Register

Figure 13-31 shows the PIMMR base address register fields.




Figure 13-31. PIMMR Base Address Configuration Register

Table 13-33 shows the bit settings of the PIMMR base address register.

Table 13-33. PIMMR Base Address Configuration Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–20 | BA | Base address. Defines the base address for the internal (on-chip) memory-mapped register space. The size of this space is 1 Mbytes. |
| 19–4 | — | Reserved |
| 3 | PRE | Prefetchable. Hard-wired to 0. |

13.3.3.17 Subsystem Vendor ID Configuration Register

Figure 13-35 shows the subsystem vendor ID fields. The subsystem vendor ID configuration register is read-only from the PCI bus, but it can be programmed from the CSB.




Figure 13-35. Subsystem Vendor ID Configuration Register

Table 13-37 shows the bit settings of the subsystem vendor ID configuration register.

Table 13-37. Subsystem Vendor ID Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 15–0 | SVID | Subsystem vendor ID. Identifies the manufacturer of the board or subsystem that contains this device. |

13.3.3.18 Subsystem Device ID Configuration Register

Figure 13-36 shows the subsystem device configuration register ID fields. The subsystem device ID configuration register is read-only from the PCI bus, but it can be programmed from the CSB.




Figure 13-36. Subsystem Device ID Configuration Register

Table 13-38 shows the bit settings of the subsystem device ID configuration register.

Table 13-38. Subsystem Device ID Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 15–0 | SDID | Subsystem device ID. This field identifies the board or subsystem that contains this device. |

13.3.3.19 Capabilities Pointer Configuration Register

The capabilities pointer register specifies the byte offset in the PCI configuration space that contains the first item in the capabilities list. Figure 13-37 shows the capabilities pointer configuration register fields.




Figure 13-37. Capabilities Pointer Configuration Register

13.3.3.20 Interrupt Line Configuration Register

Figure 13-38 shows the interrupt line configuration register fields.




Figure 13-38. Interrupt Line Configuration Register

Table 13-39 shows the bit settings of the interrupt line configuration register.

Table 13-39. Interrupt Line Configuration Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 7-0 | IL | Interrupt line. Used to communicate interrupt line routing information. The value has no effect on the operation of the PCI controller. |

13.3.3.21 Interrupt Pin Configuration Register

The interrupt pin configuration register tells which interrupt pin is used (0x01 means PCI_INTA).

Figure 13-39 shows the interrupt pin configuration register fields.




Figure 13-39. Interrupt Pin Register

13.3.3.22 Minimum Grant Configuration Register

Figure 13-40 shows the minimum grant configuration register fields.

PCI Bus Interface

Offset 0x3E




Figure 13-40. Minimum Grant Configuration Register

13.3.3.23 Maximum Latency Configuration Register

Figure 13-41 shows the maximum latency configuration register fields.

Offset 0x3F




Figure 13-41. Maximum Latency Configuration Register

13.3.3.24 PCI Function Configuration Register

Figure 13-42 shows the PCI function configuration register fields.

Offset 0x44

Access: Read/Write




Figure 13-42. PCI Function Configuration Register

Table 13-40 shows the bit settings of the PCI function configuration register.

Table 13-40. PCI Function Configuration Register Field Descriptions

| Bits | Name | Description |
|------|----------|--|
| 15–6 | — | Reserved |
| 5 | CFG_LOCK | Configuration lock. Controls access to the PCI configuration space from the PCI port. In host mode the PCI configuration space is always inaccessible, so this bit is not used. Normally, this bit will be cleared in agent mode once the configuration of the PCI controller is complete to allow an external host to access the PCI configuration space. 0 Access to the configuration spaces is permitted. 1 Any inbound PCI access to the PCI configuration space is retried. See Section 4.3.1.1, “Reset Configuration Word Source,” for more information on reset configuration. |
| 3–4 | — | Reserved |

Table 13-40. PCI Function Configuration Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 2 | TLTD | Target latency timeout disable. Determines whether the PCI controller, while acting as a PCI target, times out when the first data phase of a transaction has not completed in 16 PCI cycles. 0 Target latency timeout enabled. 1 Target latency timeout disabled. |
| 1 | MLTD | Master latency timer disable. Determines whether the PCI controller, while acting as a PCI master, terminates a transaction upon the expiration of the master latency timer. 0 Master latency timer enabled. 1 Master latency timer disabled. |
| 0 | HA | Host/Agent. Indicates whether the PCI controller is in host mode or agent mode. It provides the value of the PCI_HOST—PCI host configuration bit is sampled at the end of the reset sequence. 0 Host mode 1 Agent mode |

13.3.3.25 PCI Arbiter Control Register (PCIACR)

Figure 13-43 shows the PCI arbiter control register (PCIACR) fields.




Figure 13-43. PCI Arbiter Control Register (PCIACR)

Table 13-41 shows the bit settings of the PCIACR.

Table 13-41. PCI Arbiter Control Register (PCIACR) Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 15 | AD | Arbiter disable. Indicates whether the PCI controller functions as the arbiter for the PCI bus. It provides the value of the PCI arbiter enable configuration bit as sampled at the end of the reset sequence. See Chapter 4, “Reset, Clocking, and Initialization,” for more information on reset configuration. 0 Arbiter enabled 1 Arbiter disabled |
| 14 | PM | Parking mode. Controls which device receives a bus grant when there are no outstanding bus requests and the bus is idle. 0 The bus is parked with the last device to use the bus. 1 The bus is parked with the PCI controller. |
| 13 | — | Reserved |
| 12 | PBMD | PCI broken master disable. Determines whether the PCI controller ignores the bus requests of an initiator that requests the bus for an excessive period without using it. 0 An initiator that requests the bus and receives the grant must begin using the bus within 16 PCI clock periods after the bus becomes idle or its request is subsequently ignored. 1 No requests are ignored. |
| 11–7 | — | Reserved |

Table 13-41. PCI Arbiter Control Register (PCIACR) Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 6–4 | PRIn | Priority level for master <i>n</i> . When the PCI controller functions as the arbiter for the PCI bus, each PRIn bit determines the arbitration priority level for the PCI master connected to the REQn/GNTn pair. 0 Low priority 1 High priority |
| 3–1 | — | Reserved |
| 0 | MPRI | My priority. When the PCI controller functions as the arbiter for the PCI bus, this bit determines the arbitration priority level for the PCI controller when it acts as a PCI master. 0 Low priority 1 High priority |

13.3.3.26 Hot Swap Register Block

Figure 13-44 shows the hot swap register block fields.

Offset 0x48




Figure 13-44. Hot Swap Register Block

Table 13-42 shows the bit settings of the Hot Swap register block.

Table 13-42. Hot Swap Register Block Field Descriptions

| Bits | Name | Description |
|-------|---------|--|
| 31–24 | — | Reserved |
| 23 | INS | Insertion status. Indicates that a card has been inserted. Write 1 to clear this bit. |
| 22 | EXT | Extraction status. Indicates that a card has been extracted. Write 1 to clear this bit. |
| 21–20 | — | Reserved |
| 19 | LOO | LED On/Off. Controls the LED when the hardware is in state H2 0 LED off 1 LED on |
| 18 | — | Reserved |
| 17 | EIM | ENUM mask. This bit masks the CPCI_HS_ENUM input. 0 Enabled 1 Masked |
| 16 | — | Reserved |
| 15–8 | NXT_PTR | Next pointer—hardwired to 0x80 to point to the address of the power management capability in the PCI controller. |
| 7–0 | CAP_ID | Capability ID for hot swap (hardwired to 0x06) |

Table 13-43. PCIPMR0 Field Descriptions (continued)

| Bits | Name | Description |
|-------|--------------|--|
| 19 | PME_Clock | PME clock 0 Indicates that no PCI clock is required for the PCI controller to generate PME# 1 Indicates that PCI clock is required for the PCI controller to generate PME# |
| 18–16 | Version | PCI Power Management Interface Specification version. 011 Revision 1.2 of the PCI Power Management Interface Specification |
| 15–8 | NEXT_CAP_PTR | The next capability pointer points to the next item in the PCI controller's capability list. 0000_0000 The end of the capability list |
| 7–0 | CAP_ID | 0000_0001 Indicates the power management support capability |

13.3.3.28 PCI Power Management Register 1 (PCIPMR1)

The PCI power management register 1 (PCIPMR1), shown in [Figure 13-46](#), contains the bit fields that software uses to manage the PCI controller's power management state, as well as to enable and monitor PMEs (power management events). This register can be accessed by the host in agent mode.

Offset 0x84

Access: Read / write




Figure 13-46. PCI Power Management Register 1 (PCIPMR1)

[Table 13-44](#) describes the PCIPMR1 fields.

Table 13-44. PCIPMR1 Field Descriptions

| Bits | Name | Description |
|-------|---------|---|
| 31–24 | Data | Reports the state dependent data requested by the Data_Select field. The value of this field is scaled by the value reported by the Data_Scale field. |
| 23 | BPCC_En | Bus power/Clock control enable 0 Disable the bus power/clock control policies defined in section 4.7.1 of the PCI Bus Power Management Interface Specification Revision 1.2 1 Enable the bus power/clock control policies defined in section 4.7.1 of the PCI Bus Power Management Interface Specification Revision 1.2 Note: This bit field is not implemented, only required for all PCI-to-PCI Bridge |

Table 13-44. PCIPMR1 Field Descriptions (continued)

| Bits | Name | Description |
|-------|---------------|---|
| 22 | B2_B3 | The state of this bit determines what will happen as a direct result of programming the function to D3_hot. 0 Indicates that when the bridge function is programmed to D3_hot, its secondary bus will have its power removed (B3). 1 Indicates that when the bridge function is programmed to D3_hot, its secondary bus's PCI clock will be stopped (B2). This bit only meaningful if bit 16 (BPCC_En) is set. Note: This bit field is not implemented, only required for all PCI-to-PCI Bridge |
| 21–16 | — | Reserved |
| 15 | PME_Status | This bit set when the PCI controller would normally assert PME# signal independent of the state of the PME_En bit. Writing a value of one to this bit will clear it and cause the PCI controller to stop asserting a PME# signal. 0 Default 1 If (Wake_Up & PME_En) |
| 14–13 | Data_Scale | The scale factor to be used when interpreting the value of the data register |
| 12–9 | Data_Select | Selects which data is to be reported through the data register and Data_Scale field |
| 8 | PME_En | Enables the function to assert PME# 0 Disables the function to assert PME# 1 Enables the function to assert PME# |
| 7–4 | — | Reserved |
| 3 | No_Soft_Reset | This bit field indicates whether an internal reset occurs during the transition from D3_hot to D0. 0 The Power_State command performs an internal reset. 1 The Power_State command does not perform an internal reset. |
| 2 | — | Reserved |
| 1–0 | Power_State | Determines the current power state of the PCI controller and sets the controller into a new power state. The power state definition is as follows: 00 D0 supports all PCI function. 01 D1 disables the inbound memory space, bus mastering and functional interrupt request. 10 D2 disables the inbound memory space, bus mastering and functional interrupt request. 11 D3_hot disables the inbound memory space, bus mastering and functional interrupt request. |

13.4 Functional Description

The following sections discuss the operation of the PCI controller.

13.4.1 PCI Bus Arbitration

The PCI bus arbitration approach is access-based. Bus masters must arbitrate for each access performed on the bus. PCI uses a central arbitration scheme where each master has its own unique request (\overline{REQn}) output and grant ($GNTn$) input signal. A simple request-grant handshake is used to gain access to the bus. Arbitration for the bus occurs during the previous access so that no PCI bus cycles are consumed waiting for arbitration (except when the bus is idle).

The PCI internal arbiter supports three external masters (besides the PCI controller itself) by using the $\overline{\text{REQ}}$ signals and generating the $\overline{\text{GNT}}$ signals.

During reset, the PCI controller samples the reset configuration bit (and programs the `PCI_ARB_DIS` bit accordingly) to determine if the arbiter is enabled or disabled. See [Chapter 4, “Reset, Clocking, and Initialization,”](#) for more information. The arbiter can also be enabled or disabled by directly programming the `PCI_ARB_DIS` bit in the arbiter configuration register (see [Section 13.3.3.25, “PCI Arbiter Control Register \(PCIACR\),”](#) for more information). However, it is recommended to use the reset configuration bit to set the arbiter state because the arbiter state controls the direction of $\overline{\text{REQ0}}$ and $\overline{\text{GNT0}}$.

If the arbiter is disabled, the PCI controller uses $\overline{\text{REQ0}}$ to issue requests to an external arbiter, and uses $\overline{\text{GNT0}}$ to receive grants from the external arbiter.

13.4.1.1 Bus Parking

When no devices are requesting the bus, the bus is granted, or parked, for a specified device to prevent the `AD`, `PCI_C/ $\overline{\text{BE}}$` and `PCI_PAR` signals from floating. The PCI controller can be configured to either park on itself or park on the last master to use the bus (see [Section 13.3.3.25, “PCI Arbiter Control Register \(PCIACR\),”](#) for more information).

13.4.1.2 Arbitration Algorithm

The round-robin arbitration algorithm has two priority levels. Each of the external PCI bus masters, plus the PCI controller, are assigned either a high or a low priority level, as programmed in the arbiter configuration register (see [Section 13.3.3.25, “PCI Arbiter Control Register \(PCIACR\).”](#)) Within each priority group (high or low), the bus grant is given to the next requesting device in numerical order, with the PCI controller itself positioned before device 0. $\overline{\text{GNT}}_n$ is asserted for device n as soon as the previously granted device begins a transaction. Conceptually, the lowest priority device at any given time is the current bus master and the highest priority device is the next one to follow the current master. This is considered to be a fair algorithm because a given device cannot prevent other devices from having access to the bus—a given device automatically becomes the lowest priority device as soon as it begins to use the bus. If a master is not requesting the bus, the transaction slot is given to the next requesting device within the priority group.

The grant given to one device may be taken away and whenever a higher priority device asserts its request. If the bus is idle when a new device is to receive a grant, no device receives a grant for one clock; in the next clock, the new winner of the arbitration receives a grant. This operation allows for a turnaround clock when a device is using address stepping or when the bus is parked.

The low priority group collectively receives one bus transaction request slot in the high priority group. Therefore, if there are N high-priority devices, each high-priority device is guaranteed to get at least one of $(N+1)$ bus transactions, and the M low priority devices are guaranteed to each get at least one of $(N+1) \times M$ bus transactions, with one of the low-priority devices receiving the grant in one of $(N+1)$ bus transactions. If all devices are programmed to the same priority level or if there is only one device at the low priority, the algorithm provides each device an equal number of bus grants in a round-robin sequence.

An arbitration example with three masters in the high priority group and two in the low priority group is shown in [Figure 13-47](#). Noting that one position in the high priority group is actually a place-holder for

the low priority group, it can be seen that each high priority initiator is guaranteed at least 1 out of 3 transaction slots, and each low priority initiator is guaranteed at least 1 out of 6 slots. Assuming all devices are requesting the bus, the grant sequence (with device 1 being the current master) is as follows: 0, 2, the PCI controller, 0, 2, 1, 0, 2, the PCI controller, and so on. If, for example, device 2 is not requesting the bus, the grant sequence becomes 0, the PCI controller, 0, 1, 0, the PCI controller, and so on. If device 2 now requests the bus at a point in the sequence when device 0 is conducting a transaction and the PCI controller is the next grant, then the PCI controller's grant is removed, and the higher-priority device 2 is awarded the next grant.

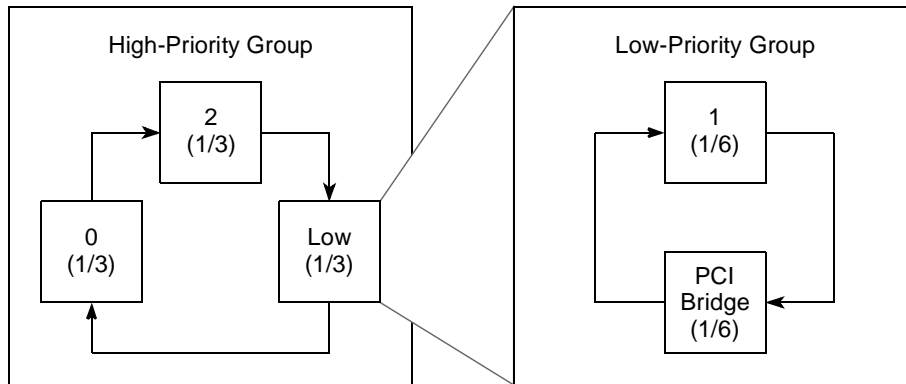


Figure 13-47. PCI Arbitration Example

13.4.1.3 Broken Master Lock-Out

The broken master feature allows the arbiter to lock out any masters that are broken or ill-behaved. This feature is controlled by programming the PCI arbiter control register. When the broken master feature is enabled, a granted device that does not assert $\overline{\text{PCI_FRAME}}$ within 16 PCI clock cycles after the bus is idle, has its grant removed and subsequent requests are ignored until its $\overline{\text{REQ}}$ is negated for at least one clock cycle. This prevents ill-behaved masters from monopolizing the bus. When the broken master feature is disabled, a device that requests the bus and receives a grant never loses its grant until and unless it begins a transaction or negates its $\overline{\text{REQ}}$ signal. Note that disabling the broken master feature is not recommended.

13.4.1.4 Master Latency Timer

The PCI controller implements the master latency timer register (see [Section 13.3.3.10, “Latency Timer Configuration Register”](#)) to prevent itself from monopolizing the bus. When the master latency timer expires, the PCI controller checks the state of its $\overline{\text{PCI_GNT}}$ signals. If the $\overline{\text{PCI_GNT}}$ signal is not asserted, the PCI controller completes one more data phase and relinquishes the bus. The master latency timer can be disabled if needed (see [Section 13.3.3.24, “PCI Function Configuration Register,”](#) for more information).

13.4.2 Bus Commands

PCI bus commands indicate the type of transaction occurring on the bus. These commands are encoded on PCI_C/BE[3:0] during the address phase of the transaction. PCI bus commands are described in [Table 13-45](#).

Table 13-45. PCI Command Definitions

| PCI_C/ BE[3:0] | Command Type | Supported as: | | Definition |
|-------------------|-----------------------------|---------------|--------|---|
| | | Initiator | Target | |
| 0b0000 | Interrupt acknowledge | Yes | No | A read implicitly addressed to the system interrupt controller. The size of the vector to be returned is indicated on the byte enables after the address phase. |
| 0b0001 | Special cycle | Yes | No | Provides a simple message broadcast mechanism. See Section 13.4.4.6, "Special Cycle Command," for more information. |
| 0b0010 | I/O read | Yes | No | Accesses agents mapped in I/O address space. |
| 0b0011 | I/O write | Yes | No | Accesses agents mapped in I/O address space. |
| 0b010x | — | — | — | Reserved. No response occurs. |
| 0b0110 | Memory read | Yes | Yes | Accesses agents mapped in memory address space. A read from prefetchable space, when seen as a target, fetches a cache line of data (32 bytes) from the starting address, even though all 32 bytes may not actually be sent to the initiator. |
| 0b0111 | Memory write | Yes | Yes | Accesses agents mapped in memory address space. Note that for inbound writes less than 4-bytes, the PCI controller splits the transaction into single byte writes to the target. Thus, the PCI interface cannot be used to perform single beat writes to 16-bit devices on the local peripheral interfaces. |
| 0b100x | — | — | — | Reserved. No response occurs. |
| 0b1010 | Configuration read | Yes | Yes | Accesses the configuration space of each agent. An agent is selected when its IDSEL signal is asserted. See Section 13.4.4.4, "Host Mode Configuration Access," for more information on configuration accesses. As a target, a configuration read is only accepted if the PCI controller is configured to be in agent mode. |
| 0b1011 | Configuration write | Yes | Yes | Accesses the configuration space of each agent. An agent is selected when its IDSEL signal is asserted. See Section 13.4.4.4, "Host Mode Configuration Access," for more information. As a target, a configuration write is only accepted if the PCI controller is configured to be in agent mode. |
| 0b1100 | Memory read multiple | Yes | Yes | Causes a prefetch of the next cache line. |
| 0b1101 | Dual address cycle | No | Yes | Transfers an 8-byte address to devices. |
| 0b1110 | Memory read line | Yes | Yes | Indicates that the initiator intends to transfer an entire cache line of data. |
| 0b1111 | Memory write and invalidate | No | Yes | Indicates that the initiator will transfer an entire cache line of data, and if PCI has any cacheable memory, this line needs to be invalidated. |

13.4.3 PCI Protocol Fundamentals

The bus transfer mechanism on the PCI bus is called a burst. A burst is comprised of an address phase and one or more data phases.

All signals are sampled on the rising edge of the PCI clock. Each signal has a setup and hold window with respect to the rising clock edge, in which transitions are not allowed. Outside this aperture, signal values or transitions have no significance.

13.4.3.1 Basic Transfer Control

PCI data transfers are controlled by the following signals:

- $\overline{\text{PCI_FRAME}}$ is driven by an initiator to indicate the beginning and end of a transaction.
- $\overline{\text{PCI_IRDY}}$ (initiator ready) is driven by an initiator, allowing it to force wait cycles.
- $\overline{\text{PCI_TRDY}}$ (target ready) is driven by a target, allowing it to force wait cycles.

The bus is idle when both $\overline{\text{PCI_FRAME}}$ and $\overline{\text{PCI_IRDY}}$ are negated. The first clock cycle in which $\overline{\text{PCI_FRAME}}$ is asserted indicates the beginning of the address phase. The address and the bus command code are transferred in that cycle. The next cycle ends the address phase and begins the data phase.

During the data phase, data is transferred in each cycle that both $\overline{\text{PCI_IRDY}}$ and $\overline{\text{PCI_TRDY}}$ are asserted. Once the PCI controller, as an initiator, has asserted $\overline{\text{PCI_IRDY}}$, it does not change $\overline{\text{PCI_IRDY}}$ or $\overline{\text{PCI_FRAME}}$ until the current data phase completes, regardless of the state of $\overline{\text{PCI_TRDY}}$. Once the PCI controller, as a target, has asserted $\overline{\text{PCI_TRDY}}$ or $\overline{\text{PCI_STOP}}$ it does not change $\overline{\text{PCI_DEVSEL}}$, $\overline{\text{PCI_TRDY}}$, or $\overline{\text{PCI_STOP}}$ until the current data phase completes.

When the PCI controller (as a master) intends to complete only one more data transfer, $\overline{\text{PCI_FRAME}}$ is negated and $\overline{\text{PCI_IRDY}}$ is asserted (or kept asserted) indicating the initiator is ready. After the target indicates it is ready ($\overline{\text{PCI_TRDY}}$ asserted) the bus returns to the idle state.

13.4.3.2 Addressing

The PCI specification defines three physical address spaces—memory, I/O, and configuration. The memory and I/O address spaces are standard for all systems. The configuration address space supports the PCI hardware configuration. Each PCI device decodes the address for each PCI transaction with each agent responsible for its own address decode.

The information contained in the two lower address bits (AD1 and AD0) depends on the address space. In the I/O address space, all 32 address/data lines provide the full byte address. AD[1:0] are used for the generation of $\overline{\text{PCI_DEVSEL}}$ and indicate the least significant valid byte involved in the transfer. Once a target has claimed an I/O access, it first determines if it can complete the entire access as indicated by the byte enable signals. If all the selected bytes are not in the address range, the entire access should not be completed; that is, the target should not transfer any data and should terminate the transaction with a target-abort operation. See [Section 13.4.3.6, “Bus Transactions,”](#) for more information.

In the configuration address space, accesses are decoded to a 4-byte address using AD[7:2]. An agent determines if it is the target of the access when a configuration command is decoded, IDSEL is asserted, and AD[1:0] are 0b00; otherwise, the agent ignores the current transaction. The PCI controller determines

a configuration access is for a device on the PCI bus by decoding a configuration command. When in agent mode, the PCI controller responds to host-generated PCI configuration cycles when its IDSEL is asserted during a configuration cycle.

For memory accesses, the address is decoded using AD[31:2]; thereafter, the address is incremented internally by 4 bytes until the end of the burst transfer. Another initiator in a memory access should drive 0b00 on AD[1:0] during the address phase to indicate a linear incrementing burst order. The PCI controller checks AD[1:0] during a memory command access and provides the linear incrementing burst order. On reads, if AD[1:0] is 0b10, which represents a cache line wrap, the PCI controller linearly increments the burst order starting at the critical 64-bit address, wraps at the end of the cache line, and disconnects after reading one cache line. If AD[1:0] is 0bx1 (a reserved encoding) and the PCI_C/BE[3:0] signals indicate a memory transaction, it executes a target disconnect after the first data phase is completed. Note that AD[1:0] are included in parity calculations.

13.4.3.3 Device Selection

As a target, the PCI controller drives $\overline{\text{PCI_DEVSEL}}$ one clock following the address phase as indicated in the configuration space status register; see [Section 13.3.3.4, “PCI Status Configuration Register,”](#) for more information. The PCI controller as a target qualifies the address/data lines with $\overline{\text{PCI_FRAME}}$ before asserting $\overline{\text{PCI_DEVSEL}}$. The $\overline{\text{PCI_DEVSEL}}$ signal is asserted at or before the clock edge at which the PCI controller enables its $\overline{\text{PCI_TRDY}}$, $\overline{\text{PCI_STOP}}$, or data (for a read). The $\overline{\text{PCI_DEVSEL}}$ signal is not negated until $\overline{\text{PCI_FRAME}}$ is negated, with $\overline{\text{PCI_IRDY}}$ asserted and either $\overline{\text{PCI_STOP}}$ or $\overline{\text{PCI_TRDY}}$ asserted. The exception to this is a target-abort; see [Section 13.4.3.8, “Transaction Termination,”](#) for more information.

As an initiator, if the PCI controller does not see the assertion of $\overline{\text{PCI_DEVSEL}}$ within 4 clocks of $\overline{\text{PCI_FRAME}}$, it terminates the transaction with a master-abort as described in [Section 13.4.3.8, “Transaction Termination,”](#) for more information.

13.4.3.4 Byte Enable Signals

The byte enable signals ($\overline{\text{BE}}[3:0]$) indicate which byte lanes carry valid data. The byte enable signals may enable different bytes for each of the data phases. The byte enable signals are valid on the edge of the clock that starts each data phase and remain valid for the entire data phase.

If the PCI controller, as a target, sees no byte enable signals asserted, it completes the current data phase with no permanent change. This implies that on a read transaction, the PCI controller expects the data not to be changed, and on a write transaction, the data is not stored.

13.4.3.5 Bus Driving and Turnaround


The turnaround-cycle is one clock cycle and is required to avoid contention. This cycle occurs at different times for different signals. $\overline{\text{PCI_IRDY}}$, $\overline{\text{PCI_TRDY}}$, and $\overline{\text{PCI_DEVSEL}}$ use the address phase as their turnaround-cycle. $\overline{\text{PCI_FRAME}}$, $\overline{\text{PCI_C/BE}}[3:0]$, and AD[31:0] use the idle cycle between transactions as their turnaround-cycle. (An idle cycle in PCI is when both $\overline{\text{PCI_FRAME}}$ and $\overline{\text{PCI_IRDY}}$ are negated).

Byte lanes not involved in the current data transfer are driven to a stable condition even though the data is not valid.

13.4.3.6 Bus Transactions

The timing diagrams in this section show the relationship of significant signals involved in bus transactions.

Note the following conventions:

- When a signal is drawn as a solid line, it is actively being driven by the current initiator or target.
- When a signal is drawn as a dashed line, no agent is actively driving it.
- Three-stated signals with slashes between the two rails have indeterminate values.
- The terms ‘edge’ and ‘clock edge’ refer to the rising edge of the clock.
- The terms ‘asserted’ and ‘negated’ refer to the globally visible state of the signal on the clock edge, and not to signal transitions.
- The symbol  represents a turnaround-cycle.

13.4.3.7 Read and Write Transactions

Both read and write transactions begin with an address phase followed by a data phase. The address phase occurs when $\overline{\text{PCI_FRAME}}$ is asserted for the first time, and the AD[31:0] signals contain a byte address and the PCI_C/ $\overline{\text{BE}}$ [3:0] signals contain a bus command. The data phase consists of the actual data transfer and possible wait cycles; the byte enable signals remain actively driven from the first clock of the data phase through the end of the data transfer.

A read transaction starts when $\overline{\text{PCI_FRAME}}$ is asserted for the first time and the PCI_C/ $\overline{\text{BE}}$ [3:0] signals indicate a read command. [Figure 13-48](#) shows an example of a single beat read transaction.

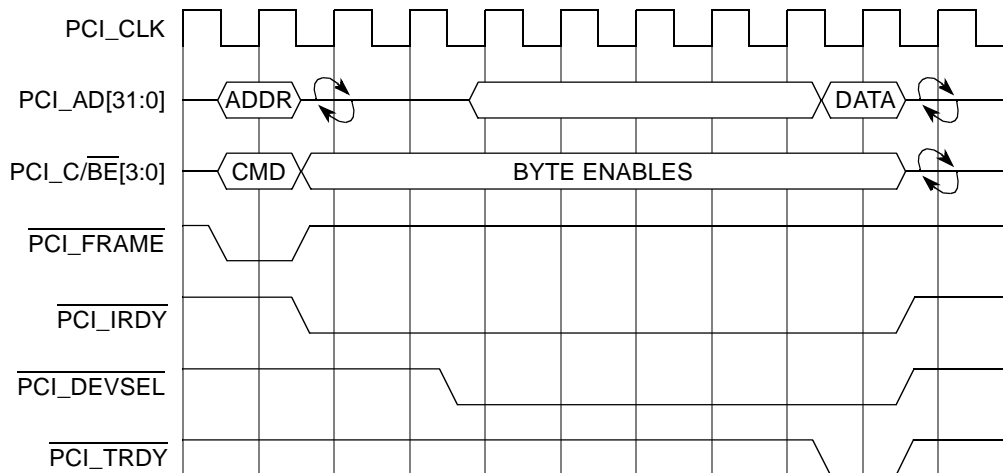


Figure 13-48. Single Beat Read Example

Figure 13-49 shows an example of a burst read transaction.

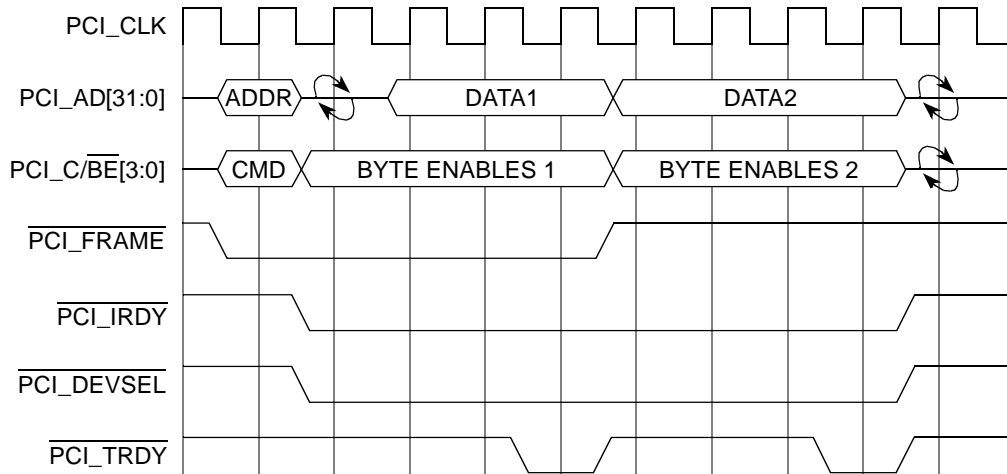


Figure 13-49. Burst Read Example

During the turnaround-cycle following the address phase, the $\overline{\text{PCI_C/BE}}[3:0]$ signals indicate which byte lanes are involved in the data phase. The turnaround-cycle must be enforced by the target with the $\overline{\text{PCI_TRDY}}$ signal if using fast $\overline{\text{PCI_DEVSEL}}$ assertion. The earliest the target can provide valid data is one cycle after the turnaround cycle. The target must drive the $\text{AD}[31:0]$ signals when $\overline{\text{PCI_DEVSEL}}$ is asserted except during the turnaround cycle.

The data phase completes when data is transferred, which occurs when both $\overline{\text{PCI_IRDY}}$ and $\overline{\text{PCI_TRDY}}$ are asserted on the same clock edge. When either is negated, a wait cycle is inserted and no data is transferred. To indicate the last data phase $\overline{\text{PCI_IRDY}}$ must be asserted when $\overline{\text{PCI_FRAME}}$ is negated.

A write transaction starts when $\overline{\text{PCI_FRAME}}$ is asserted for the first time and the $\overline{\text{PCI_C/BE}}[3:0]$ signals indicate a write command. Figure 13-50 shows an example of a single-beat write transaction.

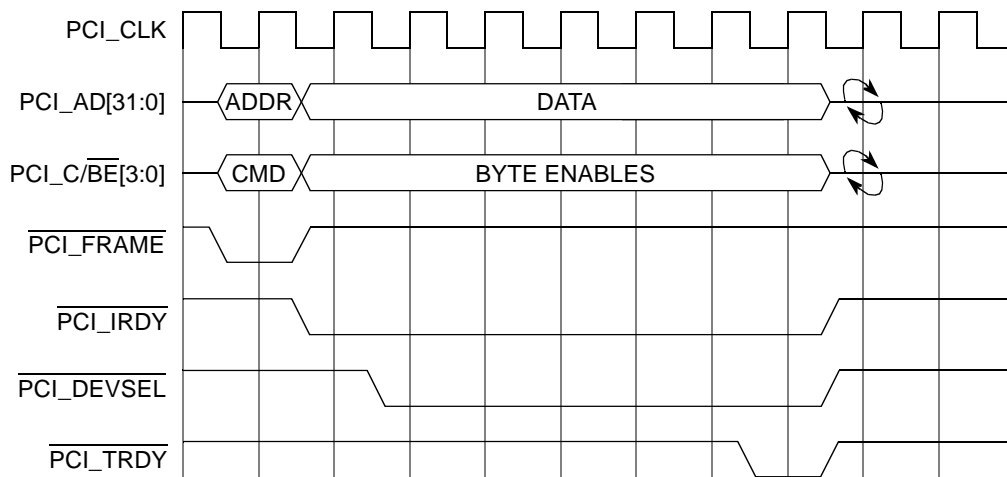


Figure 13-50. Single Beat Write Example

Figure 13-51 shows an example of a burst write transaction.

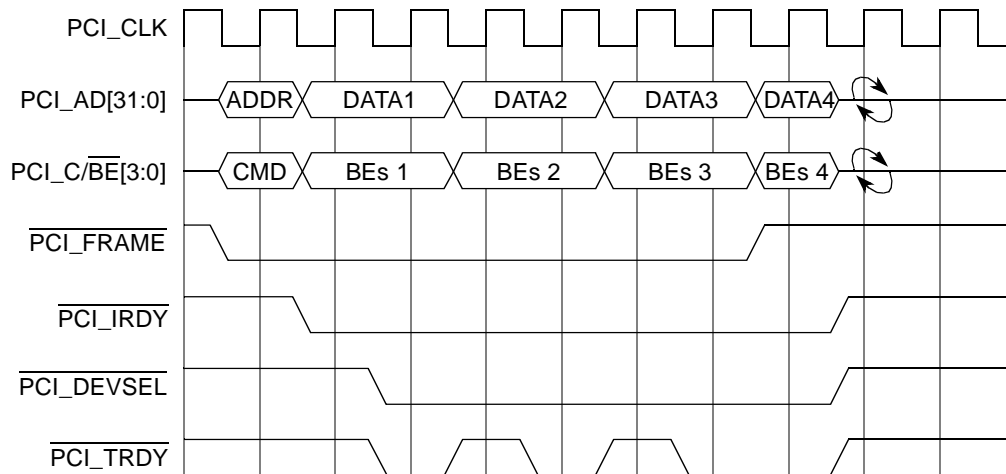


Figure 13-51. Burst Write Example

A write transaction is similar to a read transaction except no turnaround cycle is needed following the address phase because the initiator provides both address and data. Data phases are the same for both read and write transactions.

13.4.3.8 Transaction Termination

The termination of a PCI transaction is orderly and systematic, regardless of the cause of the termination. All transactions end when $\overline{\text{PCI_FRAME}}$ and $\overline{\text{PCI_IRDY}}$ are both negated, indicating the idle cycle.

The PCI controller as an initiator terminates a transaction when $\overline{\text{PCI_FRAME}}$ is negated and $\overline{\text{PCI_IRDY}}$ is asserted. This indicates that the final data phase is in progress. The final data transfer occurs when both $\overline{\text{PCI_TRDY}}$ and $\overline{\text{PCI_IRDY}}$ are asserted. A master-abort is an abnormal case of a master initiated termination. If the PCI controller detects that $\overline{\text{PCI_DEVSEL}}$ has remained negated for more than four clocks after the assertion of $\overline{\text{PCI_FRAME}}$, it negates $\overline{\text{PCI_FRAME}}$ and then, on the next clock, negates $\overline{\text{PCI_IRDY}}$. On aborted reads, the PCI controller returns 0xFFFF_FFFF. The data is lost on aborted writes.

When the PCI controller as a target needs to suspend a transaction, it asserts $\overline{\text{PCI_STOP}}$. Once asserted, $\overline{\text{PCI_STOP}}$ remains asserted until $\overline{\text{PCI_FRAME}}$ is negated. Depending on the circumstances, data may or may not be transferred during the request for termination. If $\overline{\text{PCI_TRDY}}$ and $\overline{\text{PCI_IRDY}}$ are asserted during the assertion of $\overline{\text{PCI_STOP}}$, data is transferred. This type of target-initiated termination is called a disconnect B, shown in Figure 13-52. If $\overline{\text{PCI_TRDY}}$ is asserted when $\overline{\text{PCI_STOP}}$ is asserted but $\overline{\text{PCI_IRDY}}$ is not, $\overline{\text{PCI_TRDY}}$ must remain asserted until $\overline{\text{PCI_IRDY}}$ is asserted and the data is transferred. This is called a disconnect A target-initiated termination, also shown in Figure 13-52. However, if $\overline{\text{PCI_TRDY}}$ is negated when $\overline{\text{PCI_STOP}}$ is asserted, no more data is transferred, and the initiator therefore does not have to wait for a final data transfer (see the retry diagram in Figure 13-50).

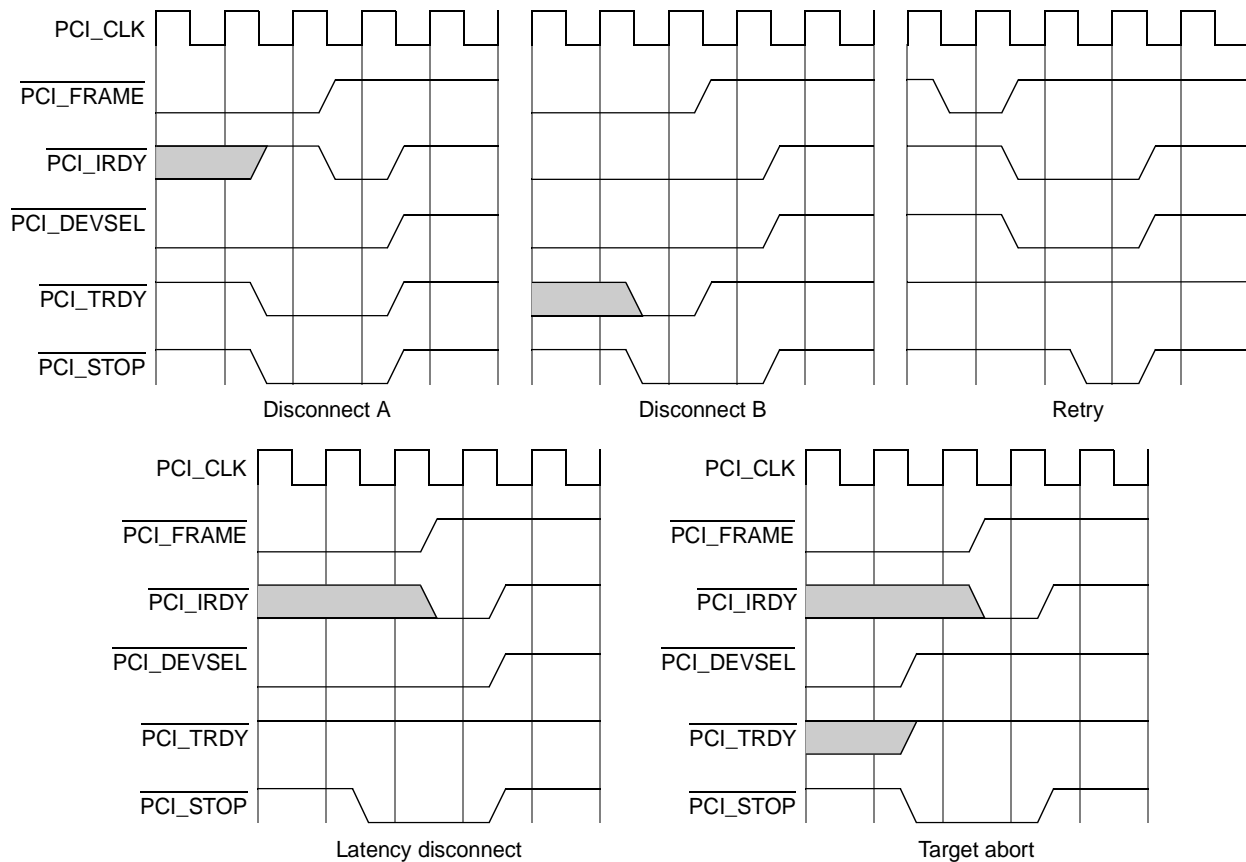


Figure 13-52. Target-Initiated Terminations

Note that when an initiator is terminated by $\overline{\text{PCI_STOP}}$, it must negate its $\overline{\text{REQn}}$ signal for a minimum of two PCI clocks (of which one clock is needed for the bus to return to the idle state). If the initiator intends to complete the transaction, it should reassert its $\overline{\text{REQn}}$ immediately following the two clocks or potential starvation may occur. If the initiator does not intend to complete the transaction, it can assert $\overline{\text{REQn}}$ whenever it needs to use the PCI bus again.

The PCI controller terminates a transaction in the following cases:

- Eight PCI clock cycles have elapsed between data phases. This is a ‘latency disconnect’ (see [Figure 13-50](#)).
- AD[1:0] is 0bx1 (a reserved burst ordering encoding) during the address phase and one data phase has completed.
- The PCI command is a configuration command and one data phase has completed.
- A streaming transaction crosses a 4-Kbyte page boundary.
- A streaming transaction runs out of I/O sequencer buffer entries.
- A cache line wrap transaction has completed a cache line transfer.

Another target-initiated termination is the retry termination. Retry refers to termination requested because the target is currently in a state where it is unable to process the transaction. This can occur because no buffer entries are available in the I/O sequencer, or the sixteen clock latency timer has expired without

transfer of the first data. The target latency timer of the PCI controller can be optionally disabled. See [Section 13.3.3.24, “PCI Function Configuration Register,”](#) for more information.

When the PCI controller is in host mode it does not respond to any PCI configuration transactions. When the PCI controller is in agent mode and the CFG_LOCK lock bit is set (see [Section 13.3.3.24, “PCI Function Configuration Register”](#)) the PCI controller retries all transactions to the PCI configuration space or the internal (on-chip) memory-mapped register space. Note that all retried accesses need to be completed. An example of a retry is shown in [Figure 13-50](#).

Note that because a target can determine whether or not data is transferred (when both $\overline{\text{PCI_IRDY}}$ and $\overline{\text{PCI_TRDY}}$ are asserted), if it wants to do only one more data transfer and then stop, it may assert $\overline{\text{PCI_TRDY}}$ and $\overline{\text{PCI_STOP}}$ at the same time.

Target-abort refers to the abnormal termination that is used when a fatal error has occurred, or when a target will never be able to respond. Target-abort is indicated when $\overline{\text{PCI_STOP}}$ is asserted and $\overline{\text{PCI_DEVSEL}}$ is negated. This indicates that the target requires the transaction to be terminated and does not want the transaction tried again. Note that any transferred data may have been corrupted.

The PCI controller terminates a transaction with target-abort in the case in which it is the intended target of a read transaction from system memory and the data from memory is corrupt. If the PCI controller is the intended target of a transaction and an address parity error occurs, or a data parity error occurs on a write transaction to system memory, it continues the transaction on the PCI bus but aborts internally. The PCI controller does not target-abort in this case.

If the PCI controller is mastering a transaction that terminates with a target-abort, undefined data is returned on a read and write data is lost. An example of a target-abort is shown in [Figure 13-50](#).

An initiator may retry any target disconnect accesses, except target-abort, at a later time starting with the address of the next non-transferred data. Retry is actually a special case of disconnect where no data transfer occurs at all and the initiator must start the entire transaction over again.

13.4.4 Other Bus Operations

The following sections provide information on additional PCI bus operations.

13.4.4.1 Fast Back-to-Back Transactions

In the two types of fast back-to-back transactions, the first type places the burden of avoiding contention on the initiator while the second places the burden on all potential targets. The PCI controller as a target supports both types of fast back-to-back transactions but does not support them as an initiator. The PCI controller as a target has the fast back-to-back enable bit hardwired to one, that is, enabled.

For the first type (governed by the initiator), the initiator may only run a fast back-to-back transaction to the same target. For the second type, when the PCI controller detects a fast-back-to-back operation and did not drive $\overline{\text{PCI_DEVSEL}}$ in the previous cycle, it delays the assertion of $\overline{\text{PCI_DEVSEL}}$ and $\overline{\text{PCI_TRDY}}$ for one cycle to allow the other target to get off the bus.

13.4.4.2 Dual Address Cycles

The PCI controller supports dual address cycle (DAC) commands (64-bit addressing on PCI bus) as a target only. DACs are different from single address cycles (SACs) in that the address phase takes two PCI beats instead of one PCI beat to transfer (64-bit vs. 32-bit addressing). Only PCI memory commands can use DAC cycles; I/O, configuration, interrupt acknowledge, and special cycle command cannot use DAC cycles. The PCI controller supports single-beat and burst DAC transactions.

13.4.4.3 Data Streaming

The PCI controller provides data streaming for PCI transactions to and from prefetchable memory. In other words, when the PCI controller is a target for a PCI initiated transaction, it supplies or accepts multiple cache lines of data without disconnecting. For PCI transactions to non-prefetchable space, the PCI controller disconnects after the first data phase so streaming cannot occur.

For PCI memory reads, streaming is achieved by performing speculative reads from memory in prefetchable space. A block of memory may be marked as prefetchable by setting the PCI configuration registers bit for the inbound address translation (see [Section 13.3.2.14, “PCI Inbound Window Attribute Registers \(PIWARn\),”](#) for more information) in the following cases:

- When reads do not alter the contents of memory (reads have no side effects)
- When reads return all bytes regardless of the byte enable signals
- When writes can be merged without causing errors

For a memory read command or a memory read line command, the PCI controller reads one cache line from memory. If the transaction crosses a cache line boundary, the PCI controller starts the read of a new cache line. For a memory read multiple command, the PCI controller reads two cache lines from memory. When the PCI transaction finishes the read for the first cache line, the PCI controller performs a speculative read of a third cache line. The PCI controller continues this prefetching until the end of the transaction.

For PCI writes to memory, streaming is achieved by buffering the transaction in the space available within the I/O sequencer. This allows PCI memory writes to execute with no wait states.

A disconnect occurs if the PCI controller runs out of buffer space on writes, or the PCI controller cannot supply consecutive data beats for reads within eight PCI bus clocks of each other. A disconnect also occurs if the transaction crosses a 4-Kbyte page boundary.

13.4.4.4 Host Mode Configuration Access

The PCI controller provides two types of configuration accesses to support hierarchical bridges. To access configuration space, a value is written to the CONFIG_ADDR register specifying which PCI bus, which device, and which configuration register to be accessed.

When the PCI controller sees an access that falls inside the 4 bytes beginning at the CONFIG_DATA address, it checks the enable bit, the device number and the bus number in the CONFIG_ADDR register. If the enable bit is set and the device number is not equal to all ones, a configuration cycle translation is performed. When the device number field is equal to all ones, it has a special meaning (see [Section 13.4.4.6, “Special Cycle Command,”](#) for more information).

There are two types of translations supported:

- Type 0 translations—For when the device is on the PCI bus connected to the PCI controller.
- Type 1 translations—For when the device is on another bus somewhere behind the PCI controller.

For type 0 translations, the PCI controller decodes the device number field to assert the appropriate IDSEL line and perform a configuration cycle on the PCI bus with AD[1:0] as 0b00. All 21 IDSEL bits are decoded, starting with bit AD11. That is, if the device number field contains 0b01011, AD11 on the PCI bus is set. The IDSEL lines are bit-wise associated with increasing values for the device number such that AD12 corresponds to 0b01100, and so on up to bit 30 as shown in [Table 13-41](#). AD31 is selected with 0b01010. A device number of 0b11111 indicates a special cycle. Device number 0b00000 is used for configuring the PCI controller itself. Bits 10 through 8 are copied to the PCI bus as an encoded value for components which contain multiple functions. Bits 7 through 2 are also copied onto the PCI bus. The PCI controller implements address stepping on configuration cycles so that the target's PCI_IDSEL, which is connected directly to one of the AD lines, reaches a stable value. This means that a valid address and command are driven on the AD and PCI_C/ $\overline{\text{BE}}$ lines one cycle before the assertion of $\overline{\text{PCI_FRAME}}$.

For type 1 translations, the PCI controller copies the contents of the CONFIG_ADDR register directly onto the PCI address/data lines during the address phase of a configuration cycle, with the exception that AD[1-0] contains 0b01 (not 0b00 as in Type 0 translations).

When the PCI controller is configured as a host device, a local master sometimes needs to perform configuration reads from unpopulated PCI slots (as part of the system configuration). To avoid getting a machine check interrupt, the following steps should be taken:

1. Mask the NORSP bit in the error mask register. See [Section 13.3.2.9, “PCI Error Control Register \(PCI_ECR\).”](#)
2. Perform the PCI configuration reads.
3. Clear the NORSP bit in the error status register.
4. Unmask (write 1) the NORSP bit in the error mask register. See [Section 13.3.2.3, “PCI Error Enable Register \(PCI_EER\).”](#)

13.4.4.5 Agent Mode Configuration Access

When the PCI controller is configured as an agent device, it responds to remote host generated PCI configuration accesses to the PCI interface. This is indicated by decoding the configuration command along with the PCI controller's IDSEL being asserted. A remote host can access the 256-byte PCI configuration area and the memory-mapped configuration registers within the PCI controller.

13.4.4.6 Special Cycle Command

A special cycle command contains no explicit destination address but is broadcast to all PCI agents. Each receiving agent must determine whether the message is applicable to itself. No assertion of $\overline{\text{PCI_DEVSEL}}$ in response to a special cycle command is necessary.

A special cycle command is like any other bus command in that it has an address phase and a data phase. The address phase starts like all other commands with the assertion of $\overline{\text{PCI_FRAME}}$ and completes when

$\overline{\text{PCI_FRAME}}$ and $\overline{\text{PCI_IRDY}}$ are negated. Special cycles terminate with a master-abort. (In the special cycle case, the received-master-abort bit in the configuration status register is not set.)

The address phase contains no valid information other than the command field. Even though there is no explicit address, the address/data lines are driven to a stable state and parity is generated. During the data phase, the address/data lines contain the message type and an optional data field. The message is encoded on the sixteen least-significant bits (AD[15:0]). The data field is encoded on AD[31:16]. When running a special cycle, the message and data are valid on the first clock $\overline{\text{PCI_IRDY}}$ is asserted.

When the PCI_CONFIG_ADDRESS register is written with a value so that the bus number matches the bridge bus, the device number is all ones, the function number is all ones, and the register number is zero. The next time the PCI_CONFIG_DATA register is accessed, the PCI controller executes either a special cycle or an interrupt acknowledge command. When the PCI_CONFIG_DATA register is written, the PCI controller generates a special cycle encoding on the command/byte enable lines during the address phase and drives the data from the PCI_CONFIG_DATA register onto the address/data lines during the first data phase.

If the bus number field of the PCI_CONFIG_ADDRESS does not match one of the PCI controller bus numbers, the PCI controller passes the write to PCI_CONFIG_DATA through to the PCI bus as a type 1 configuration cycle as it does any other time the bus number field does not match.

Table 13-46. Special Cycle Commands

| Address (AD[15-0]) | Message Type | Description |
|--------------------|------------------|---|
| 0x0000 | SHUTDOWN (SLEEP) | Indicates the processor is entering its most power saving mode |
| 0x0001 | HALT (DOZE) | Indicates the processor is entering a power save mode where address decoding is still available |
| 0x0002–0xFFFF | — | Reserved for future commands |

13.4.4.7 Interrupt Acknowledge

When the PCI_CONFIG_ADDRESS register is written with a value such that the bus number is 0x00, the device number is all ones, the function number is all ones, and the register number is zero, the next time the PCI_CONFIG_DATA register is accessed the PCI controller does either a special cycle command or an interrupt acknowledge command. When the PCI_CONFIG_DATA register is read, the PCI controller generates an interrupt acknowledge command encoding on the command/byte enable lines during the address phase. During the address phase, AD[31:0] do not contain a valid address but are driven with stable data and valid parity (PCI_PAR). During the data phase, the byte enable signals determine which bytes are involved in the transaction. The interrupt vector must be returned when $\overline{\text{PCI_TRDY}}$ is asserted.

An interrupt acknowledge transaction can also be issued on the PCI bus by reading from the PCI_INT_ACK register.

13.4.5 Error Functions

This section describes PCI bus errors.

13.4.5.1 Parity

During valid 32-bit address and data transfers, parity covers all 32 address/data lines and the 4 command/byte enable lines regardless of whether or not all lines carry meaningful information. Byte lanes not actually transferring data are driven with stable (albeit meaningless) data and are included in the parity calculation. During configuration, special cycle or interrupt acknowledge commands, some address lines are not defined but are still driven to stable values and included in the parity calculation.

Even parity is calculated for all PCI operations: the value of PCI_PAR is generated such that the number of ones on PCI_AD[31:0], PCI_C/ $\overline{\text{BE}}$ [3:0] and PCI_PAR equals an even number. The PCI_PAR signal is driven when the address/data lines are driven and follow the corresponding address or data by one clock.

The PCI controller checks the parity after all valid address phases (the assertion of $\overline{\text{PCI_FRAME}}$) and for valid data transfers ($\overline{\text{PCI_IRDY}}$ and $\overline{\text{PCI_TRDY}}$ asserted) involving the PCI controller. When an address or data parity error is detected, the detected-parity-error bit in the configuration space status register is set (see [Section 13.3.3.4, “PCI Status Configuration Register.”](#))

13.4.5.2 Error Reporting

Except for setting the detected-parity-error bit, all parity error reporting and response is controlled by the parity-error-response bit (see [Section 13.3.3.3, “PCI Command Configuration Register,”](#) for more information). If the parity-error-response bit is cleared, the PCI controller completes all transactions regardless of parity errors (address or data). If the bit is set, the PCI controller asserts $\overline{\text{PCI_PERR}}$ two clocks after the actual data transfer in which a data parity error is detected, and keeps $\overline{\text{PCI_PERR}}$ asserted for one clock. When acting as an initiator during a read transaction or as a target involved in a write to system memory the PCI controller asserts $\overline{\text{PCI_PERR}}$.

Figure 13-53 shows the possible assertion points for $\overline{\text{PCI_PERR}}$ if the PCI controller detects a data parity error.

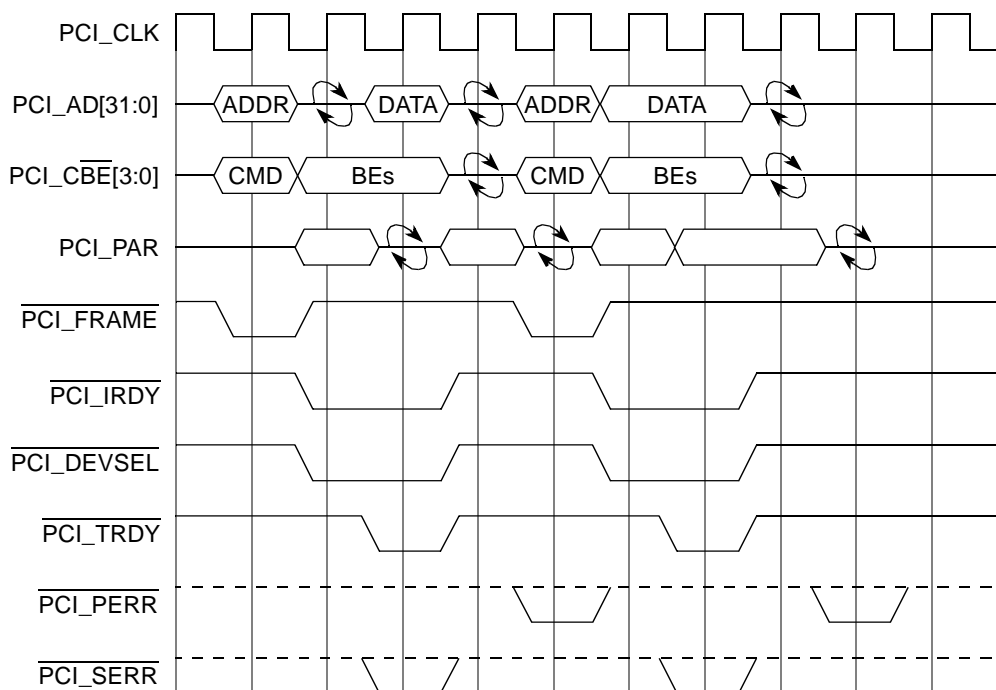


Figure 13-53. PCI Parity Operation

As an initiator, the PCI controller attempts to complete the transaction on the PCI bus if a data parity error is detected and sets the data-parity-reported bit in the configuration space status register. If a data parity error occurs on a read transaction, the PCI controller aborts the transaction internally. As a target, the PCI controller completes the transaction on the PCI bus even if a data parity error occurs. If parity error occurs during a write to system memory, the transaction completes on the PCI bus, but is aborted internally, insuring that potentially corrupt data does not go to memory.

When the PCI controller asserts $\overline{\text{PCI_SERR}}$, it sets the signaled-system-error bit in the configuration space status register. Additionally, if the error is an address parity error, the parity-error-detected bit is set; reporting an address parity error on $\overline{\text{PCI_SERR}}$ is conditioned on the parity-error-response bit being enabled in the command register. $\overline{\text{PCI_SERR}}$ is asserted when the PCI controller detects an address parity error while acting as a target. The system error is passed to the PCI controller's interrupt processing logic to assert $\overline{\text{MCP}}$. Figure 13-53 shows where the PCI controller could detect an address parity error and assert $\overline{\text{PCI_SERR}}$ or where the PCI controller, acting as an initiator, checks for the assertion of $\overline{\text{PCI_SERR}}$ signaled by the target detecting an address parity error.

As a target that asserts $\overline{\text{PCI_SERR}}$ on an address parity, the PCI controller completes the transaction on the PCI bus, aborting internally if the transaction is a write to system memory. If $\overline{\text{PCI_PERR}}$ is asserted during a PCI controller write to PCI, the PCI controller attempts to continue the transfer, allowing the target to abort/disconnect if desired. If the PCI controller detects a parity error on a read from PCI, the PCI controller aborts the transaction internally and continues the transfer on the PCI bus, allowing the target to abort/disconnect if desired.

In all cases of parity errors on the PCI bus, regardless of the parity-error-response bit, information about the transaction is logged in the PCI error control capture register, the PCI error address capture register and the PCI error data capture register; \overline{MCP} is also asserted to the core as an option.

13.4.6 PCI Inbound Address Translation

For inbound transactions (transactions generated by an external master on the PCI bus where the PCI controller responds as a slave device), the PCI controller only responds to PCI addresses within the windows mapped by the PCI inbound base address registers (PIBARs). If there is an address hit in one of the PIBARs, the PCI address is translated from PCI space to local memory space through the associated PCI inbound translation address registers (PITARs). This allows an external master to access local memory. Each PIBAR register is associated with a PITAR and PIWAR which are located in the PCI controller’s PCI CSR space. [Figure 13-54](#) shows an example translation window for inbound memory accesses.

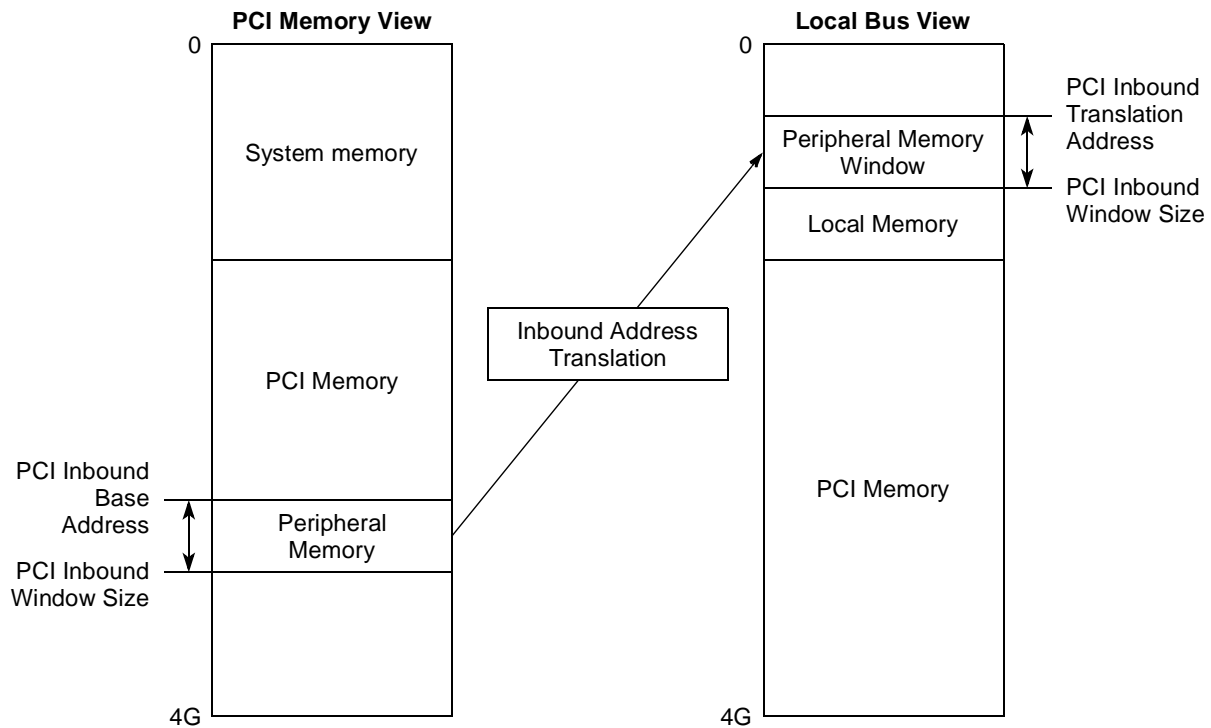


Figure 13-54. Inbound PCI Memory Address Translation

There are three full sets of inbound translation registers, in addition to the PIMMR base address register, allowing four simultaneous translation windows, one to a fixed destination and three programmable. Only two of the programmable windows can be mapped anywhere in the 64-bit PCI address space. Window 0 can only be mapped within the lowest 4-Gbyte space. Software can move the programmable translation base addresses during run-time to access different portions of local memory, but the PCI inbound translation windows may not overlap.

The translation windows are disabled after reset, that is, after reset, the PCI controller does not acknowledge externally mastered transactions on the PCI bus by asserting $\overline{PCI_DEVSEL}$ until the inbound translation windows are enabled.

13.4.7 CompactPCI Hot Swap Specification Support

CompactPCI is an open specification supported by the PCI Industrial Computer Manufacturers Group (PICMG) and is intended for embedded applications using PCI. CompactPCI Hot Swap is an extension of the CompactPCI specification and allows the insertion and extraction (or “hot swapping”) of boards without adversely affecting system operation. The hot swap specification defines the following levels of support:

- Hot swap capable
- Hot swap friendly
- Hot swap ready

The PCI controller is hot swap friendly, meaning that it supports the hardware and software connection processes as defined in the hot swap specification. This level of support allows the board and system designers to build full Hot Swap and high availability systems based on the PCI controller as a PCI target device. For details on the hot swap process, refer to the *Hot Swap Specification PICMG 2.1, R1.0, August 3, 1998*.

13.4.8 Byte Ordering

Whenever data must cross a bridge between two busses, the byte ordering of data on the source and destination buses must be considered. The internal platform bus of this device is inherently big endian and the PCI bus interface is inherently little endian.

There are two methods to handle ordering of data as it crosses a bridge—address invariance and data invariance. Address invariance preserves the addressing of bytes within a scalar data element, but not the relative significance of the bytes within that scalar. Conversely, data invariance preserves the relative significance of bytes within a scalar, but not the addressing of the individual bytes that make up a scalar.

This device uses address invariance as its byte ordering policy

As stated above, address invariance preserves the byte address of each byte on an I/O interface as it is placed in memory or moved into a register. This policy can have the effect of reversing the significance order of bytes (most significant to least significant and vice versa), but it has the benefit of preserving the format of general data structures. Provided that software is aware of the endianness and format of the data structure, it can correctly interpret the data on either side of the bridge.

Figure 13-55 shows the transfer of a 4-byte scalar, 0x4142_4344, from a big endian source across an address invariant bridge to a little endian destination.




Figure 13-55. Address Invariant Byte Ordering—4 bytes Outbound

Note that although the significance of the bytes within the scalar have changed, the address of the individual bytes that make up the scalar have not changed. As long as software is aware that the source of the data used a big endian format, the data can be interpreted correctly.

Figure 13-56 shows data flowing the other way, from a little endian source to a big endian destination.

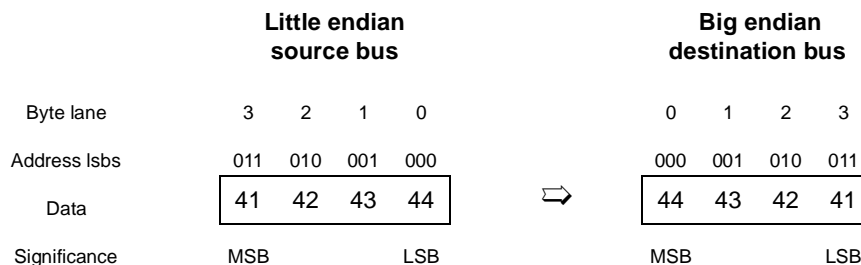


Figure 13-56. Address Invariant Byte Ordering—4 bytes Inbound

Figure 13-57 shows an outbound transfer of an 8-byte scalar, 0x5455_1617_CDCE_2728, using address invariance.




Figure 13-57. Address Invariant Byte Ordering—8 bytes Outbound

Figure 13-58 shows an inbound transfer of a 2-byte scalar, 0x5837, using address invariance.




Figure 13-58. Address Invariant Byte Ordering—2 bytes Inbound

Note that in all of these examples, the original addresses of the individual bytes within the scalars (as created by the source) have been preserved.

13.4.8.1 Byte Order for Configuration Transactions

All internal memory-mapped registers in the CCSR space use big endian byte ordering. However, the PCI specification defines PCI configuration registers as little endian. All accesses to the PCI configuration port, CFG_DATA, including the those targeting the internal PCI configuration registers, use the address invariance policy as shown in Figure 13-59. Therefore, software must access CFG_DATA with

little-endian formatted data—either using the **lwbrx/stwbrx** instructions or by manipulating the data before writing to and after reading from **CFG_DATA**.




Figure 13-59. CFG_DATA Byte Ordering

13.5 Initialization/Application Information

The following sections describe initialization sequences for host and agent modes.

13.5.1 Initialization Sequence for Host Mode

The following sequence must be followed in host mode:

1. Enable PCI output clocks and select desired frequency ratios. See [Section 4.4.1, “Clocking in PCI Host Mode.”](#)
2. Wait for at least 1 ms to enable stable clocks into agent devices
3. Deactivate **PCI_RESET_OUT** signal for PCI. See [Table 13-3](#) for more information on **PCI_RESET_OUT** signal.
4. Wait for at least 1 ms to enable devices to complete the powerup sequence.
5. Configure PCI internal registers and PCI agents to desired modes of operation

13.5.2 Initialization Sequence for Agent Mode

The following sequence must be followed in agent mode:

1. Optionally initialize subsystem vendor ID/device ID
 - a) Initialize PCI inbound window size in **PIWAR[1:3]** desired window size
 - b) Unlock configuration lock in PCI function configuration register

Chapter 14

PCI Express Interface Controller

The MPC8315E and MPC8314E PCI Express interface is compatible with the *PCI Express™ Base Specification*, Revision 1.0a (available from <http://www.pcisig.org>). It is beyond the scope of this manual to document the intricacies of the PCI Express protocol. This chapter describes the PCI Express controller of this device and provides a basic description of the PCI Express protocol. The specific emphasis is directed at how the device implements the PCI Express specification. Designers of systems incorporating PCI Express devices should refer to the specification for a thorough description of PCI Express.

NOTE

Much of the available PCI Express literature refers to a 16-bit quantity as a WORD and a 32-bit quantity as a DWORD. This is inconsistent with the terminology in the rest of this manual where the terms ‘word’ and ‘double word’ refer to a 32-bit and 64-bit quantity, respectively. Where necessary to avoid confusion, the precise number of bits or bytes is specified.

NOTE

The PCI Express engine does not support misaligned byte transfers. It must be DWORD aligned to the CSB bus.

14.1 Introduction

The PCI Express controller is a mechanism for communicating with PCI Express devices. The controller contains three major parts:

- PCI Express core—Handles the transaction, data link and MAC layers and contains the configuration header and control registers.
- CSB bridge—Controls the transfer of the transactions between the PCI Express transaction layer and the CSB, and include Write and Read DMA engines, a message manager and a set of configuration registers.
- SerDes—Controls the transfer between the PCI Express MAC layer and the physical link, and includes another set of configuration registers (described in [Chapter 16, “SerDes PHY”](#)).

Figure 14-1 is a high-level block diagram of the PCI Express controller.




Figure 14-1. PCI Express Controller Block Diagram

The PCI Express controller connects the coherent system bus (CSB) to the PCI Express bus, which is a 2.5-GHz serial interface that supports up to a x1 lane. As both a master (initiator) and a target device, the PCI Express interface is capable of high bandwidth data transfer and is designed to support the next generation of I/O devices. When it comes out of reset, the PCI Express interface performs link width negotiation and exchanges flow control credits with its link partner. When link auto-negotiation finishes, the controller is ready for operation.

Internally, the design contains queues to keep track of inbound and outbound transactions. Control logic handles buffer management, bus protocol, transaction spawning, and tag generation. In addition, memory blocks store inbound and outbound data.

The device can be configured to operate in either root complex (RC) or endpoint (EP) mode. An RC device connects the CPU/memory subsystem to the I/O devices, while an EP device typically denotes a peripheral or I/O device. In RC mode, a type 1 configuration header is used. In EP mode, a type 0 configuration header is used.

As an initiator, the device supports memory read and write operations with a maximum payload of 128 bytes. In addition, outbound configuration are supported if the device is in RC mode. As a target interface, the device accepts read and write operations to local memory space. Furthermore, as an EP device, the device accepts configuration transactions to the internal PCI Express configuration registers. Inbound I/O transactions are not supported.

14.1.1 MPC8315E/MPC8314E as a PCI Express Initiator

Outbound CSB transactions to PCI Express are first mapped to a translation window to determine which PCI Express transactions are to be issued. A transaction from the CSB can become a memory, I/O, or configuration transaction on the PCI Express bus depending on the window attributes. A transaction can be broken up into smaller transactions depending on the original request size, transaction type, PCI Express Device Control register's Max_Payload_Size field (for writes) and PCI Express Device Control Register's Max_Read_Request_Size field (for reads). The device performs PCI Express ordering rule checks to determine the next transaction to be sent on the PCI Express bus. In general, transactions are serviced in the order they are received from the CSB. The device allows reordering of higher-priority transactions to bypass lower-priority transactions only when there is a stall condition. For posted write transactions, after all data is received on the CSB, the data is forwarded to the PCI Express bus and the transaction is considered to be complete. For non-posted read transactions, the device waits for all completion packets to return from the link partner and then forwards all data back to the CSB before terminating the transaction.

There are two methods of generating PCI Express outbound transactions:

- One of the CSB masters, such as the e300 host, directly initiates a transaction. This is referred to as "PIO."
- The write or read DMA engines, which are part of the PCI Express controller CSB bridge, is used.

The DMA method is more efficient for transferring large chunks of data. The outbound windows are used and shared by both methods.

14.1.2 MPC8315E/MPC8314E as a PCI Express Target

Inbound PCI Express transactions to the CSB are first mapped to the CSB address space through a translation window. A transaction can be broken up into smaller transactions when it is sent to the CSB depending on the original size, byte enables and starting/ending addresses. The device performs PCI Express ordering rule checks to determine the next transaction to be sent to the CSB. In general, transactions are serviced in the order they are received from the PCI Express bus. The device allows reordering of higher-priority transactions to bypass lower-priority transactions only when there is a stall condition. For posted write transactions, after all data is received on the PCI Express bus, the data is forwarded to the CSB and the transaction is considered to be complete. For non-posted read transactions, the device waits for enough completion packets (dependent on the packet length) to return and then forwards data back to the PCI Express bus. This process continues until there are no more completion packets left to be sent.

14.1.3 Features

The following is a list of PCI Express controller features:

- Designed to be compatible with the *PCI Express Base Specification, Version 1.0a*
- Root complex (RC) and endpoint (EP) configurations
- 32- and 64-bit address support
- Two PCI Express links of x1 lane each
- Access to all PCI Express memory
- Access to I/O address spaces as requestor only in RC mode
- Posting of processor-to-PCI Express and PCI Express-to-memory writes
- Strong and relaxed transaction ordering rules
- PCI Express configuration registers
- Baseline and advanced error reporting
- One virtual channel (VC0) per controller
- 128-byte maximum payload size (Max_Payload_Size) for memory read and write operations
- Four inbound general-purpose translation windows per controller
- Four outbound translation windows per controller
- Up to four outstanding PCI Express transactions from each controller (posted or non-posted)
- Credit-based flow control management
- PCI Express messages and interrupts
- Maximum 32-byte payload transactions from the CSB
- Interrupt generation from messages or upon detection of errors
- Read and Write DMA engines per controller

14.1.4 Modes of Operation

This section describes how some parameters that affect the PCI Express controller operating modes are determined at power-on reset (POR) by the reset configuration words.

14.1.4.1 Root Complex/Endpoint Modes

The PCI Express controller can function as either a root complex (RC) or an endpoint (EP) device. The PCI Express control registers 1 and 2 determine the RC/EP mode; see [Section 5.3.2.10, “PCI Express Control Registers \(PECR1 and PECR2\).”](#)

14.1.4.2 Link Width

The link width of each PCI Express controller is x1.

14.1.4.3 Reference Clock

The reference clock for the PCI Express PHY is determined based on the TSEC1M and TSEC2M fields in the RCWH. If the value of any of these fields is different from 0b110, the PHY is programmed to operate with a reference clock frequency of 100 MHz. To change the reference clock frequency after POR, it is required to program the SRDSCR4 (see [Section 16.3.5, “SerDes Control Register 4 \(SRDSCR4\)”](#)) and initiate a SerDes PHY reset sequence..

14.2 External Signal Descriptions

PCI Express defines the connection between two devices as a link, which can be composed of a single lane or multiple lanes. Each lane consists of a differential pair for transmitting (TX_n and \overline{TX}_n) and a differential pair for receiving (RX_n and \overline{RX}_n) with an embedded data clock.

[Table 14-1](#) describes the external PCI Express interface signals.

Table 14-1. PCI Express Interface Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|--|-----|---|
| $\overline{RXA}, \overline{RXB}/RXA/RXB$ | I | Receive data. Receive data differential signal pair carry PCI Express packet information. |
| $\overline{TXA}, \overline{TXB}/TXA/TXB$ | O | Transmit data. The transmit data differential signal pair carry PCI Express packet information. |

14.3 Memory Map/Register Definitions

The PCI Express interface supports the following register types:

- Memory-mapped registers—Control PCI Express address translation, PCI error management, and PCI Express configuration register access on the device. These registers are described in [Section 14.3.1, “PCI Express Memory Map.”](#)
- PCI Express configuration registers within the PCI Express configuration header—Specified by the PCI Express specification for every PCI Express device. They are described in [Section 14.4.1, “Common PCI-Compatible Configuration Header Registers.”](#)

From the PCI Express side, the configuration header registers can be accessed through configuration access, and the memory-mapped registers can be accessed through memory transactions after the inbound translation window is programmed. From the CSB side, all these registers are memory-mapped.

14.3.1 PCI Express Memory Map

The PCI Express memory-mapped registers, listed in [Table 14-3](#), are accessed by reading and writing to an address composed of the base address (specified in the IMMRBAR on the CSB side, or the ATMU windows on the PCI Express side) plus the offset of the specific register to be accessed. In this table and in the register figures and fields description, the following access definitions apply:

- Reserved fields are always ignored for the purposes of determining access type.
- R/W, R, and W (read/write, read only, and write only) indicate that all the non-reserved fields in a register have the same access type.

- w1c indicates that all of the non-reserved fields in a register are cleared by writing ones to them.
- Mixed indicates a combination of access types.
- Special is used when no other category applies. In this case the register figure and field description table should be read carefully.

Also note that although the table explicitly lists only the registers for the PCI Express Controller 1, the register map for PCI Express Controller 2 is the same except the for the block base address. Memory-mapped registers for PCI Express Controller 1 begin at block base address 0x0_9000 and the memory-mapped registers for PCI Express Controller 2 begin at block base address 0x0_A000.

Table 14-2 below lists the address ranges for each type of register. Undefined address spaces are reserved.

Table 14-2. PCI Express Controller Register Groups

| Register | Offset Range | Section/Page |
|--|--------------|--------------------------------|
| PCI Express Core Registers | | |
| Common PCI-Compatible Configuration Header Registers | 0x000–0x3FF | 14.4.1/14-15 |
| PCI Express Core Control and Status Registers (CSRs) | 0x400–0x4CF | 14.4.6/14-61 |
| PCI Express BAR Configuration Registers (EP Mode) | 0x4D8–0x4FF | 14.4.7/14-71 |
| PCI Express Extended Status and Control Registers | 0x590–0x7FF | 14.4.8/14-74 |
| PCI Express CSB Bridge Registers | | |
| Global Registers | 0x800–0x83F | 14.5.2/14-77 |
| PCI Express Outbound PIO Registers | 0x840–0x8DF | 14.5.3/14-79 |
| PCI Express Inbound PIO Registers | 0x8E0–0x87F | 14.5.4/14-81 |
| PCI Express DMA Registers | 0x990–0xADF | 14.5.5/14-82 |
| Mailbox Registers | 0xB20–0xB9F | 14.5.6/14-86 |
| PCI Express Host Interrupt Registers | 0xBA0–0XBDF | 14.5.7/14-88 |
| CSB System Interrupt Registers | 0xBE0–0xC1F | 14.5.8/14-93 |
| PCI Express Outbound Address Mapping Registers | 0xCA0–0xDDF | 14.5.10/14-102 |
| PCI Express EP Inbound Address Translation Registers | 0xDE0–0xE5F | 14.5.11/14-105 |
| PCI Express RC Inbound Address Mapping Registers | 0xE60–0xFFF | 14.5.12/14-106 |

Table 14-3. PCI Express Memory Map

| Offset | Register | Access | Reset | Section/Page |
|--|--------------------------------|--------|-----------------|--------------------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| PCI Express Controller 1 Registers | | | | |
| PCI Express1 Core Configuration Header Registers | | | | |
| 0x000 | PCI Express Vendor ID Register | R | 0x1957 | 14.4.1.1/14-15 |
| 0x002 | PCI Express Device ID Register | R | Device-specific | 14.4.1.2/14-16 |

Table 14-3. PCI Express Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--|---|---------|----------------------------------|----------------------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| 0x004 | PCI Express Command Register | Mixed | All zeros | 14.4.1.3/14-16 |
| 0x006 | PCI Express Status Register | Mixed | 0x0010 | 14.4.1.4/14-17 |
| 0x008 | PCI Express Revision ID Register | R | Revision-specific | 14.4.1.5/14-18 |
| 0x009 | PCI Express Class Code Register | Mixed | 0x0B20 | 14.4.1.6/14-19 |
| 0x00C | PCI Express Cache Line Size Register | R/W | 0x00 | 14.4.1.7/14-19 |
| 0x00D | PCI Express Latency Timer Register | R | 0x00 | 14.4.1.8/14-20 |
| 0x00E | PCI Express Header Type Register | R | 0x00 (EP mode) 0x01 (RC mode) | 14.4.1.9/14-20 |
| 0x00F | PCI Express BIST Register | R | 0x00 | 14.4.1.10/14-21 |
| 0x010– 0x014 | Base Address Registers 0 and 1 (BAR0/BAR1) (EP mode only) | Mixed | 0x0008 | 14.4.2.1.1/14-22 |
| 0x018– 0x020 | Base Address Registers 2 and 4 (BAR2/BAR4) (EP mode only) | Mixed | 0x0000_000C | 14.4.2.1.2/14-23 |
| 0x01C– 0x024 | Base Address Registers 3 and 5 (BAR3/BAR5) (EP mode only) | R/W | All zeros | 14.4.2.1.3/14-23 |
| 0x02C | PCI Express Subsystem Vendor ID Register (EP mode only) | Special | All zeros | 14.4.2.2/14-24 |
| 0x02E | PCI Express Subsystem ID Register (EP mode only) | Special | All zeros | 14.4.2.3/14-24 |
| 0x034 | PCI Express Capabilities Pointer Register | R | 0x0044 | 14.4.2.4/14-25 |
| 0x03C | PCI Express Interrupt Line Register (EP mode only) | R/W | All zeros | 14.4.2.5/14-25 |
| 0x03D | PCI Express Interrupt Pin Register | R | 0x0001 | 14.4.2.6/14-26 |
| 0x03E | PCI Express Minimum Grant Register (EP mode only) | R | All zeros | 14.4.2.7/14-26 |
| 0x03F | PCI Express Maximum Latency Register (EP mode only) | R | All zeros | 14.4.2.8/14-26 |
| 0x018 | PCI Express Primary Bus Number Register (RC mode only) | R/W | All zeros | 14.4.3.1/14-27 |
| 0x019 | PCI Express Secondary Bus Number Register (RC mode only) | R/W | All zeros | 14.4.3.2/14-28 |
| 0x01A | PCI Express Subordinate Bus Number Register (RC mode only) | R/W | All zeros | 14.4.3.3/14-28 |
| 0x01B | Secondary Latency Timer Register 2 (RC mode only) | — | All zeros | — |
| 0x01C | PCI Express I/O Base Register (RC mode only) | R | All zeros | 14.4.3.5/14-29 |
| 0x01D | PCI Express I/O Limit Register (RC mode only) | R | All zeros | 14.4.3.6/14-29 |
| 0x01E | PCI Express Secondary Status Register (RC mode only) | Mixed | All zeros | 14.4.3.7/14-30 |
| 0x020 | PCI Express Memory Base Register (RC mode only) | R/W | All zeros | 14.4.3.8/14-31 |
| 0x022 | PCI Express Memory Limit Register (RC mode only) | R/W | All zeros | 14.4.3.9/14-31 |
| 0x024 | PCI Express Prefetchable Memory Base Register (RC mode only) | R/W | All zeros | 14.4.3.10/14-32 |
| 0x026 | PCI Express Prefetchable Memory Limit Register (RC mode only) | R/W | All zeros | 14.4.3.11/14-32 |

Table 14-3. PCI Express Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--|---|--------|-------------|---------------------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| 0x028 | PCI Express Prefetchable Base Upper 32-Bit Register (RC mode only) | R/W | All zeros | 14.4.3.12/14-33 |
| 0x02C | PCI Express Prefetchable Limit Upper 32-Bit Register (RC mode only) | R/W | All zeros | 14.4.3.13/14-33 |
| 0x030 | PCI Express I/O Base Upper 16-Bit Register (RC mode only) | R | All zeros | 14.4.3.14/14-33 |
| 0x032 | PCI Express I/O Limit Upper 16-Bit Register (RC mode only) | R' | All zeros | 14.4.3.15/14-34 |
| 0x034 | PCI Express Capabilities Pointer Register | R | 0x044 | 14.4.3.16/14-34 |
| 0x03C | PCI Express Interrupt Line Register | R/W | All zeros | 14.4.3.17/14-35 |
| 0x03D | PCI Express Interrupt Pin Register | R | 0x0001 | 14.4.3.18/14-35 |
| 0x03E | PCI Express Bridge Control Register (RC mode only) | R/W | All zeros | 14.4.3.19/14-35 |
| 0x044 | PCI Express Power Management Capability ID Register | R | 0x01 | 14.4.4.1/14-38 |
| 0x045 | PCI Express Power Management Next Capabilities Pointer Register | R | 0x4C | 14.4.4.2/14-38 |
| 0x046 | PCI Express Power Management Capabilities Register | R | 0x7E02 | 14.4.4.3/14-39 |
| 0x048 | PCI Express Power Management Status and Control Register | Mixed | All zeros | 14.4.4.4/14-39 |
| 0x04B | PCI Express Power Management Data Register | R | All zeros | 14.4.4.5/14-40 |
| 0x04C | PCI Express Capability ID Register | R | 0x10 | 14.4.4.6/14-40 |
| 0x04D | PCI Express Next Capabilities Pointer Register | R | 0x70 | 14.4.4.7/14-41 |
| 0x04E | PCI Express Capabilities Register | R | 0x00n1 | 14.4.4.8/14-41 |
| 0x050 | PCI Express Device Capabilities Register | R | All zeros | 14.4.4.9/14-42 |
| 0x054 | PCI Express Device Control Register | R/W | 0x2810 | 14.4.4.10/14-43 |
| 0x056 | PCI Express Device Status Register | Mixed | All zeros | 14.4.4.11/14-44 |
| 0x058 | PCI Express Link Capabilities Register | R | 0x0003_B481 | 14.4.4.12/14-44 |
| 0x05C | PCI Express Link Control Register | R/W | All zeros | 14.4.4.13/14-45 |
| 0x05E | PCI Express Link Status Register | R | 0x0011 | 14.4.4.14/14-46 |
| 0x060 | PCI Express Slot Capabilities Register | R | 0x000007c0 | 14.4.4.15/14-46 |
| 0x064 | PCI Express Slot Control Register | R/W | All zeros | 14.4.4.16/14-47 |
| 0x066 | PCI Express Slot Status Register | Mixed | 0x0040 | 14.4.4.17/14-48 |
| 0x068 | PCI Express Root Control Register (RC mode only) | R/W | All zeros | 14.4.4.18/14-48 |
| 0x06C | PCI Express Root Status Register (RC mode only) | Mixed | All zeros | 14.4.4.19/14-49 |
| 0x070 | PCI Express MSI Message Capability ID Register (EP mode only) | R | 0x05 | 14.4.4.20/14-49 |
| 0x072 | PCI Express MSI Message Control Register (EP mode only) | Mixed | 0x0088 | 14.4.4.21/14-50 |
| 0x074 | PCI Express MSI Message Address Register (EP mode only) | R/W | All zeros | 14.4.4.22/14-50 |

Table 14-3. PCI Express Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--|--|--------|-------------|-----------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| 0x078 | PCI Express MSI Message Upper Address Register (EP mode only) | R/W | All zeros | 14.4.4.23/14-51 |
| 0x07C | PCI Express MSI Message Data Register (EP mode only) | R/W | 0x0000 | 14.4.4.24/14-51 |
| 0x100 | PCI Express Advanced Error Reporting Capability ID Register | R | 0x1381_0001 | 14.4.5.1/14-53 |
| 0x104 | PCI Express Uncorrectable Error Status Register | R/W | All zeros | 14.4.5.2/14-53 |
| 0x108 | PCI Express Uncorrectable Error Mask Register | R/W | All zeros | 14.4.5.3/14-54 |
| 0x10C | PCI Express Uncorrectable Error Severity Register | R/W | 0x0006_2010 | 14.4.5.4/14-55 |
| 0x110 | PCI Express Correctable Error Status Register | w1c | All zeros | 14.4.5.5/14-56 |
| 0x114 | PCI Express Correctable Error Mask Register | R/W | All zeros | 14.4.5.6/14-57 |
| 0x118 | PCI Express Advanced Error Capabilities and Control Register | R/W | 0x0000_00A0 | 14.4.5.7/14-57 |
| 0x11C | PCI Express Header Log Register | R | All zeros | 14.4.5.8/14-58 |
| 0x120 | PCI Express Header Log Register | R | All zeros | |
| 0x124 | PCI Express Header Log Register | R | All zeros | |
| 0x128 | PCI Express Header Log Register | R | All zeros | |
| 0x12C | PCI Express Root Error Command Register | R/W | All zeros | 14.4.5.9/14-59 |
| 0x130 | PCI Express Root Error Status Register | Mixed | All zeros | 14.4.5.10/14-59 |
| 0x134 | PCI Express Error Source Identification Register | R | All zeros | 14.4.5.11/14-60 |
| PCI Express Core Control and Status Registers (CSRs) | | | | |
| 0x404 | PCI Express LTSSM State Status Register (PEX_LTSSM_STAT) | R | All zeros | 14.4.6.1/14-61 |
| 0x41C | PCI Express N_FTS Control Register (PEX_NFTS_CTRL) | R/W | 0x0000_4040 | 14.4.6.2/14-62 |
| 0x438 | PCI Express ACK Replay Timeout Register (PEX_ACKRPLY_TO) | R/W | 0x005B_2090 | 14.4.6.3/14-63 |
| 0x440 | PCI Express Core Clock Ratio Register (PEX_GCLK_RATIO) | R/W | 0x0000_0010 | 14.4.6.4/14-64 |
| 0x450 | PCI Express Power Management Timer Register (PEX_PM_TIMER) | R/W | 0x000A_63E4 | 14.4.6.5/14-65 |
| 0x454 | PCI Express PME Time-Out Register (PEX_PME_TIMEOUT) | R/W | 0x00FD_4BC0 | 14.4.6.6/14-66 |
| 0x45C | PCI Express ASPM Request Timer Register (PEX_ASPM_REQTMR) (RC mode only) | R/W | 0x0000_0629 | 14.4.6.7/14-66 |
| 0x478 | PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE) | R/W | All zeros | 14.4.6.8/14-67 |
| 0x47C | PCI Express Device Capabilities Update Register (PEX_DEVCAP_UPDATE) | R/W | All zeros | 14.4.6.9/14-68 |
| 0x480 | PCI Express Link Capabilities Update Register (PEX_LINKCAP_UPDATE) | R/W | 0x0000_3D41 | /14-68 |
| 0x490 | PCI Express Slot Capabilities Update Register (PEX_SLCAP_UPDATE) | R/W | 0x0000_07C0 | 14.4.6.11/14-70 |

Table 14-3. PCI Express Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--|---|--------|-------------|---------------------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| 0x4B0 | PCI Express Configuration Ready Register (PEX_CFG_READY) | R/W | All zeros | 14.4.6.12/14-71 |
| PCI Express BAR Configuration Registers (EP Mode) | | | | |
| 0x4D4 | PCI Express BAR Enable Register (PEX_BAR_ENABLE) | R/W | 0x0000_000F | 14.4.7.1/14-71 |
| 0x4D8 | PCI Express BAR Size Low Configuration Register (PEX_BAR_SIZEL) | R/W | 0xFC00_0000 | 14.4.7.2/14-72 |
| 0x4DC | Reserved | R/W | 0xFFFF_FFFF | 14.4.7.3/14-73 |
| 0x4E0 | PCI Express BAR Select Configuration Register (PEX_BAR_SEL) | R/W | 0x0000_0000 | 14.4.7.3/14-73 |
| 0x504 | PCI Express BAR Prefetch Configuration Register (PEX_BAR_PF) | R/W | 0x0000_0000 | 14.4.7.4/14-73 |
| PCI Express Extended Status and Control Register | | | | |
| 0x590 | PCI Express PME_To_Ack Timeout Register (PEX_PME_TO_ACK_TOR) | R/W | 0x00FD_4BC0 | 14.4.8.1/14-74 |
| 0x594 | PCI Express PME_To_Ack Status Register (PEX_PME_TO_ACK_SR) | w1c | All zeros | 14.4.8.2/14-75 |
| 0x5A0 | PCI Express PCI Interrupt Mask Register (PEX_SS_INTR_MASK) | R/W | 0x0000_003F | 14.4.8.3/14-76 |
| PCI Express CSB Bridge Registers | | | | |
| Global Registers | | | | |
| 0x800 | Reserved | R | 0x0110_1010 | — |
| 0x804 | Reserved | R | 0x0003_249F | — |
| 0x808 | PCI Express CSB Bridge Control register (PEX_CSB_CTRL) | R/W | 0x0000_0130 | 14.5.2.1/14-77 |
| 0x80C | Reserved | R | All zeros | — |
| 0x814 | PCI Express DMA Descriptor Timer Register (PEX_DMA_DSTMR) | R/W | All zeros | 14.5.2.2/14-78 |
| 0x818 | Reserved | R | All zeros | — |
| 0x81C | PCI Express CSB Bridge Status register (PEX_CSB_STAT) | RO | All zeros | 14.5.2.3/14-78 |
| 0x820 | Reserved | R | All zeros | — |
| PCI Express Outbound PIO Registers | | | | |
| 0x840 | PCI Express Outbound PIO Control Register (PEX_CSB_OBCTRL) | R/W | All zeros | 14.5.3.1/14-79 |
| 0x844 | PCI Express Outbound PIO Status Register (PEX_CSB_OBSTAT) | w1c | All zeros | 14.5.3.2/14-80 |
| 0x848 | Reserved | R | All zeros | — |
| PCI Express Inbound PIO Registers | | | | |

Table 14-3. PCI Express Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--|--|--------|-----------|--------------------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| 0x8E0 | PCI Express Inbound PIO Control Register (PEX_CSB_IBCTRL) | R/W | All zeros | 14.5.4.1/14-81 |
| 0x8E4 | PCI Express Inbound PIO Status Register (PEX_CSB_IBSTAT) | w1c | All zeros | 14.5.4.2/14-81 |
| 0x8E8 | Reserved | R | All zeros | — |
| PCI Express DMA Registers | | | | |
| 0x990 | Reserved | R | All zeros | — |
| 0x9A0 | PCI Express Write DMA Control Register (PEX_WDMA_CTRL) | R/W | All zeros | 14.5.5.1/14-82 |
| 0x9A4 | PCI Express Write DMA first Address Register (PEX_WDMA_ADDR) | R/W | All zeros | 14.5.5.2/14-83 |
| 0x9A8 | PCI Express Write DMA Status Register (PEX_WDMA_STAT) | w1c | All zeros | 14.5.5.3/14-83 |
| 0x9AC | Reserved | R | All zeros | — |
| 0xA40 | PCI Express Read DMA Control Register (PEX_RDMA_CTRL) | R/W | All zeros | 14.5.5.4/14-84 |
| 0xA44 | PCI Express Read DMA first Address Register (PEX_RDMA_ADDR) | R/W | All zeros | 14.5.5.5/14-85 |
| 0xA48 | PCI Express Read DMA Status Register (PEX_RDMA_STAT) | w1c | All zeros | 14.5.5.6/14-85 |
| Mailbox Registers | | | | |
| 0xB20 | PCI Express Outbound Mailbox Control Register (PEX_OMBCR) | R/W | All zeros | 14.5.6.1/14-86 |
| 0xB24 | PCI Express Outbound Mailbox Data Register (PEX_OMBDR) | R/W | All zeros | 14.5.6.2/14-87 |
| 0xB60 | PCI Express Inbound Mailbox Control Register (PEX_IMBCR) | R/W | All zeros | 14.5.6.3/14-87 |
| 0xB64 | PCI Express Inbound Mailbox Data Register (PEX_IMBDR) | R/W | All zeros | 14.5.6.4/14-88 |
| PCI Express Host Interrupts Registers | | | | |
| 0xBA0 | PCI Express Host Interrupt Enable Register (PEX_HIER) | R/W | All zeros | 14.5.7.1/14-89 |
| 0xBA4 | PCI Express Host Interrupt Status Register (PEX_HISR) | w1c | All zeros | 14.5.7.2/14-90 |
| 0xBA8 | PCI Express Host Outbound PIO Interrupt Vector Register (PEX_HOPIVR) | R/W | All zeros | 14.5.7.3/14-91 |
| 0xBC0 | PCI Express Host Inbound PIO Interrupt Vector Register (PEX_HIPIVR) | R/W | All zeros | 14.5.7.4/14-91 |
| 0xBC8 | PCI Express Host Write DMA Interrupt Vector Register (PEX_HWDIVR) | R/W | All zeros | 14.5.7.5/14-92 |
| 0xBD0 | PCI Express Host Read DMA Interrupt Vector Register (PEX_HRDIVR) | R/W | All zeros | 14.5.7.6/14-92 |
| 0xBD8 | PCI Express Host Miscellaneous Interrupt Vector Register (PEX_HMIVR) | R/W | All zeros | 14.5.7.7/14-93 |
| CSB System Interrupts Registers | | | | |
| 0xBE0 | CSB System PIO Interrupt Enable Register (PEX_CSPIER) | R/W | All zeros | 14.5.8.1/14-93 |

Table 14-3. PCI Express Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--|---|--------|-------------|----------------------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| 0xBE4 | CSB System Write DMA Interrupt Enable Register (PEX_CSWDIER) | R/W | All zeros | 14.5.8.2/14-94 |
| 0xBE8 | CSB System Read DMA Interrupt Enable Register (PEX_CSRDIER) | R/W | All zeros | 14.5.8.3/14-95 |
| 0xBEC | CSB System Miscellaneous Interrupt Enable Register (PEX_CSMIER) | R/W | 0x0000_0002 | 14.5.8.4/14-95 |
| 0xBF0 | CSB System PIO Interrupt Status Register (PEX_CSPISR) | w1c | All zeros | 14.5.8.5/14-97 |
| 0xBF4 | CSB System Write DMA Interrupt Status Register (PEX_CSWDISR) | w1c | All zeros | 14.5.8.6/14-97 |
| 0xBF8 | CSB System Read DMA Interrupt Status Register (PEX_CSRDISR) | w1c | All zeros | 14.5.8.7/14-98 |
| 0xBFC | CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR) | w1c | All zeros | 14.5.8.8/14-99 |
| Power Management Registers | | | | |
| 0xC80 | PCI Express PM Control Register (PEX_PM_CTRL) | R/W | All zeros | 14.5.9.1/14-100 |
| 0xC88 | PCI Express slot control misc register | R/W | All zeros | 14.5.9.2/14-101 |
| PCI Express Outbound Address Mapping Registers | | | | |
| 0xCA0 | PCI Express Outbound Window Attributes Register 0 (PEX_OWAR0) | R/W | All zeros | 14.5.10.1/14-102 |
| 0xCA4 | PCI Express Outbound Window Base Address Register 0 (PEX_OWBAR0) | R/W | All zeros | 14.5.10.2/14-103 |
| 0xCA8 | PCI Express Outbound Window Translation Address Register Low 0 (PEX_OWTARL0) | R/W | All zeros | 14.5.10.3/14-103 |
| 0xCAC | PCI Express Outbound Window Translation Address Register High 0 (PEX_OWTARH0) | R/W | All zeros | 14.5.10.4/14-104 |
| 0xCB0 | PCI Express Outbound Window Attributes Register 1 (PEX_OWAR1) | R/W | All zeros | 14.5.10.1/14-102 |
| 0xCB4 | PCI Express Outbound Window Base Address Register 1 (PEX_OWBAR1) | R/W | All zeros | 14.5.10.2/14-103 |
| 0xCB8 | PCI Express Outbound Window Translation Address Register Low 1 (PEX_OWTARL1) | R/W | All zeros | 14.5.10.3/14-103 |
| 0xCBC | PCI Express Outbound Window Translation Address Register High 1 (PEX_OWTARH1) | R/W | All zeros | 14.5.10.4/14-104 |
| 0xCC0 | PCI Express Outbound Window Attributes Register 2 (PEX_OWAR2) | R/W | All zeros | 14.5.10.1/14-102 |
| 0xCC4 | PCI Express Outbound Window Base Address Register 2 (PEX_OWBAR2) | R/W | All zeros | 14.5.10.2/14-103 |
| 0xCC8 | PCI Express Outbound Window Translation Address Register Low 2 (PEX_OWTARL2) | R/W | All zeros | 14.5.10.3/14-103 |

Table 14-3. PCI Express Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--|---|--------|-----------|------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| 0xCCC | PCI Express Outbound Window Translation Address Register High 2 (PEX_OWTARH2) | R/W | All zeros | 14.5.10.4/14-104 |
| 0xCD0 | PCI Express Outbound Window Attributes Register 3 (PEX_OWAR3) | R/W | All zeros | 14.5.10.1/14-102 |
| 0xCD4 | PCI Express Outbound Window Base Address Register 3 (PEX_OWBAR3) | R/W | All zeros | 14.5.10.2/14-103 |
| 0xCD8 | PCI Express Outbound Window Translation Address Register Low 3 (PEX_OWTARL3) | R/W | All zeros | 14.5.10.3/14-103 |
| 0xCDC | PCI Express Outbound Window Translation Address Register High 3 (PEX_OWTARH3) | R/W | All zeros | 14.5.10.4/14-104 |
| PCI Express EP Inbound Address Translation Registers | | | | |
| 0xDE0 | PCI Express EP Inbound Window Translation Address Register 0 (PEX_EPIWTAR0) | R/W | All zeros | 14.5.11.1/14-105 |
| 0xDE4 | PCI Express EP Inbound Window Translation Address Register 1 (PEX_EPIWTAR1) | R/W | All zeros | 14.5.11.1/14-105 |
| 0xDE8 | PCI Express EP Inbound Window Translation Address Register 2 (PEX_EPIWTAR2) | R/W | All zeros | 14.5.11.1/14-105 |
| 0xDEC | PCI Express EP Inbound Window Translation Address Register 3 (PEX_EPIWTAR3) | R/W | All zeros | 14.5.11.1/14-105 |
| PCI Express RC Inbound Address Mapping Registers | | | | |
| 0xE60 | PCI Express RC Inbound Window Attributes Register 0 (PEX_RCIWAR0) | R/W | All zeros | 14.5.12.1/14-106 |
| 0xE64 | PCI Express RC Inbound Window Translation Address Register 0 (PEX_RCIWTAR0) | R/W | All zeros | 14.5.12.2/14-107 |
| 0xE68 | PCI Express RC Inbound Window Base Address Register Low 0 (PEX_RCIWBARL0) | R/W | All zeros | 14.5.12.3/14-108 |
| 0xE6C | PCI Express RC Inbound Window Base Address Register High 0 (PEX_RCIWBARH0) | R/W | All zeros | 14.5.12.4/14-108 |
| 0xE70 | PCI Express RC Inbound Window Attributes Register 1 (PEX_RCIWAR1) | R/W | All zeros | 14.5.12.1/14-106 |
| 0xE74 | PCI Express RC Inbound Window Translation Address Register 1 (PEX_RCIWTAR1) | R/W | All zeros | 14.5.12.2/14-107 |
| 0xE78 | PCI Express RC Inbound Window Base Address Register Low 1 (PEX_RCIWBARL1) | R/W | All zeros | 14.5.12.3/14-108 |
| 0xE7C | PCI Express RC Inbound Window Base Address Register High 1 (PEX_RCIWBARH1) | R/W | All zeros | 14.5.12.4/14-108 |
| 0xE80 | PCI Express RC Inbound Window Attributes Register 2 (PEX_RCIWAR2) | R/W | All zeros | 14.5.12.1/14-106 |
| 0xE84 | PCI Express RC Inbound Window Translation Address Register 2 (PEX_RCIWTAR2) | R/W | All zeros | 14.5.12.2/14-107 |

Table 14-3. PCI Express Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|--|--|--------|-----------|------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| 0xE88 | PCI Express RC Inbound Window Base Address Register Low 2 (PEX_RCIWBARL2) | R/W | All zeros | 14.5.12.3/14-108 |
| 0xE8C | PCI Express RC Inbound Window Base Address Register High 2 (PEX_RCIWBARH2) | R/W | All zeros | 14.5.12.4/14-108 |
| 0xE90 | PCI Express RC Inbound Window Attributes Register 3 (PEX_RCIWAR3) | R/W | All zeros | 14.5.12.1/14-106 |
| 0xE94 | PCI Express RC Inbound Window Translation Address Register 3 (PEX_RCIWTAR3) | R/W | All zeros | 14.5.12.2/14-107 |
| 0xE98 | PCI Express RC Inbound Window Base Address Register Low 3 (PEX_RCIWBARL3) | R/W | All zeros | 14.5.12.3/14-108 |
| 0xE9C | PCI Express RC Inbound Window Base Address Register High 3 (PEX_RCIWBARH3) | R/W | All zeros | 14.5.12.4/14-108 |
| PCI Express Controller 2 Memory-Mapped Registers | | | | |
| 0x000–0xFFC | PCI Express Controller 2 registers Note: All registers defined for PCI Express controller 1 are also defined for PCI Express controller 2; the offsets of the PCI Express controller 2 registers are the same except they have a different block base address of 0x0_A000. | | | |

14.4 PCI Express Core Configuration Header Registers

The PCI Express core implements a standard type 0/type 1 configuration space, which consists of a 64-byte type 0 configuration space header and capability structures listed in PCI/PCI Express specification.

The various capabilities supported are as follows:

- Power management (PM)
- PCI Express (PCI_EX)
- Message signaled interrupt (MSI) (not present for RC)
- Vital product data (VPD) (not present for RC)
- Subsystem ID and subsystem vendor ID (SSID/SSVID) (optional capability only for type-1 header devices)

The supported PCI Express extended capabilities are as follows:

- Advanced error reporting
- Device serial number
- Power budgeting VC capability (including VC arbitration table for WRR-32/port arbitration table)
- Vendor specific capability (VSEC)

14.4.1 Common PCI-Compatible Configuration Header Registers

The first 64 bytes of the 256-byte PCI-compatible configuration space consists of a predefined header that every PCI device must support. The first 16 bytes of the predefined header are defined the same way for all PCI Express devices. These common registers are shown in [Figure 14-2](#). They are common to both type 0 and type 1 configuration headers.

| Reserved | | | | Address Offset (Hex) |
|------------|-------------|---------------|-----------------|-------------------------|
| Device ID | | Vendor ID | | 00 |
| Status | | Command | | 04 |
| Class Code | | | Revision ID | 08 |
| BIST | Header Type | Latency Timer | Cache Line Size | 0C |

Figure 14-2. PCI Express PCI-Compatible Configuration Header Common Registers

The remaining 48 bytes of the header may have differing layouts depending on the function of the device. Two header types apply to PCI Express. Type 0 headers, described in [Section 14.4.2, “Type 0 PCI-Compatible Configuration Header Registers,”](#) are typically used by endpoints; Type 1 headers described in [Section 14.4.3, “Type 1 PCI-Compatible Configuration Header Registers,”](#) are used by root complexes and switches/bridges.

NOTE

The registers described in this section use little-endian byte ordering. Software running on the local processor in big-endian mode must byte-swap the data. No byte swapping occurs when the registers are accessed from the PCI Express bus.

14.4.1.1 PCI Express Vendor ID Register

The vendor ID register, shown in [Figure 14-3](#), identifies the manufacturer of the device.




Figure 14-3. PCI Express Vendor ID Register

[Table 14-4](#) describes the vendor ID register fields.

Table 14-4. PCI Express Vendor ID Register Field Description

| Bits | Name | Description |
|------|-----------|--------------------|
| 15–0 | Vendor ID | 0x1957 (Freescale) |

14.4.1.2 PCI Express Device ID Register

The device ID register, shown in [Figure 14-4](#), identifies the device.




Figure 14-4. PCI Express Device ID Register

[Table 14-5](#) describes the device ID register fields.

Table 14-5. PCI Express Device ID Register Field Description

| Bits | Name | Description |
|------|-----------|---|
| 15–0 | Device ID | Device ID. This field identifies the device. 00B4 MPC8315E 00B5 MPC8315 00B6 MPC8314E 00B7MPC8314 |

14.4.1.3 PCI Express Command Register

The PCI Express command register, shown in [Figure 14-5](#), controls the ability to generate and respond to PCI Express cycles. The error control and status bits in the command and status registers control PCI-compatible error reporting. Note that PCI Express advanced error reporting is controlled by the PCI Express device control register described in [Section 14.4.4.10, “PCI Express Device Control Register,”](#) and the advance error reporting capability structure described in [Section 14.4.5.1, “PCI Express Advanced Error Reporting Capability ID Register,”](#) through [Section 14.4.5.11, “PCI Express Error Source Identification Register.”](#)




Figure 14-5. PCI Express Command Register

[Table 14-6](#) describes the bits of the command register.

Table 14-6. PCI Express Command Register Field Descriptions

| Bits | Name | Description |
|-------|-------------------|---|
| 15–11 | — | Reserved |
| 10 | Interrupt Disable | Controls the ability to generate INTx interrupt messages. 0 Enables messages 1 Disables messages Any INTx emulation interrupts already asserted by this device must be negated when this bit is set. |

Table 14-6. PCI Express Command Register Field Descriptions (continued)

| Bits | Name | Description |
|------|-----------------------|---|
| 9 | — | Reserved |
| 8 | SERR | Controls the reporting of fatal and non-fatal errors detected by the device to the root complex. 0 Disables reporting 1 Enables reporting |
| 7 | — | Reserved |
| 6 | Parity error response | Controls whether this PCI Express controller responds to parity errors. 0 Parity errors are ignored and normal operation continues. 1 Parity errors cause the appropriate bit in the PCI Express status register to be set. However, note that errors are reported based on the values set in the PCI Express error enable and detection registers. |
| 5–3 | — | Reserved |
| 2 | Bus master | Enables/disables this PCI Express device to behave as a PCI Express bus master. 0 Disables the ability to generate PCI Express accesses. 1 Enables this PCI Express controller to behave as a bus master. Clearing this bit prevent the device from issuing any memory or I/O transactions. Because MSI interrupts are effectively memory writes, clearing this bit also disables the ability of the device to issue MSI interrupts. |
| 1 | Memory space | Controls whether this PCI Express device (as a target) responds to memory accesses. 0 Device does not respond to PCI Express memory space accesses. 1 Device responds to PCI Express memory space accesses. Clearing this bit prevents the device from accepting any memory transaction. It does not affect outbound memory transactions. |
| 0 | I/O space | I/O space. This bit is hard-wired to 0. |

14.4.1.4 PCI Express Status Register

The status register, shown in [Figure 14-6](#), records status information for PCI Express events.



Figure 14-6. PCI Express Status Register

Table 14-7 describes the PCI Express status register bits.

Table 14-7. PCI Express Status Register Field Descriptions

| Bits | Name | Description |
|------|--|---|
| 15 | Detected parity error ¹ | Set when a device receives a poisoned TLP regardless of the state of bit 6 in the command register. |
| 14 | Signaled system error ¹ | Set when a device sends a ERR_FATAL or ERR_NONFATAL message and the SERR enable bit in the command register is set. |
| 13 | Received master-abort ¹ | Set when a requestor receives a completion with unsupported request completion status. |
| 12 | Received target-abort ¹ | Set when a device receives a completion with completer abort completion status. |
| 11 | Signaled target-abort ¹ | Set when a device completes a request using completer abort completion status. |
| 10–9 | — | Reserved |
| 8 | Master data parity error detected ¹ | Set by the requestor (primary side for Type1 headers) when either the requestor receives a completion marked poisoned or the requestor poisons a write request. Note that the parity error enable bit (bit 6) in the command register must be set before this bit can be set. |
| 7–5 | — | Reserved. |
| 4 | Capabilities List | All PCI Express devices are required to implement the PCI Express capability structure. |
| 3 | Interrupt Status | Set when an INTx interrupt message is pending internally to the device. Note that this bit is associated with INTx messages and not Message Signaled Interrupts. |
| 2–0 | — | Reserved. |

¹ The error control and status bits in the command and status registers control PCI-compatible error reporting. PCI Express advanced error reporting is controlled by the PCI Express device control register described in Section 14.4.4.10, “PCI Express Device Control Register,” and the advance error reporting capability structure described in Section 14.4.5.1, “PCI Express Advanced Error Reporting Capability ID Register,” through Section 14.4.5.11, “PCI Express Error Source Identification Register.”

14.4.1.5 PCI Express Revision ID Register

The revision ID register, shown in Figure 14-7, identifies the revision of the device.




Figure 14-7. PCI Express Revision ID Register

Table 14-8 describes the revision ID register fields.

Table 14-8. PCI Express Revision ID Register Field Descriptions

| Bits | Name | Description |
|------|-------------|--|
| 7–0 | Revision ID | Revision specific. The value is 0x10. |

14.4.1.6 PCI Express Class Code Register

The PCI Express class code register, shown in Figure 14-8, is composed of three single-byte fields—base class (offset 0x00B), sub-class (offset 0x00A), and programming interface (offset 0x009)—that indicate the basic functionality.




Figure 14-8. PCI Express Class Code Register

Table 14-9 describes the class code register fields.

Table 14-9. PCI Express Class Code Register Field Descriptions

| Bits | Name | Description |
|-------|-----------------------|------------------------------|
| 23–16 | Base Class | 0x0B—Processor |
| 15–8 | Sub-Class | 0x20—PowerPC |
| 7–0 | Programming Interface | 0x00—RC mode 0x00—EP mode |

14.4.1.7 PCI Express Cache Line Size Register

The cache line size register, shown in Figure 14-9, is provided for legacy compatibility (PCI 2.3); it is not used for PCI Express device functionality.




Figure 14-9. PCI Express Bus Cache Line Size Register

Table 14-10 describes the cache line size register.

Table 14-10. PCI Express Bus Cache Line Size Register Field Descriptions

| Bits | Name | Description |
|------|-----------------|---|
| 7–0 | Cache Line Size | Represents the cache line size of the processor in terms of 32-bit words (eight 32-bit words = 32 bytes). Note that for PCI Express operation this register is ignored. |

14.4.1.8 PCI Express Latency Timer Register

The latency timer register, shown in Figure 14-10, is provided for legacy compatibility (PCI 2.3); it is not used for PCI Express device functionality.




Figure 14-10. PCI Express Latency Timer Register

Table 14-11 describes the PCI Express latency timer register (PLTR).

Table 14-11. PCI Express Latency Timer Register Field Descriptions

| Bits | Name | Description |
|------|---------------|---|
| 7–0 | Latency Timer | Note that for PCI Express operation this register is ignored. |

14.4.1.9 PCI Express Header Type Register

The PCI Express header type register, shown in Figure 14-9, identifies the layout of the PCI-compatible header.




Figure 14-11. PCI Express Header Type Register

Table 14-12 describes the PCI Express header type register.

Table 14-12. PCI Express Header Type Register Field Descriptions

| Bits | Name | Description |
|------|---------------|---|
| 7 | Multifunction | Identifies whether a device supports multiple functions 0 Single-function device 1 Multiple-function device |
| 6–0 | Header Layout | 0x00 Endpoint. See Figure 14-12 for type 0 layout. 0x01 Root Complex. See Figure 14-23 for type 1 layout. All other encodings are reserved. |

14.4.1.10 PCI Express BIST Register

The BIST register is optional and reserved on the PCI Express controller.

14.4.2 Type 0 PCI-Compatible Configuration Header Registers

The type 0 header is shown in [Figure 14-12](#).

| <div style="display: inline-block; width: 15px; height: 10px; background-color: #cccccc; border: 1px solid black; margin-right: 5px;"></div> Reserved | | | | Address Offset (Hex) |
|---|-------------|---------------------|----------------------|-------------------------|
| Device ID | | Vendor ID | | 00 |
| Status | | Command | | 04 |
| Class Code | | | Revision ID | 08 |
| BIST | Header Type | Latency Timer | Cache Line Size | 0C |
| Base Address Registers | | | | 10 |
| | | | | 14 |
| | | | | 18 |
| | | | | 1C |
| | | | | 20 |
| | | | | 24 |
| | | | | 28 |
| Subsystem ID | | Subsystem Vendor ID | | 2C |
| | | | | 30 |
| | | | Capabilities Pointer | 34 |
| | | | | 38 |
| Expansion ROM Base Address | | | | 38 |
| MAX_LAT | MIN_GNT | Interrupt Pin | Interrupt Line | 3C |

Figure 14-12. PCI Express PCI-Compatible Configuration Header—Type 0

Section 14.4.1, “Common PCI-Compatible Configuration Header Registers,” describes the registers in the first 16 bytes of the header. This section describes the registers that are unique to the type 0 header beginning at offset 0x010.

14.4.2.2 PCI Express Subsystem Vendor ID Register (EP Mode Only)

The PCI Express subsystem vendor ID register, shown in [Figure 14-16](#), identifies the subsystem.




Figure 14-16. PCI Express Subsystem Vendor ID Register

[Table 14-16](#) describes the PCI Express subsystem vendor ID register fields.

Table 14-16. PCI Express Subsystem Vendor ID Register Field Descriptions

| Bits | Name | Description |
|------|---------------------|--|
| 15–0 | Subsystem Vendor ID | Subsystem Vendor ID. The value of this register can be set by programming the SSVID field of the PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE). See Section 14.4.6.8, “PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE)” . This value has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration. |

14.4.2.3 PCI Express Subsystem ID Register (EP Mode Only)

The PCI Express subsystem ID register identifies the subsystem.




Figure 14-17. PCI Express Subsystem ID Register

[Table 14-17](#) describes the PCI Express subsystem ID register fields.

Table 14-17. PCI Express Subsystem ID Register Field Descriptions

| Bits | Name | Description |
|------|--------------|--|
| 15–0 | Subsystem ID | Subsystem ID. The value of this register can be set by programming the SSID field of the PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE). See Section 14.4.6.8, “PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE)” . This value has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration. |

14.4.2.4 PCI Express Capabilities Pointer Register

The PCI Express capabilities pointer identifies additional functionality supported by the device.




Figure 14-18. PCI Express Capabilities Pointer Register

Table 14-18. PCI Express Capabilities Pointer Register Field Descriptions

| Bits | Name | Description |
|------|----------------------|---|
| 7–0 | Capabilities Pointer | The capabilities pointer provides the offset (0x44) for additional PCI-compatible registers above the common 64-byte header. Refer to Section 14.4.4, “PCI Compatible Device-Specific Configuration Space Registers.” |

14.4.2.5 PCI Express Interrupt Line Register (EP-Mode Only)

The PCI Express interrupt line register is used by device drivers and OS software to communicate interrupt line routing information. Values in this register are programmed by system software and are system-specific.




Figure 14-19. PCI Express Interrupt Line Register

Table 14-19. PCI Express Interrupt Line Register Field Descriptions

| Bits | Name | Description |
|------|----------------|--|
| 7–0 | Interrupt Line | Communicates interrupt line routing information. |

14.4.2.6 PCI Express Interrupt Pin Register

The interrupt pin register identifies the legacy interrupt (INTx) messages the device uses.

Offset 0x03D

Access: Read-only




Figure 14-20. PCI Express Interrupt Pin Register

Table 14-20. PCI Express Interrupt Pin Register Field Descriptions

| Bits | Name | Description |
|------|---------------|--|
| 7-0 | Interrupt pin | Legacy INTx message used by this device. 0x01 INTA only is supported by this device. |

14.4.2.7 PCI Express Minimum Grant Register (EP Mode Only)

This register does not apply to PCI Express. It is present for legacy purposes.

Offset 0x03E (EP mode only)

Access: Read-only




Figure 14-21. PCI Express Minimum Grant Register (MAX_GNT)

Table 14-21. PCI Express Mlnimum Grant Register Field Descriptions

| Bits | Name | Description |
|------|---------|---------------------------------|
| 7-0 | MIN_GNT | Does not apply for PCI Express. |

14.4.2.8 PCI Express Maximum Latency Register (EP Mode Only)

This register does not apply to PCI Express. It is present for legacy purposes.

Offset 0x03F (EP mode only)

Access: Read-only




Figure 14-22. PCI Express Maximum Latency Register (MAX_LAT)

Table 14-22. PCI Express Maximum Latency Register Field Descriptions

| Bits | Name | Description |
|------|---------|---------------------------------|
| 7–0 | MAX_LAT | Does not apply for PCI Express. |

14.4.3 Type 1 PCI-Compatible Configuration Header Registers

The type 1 header is shown in [Figure 14-23](#).

| Reserved | | | | Address Offset (Hex) |
|----------------------------------|------------------------|--------------------------|----------------------|-------------------------|
| Device ID | | Vendor ID | | 00 |
| Status | | Command | | 04 |
| Class Code | | | Revision ID | 08 |
| BIST | Header Type | Latency Timer | Cache Line Size | 0C |
| | | | | 10 |
| | | | | 14 |
| Secondary Latency Timer | Subordinate Bus Number | Secondary Bus Number | Primary Bus Number | 18 |
| Secondary Status | | I/O Limit | I/O Base | 1C |
| Memory Limit | | Memory Base | | 20 |
| Prefetchable Memory Limit | | Prefetchable Memory Base | | 24 |
| Prefetchable Base Upper 32 Bits | | | | 28 |
| Prefetchable Limit Upper 32 Bits | | | | 2C |
| I/O Limit Upper 16 Bits | | I/O Base Upper 16 Bits | | 30 |
| | | | Capabilities Pointer | 34 |
| Expansion ROM Base Address | | | | 38 |
| Bridge Control | Interrupt Pin | | Interrupt Line | 3C |

Figure 14-23. PCI Express PCI-Compatible Configuration Header—Type 1

[Section 14.4.1, “Common PCI-Compatible Configuration Header Registers,”](#) describes the registers in the first 16 bytes of the header. This section describes the registers that are unique to the type 1 header beginning at offset 0x010.

14.4.3.1 PCI Express Primary Bus Number Register (RC Mode Only)

The primary bus number register is shown in [Figure 14-24](#).



Figure 14-24. PCI Express Primary Bus Number Register

Table 14-23 describes the primary bus number register fields.

Table 14-23. PCI Express Primary Bus Number Register Field Descriptions

| Bits | Name | Description |
|------|--------------------|--|
| 7-0 | Primary Bus Number | Bus that is connected to the upstream interface. Note that this register is programmed during system enumeration; in RC mode this register should remain 0x00. |

14.4.3.2 PCI Express Secondary Bus Number Register (RC Mode Only)

The secondary bus number register is shown in Figure 14-25.




Figure 14-25. PCI Express Secondary Bus Number Register

Table 14-24 describes the secondary bus number register fields.

Table 14-24. PCI Express Secondary Bus Number Register Field Descriptions

| Bits | Name | Description |
|------|----------------------|---|
| 7-0 | Secondary Bus Number | Bus that is directly connected to the downstream interface. Note that this register is programmed during system enumeration; in RC mode, this register is typically programmed to 0x01. |

14.4.3.3 PCI Express Subordinate Bus Number Register (RC Mode Only)

The subordinate bus number register is shown in Figure 14-26.




Figure 14-26. PCI Express Subordinate Bus Number Register

Table 14-25 describes the subordinate bus number register fields.

Table 14-25. PCI Express Subordinate Bus Number Register Field Descriptions

| Bits | Name | Description |
|------|------------------------|---|
| 7-0 | Subordinate Bus Number | Highest bus number that is on the downstream interface. |

14.4.3.4 PCI Express Secondary Latency Timer Register (RC Mode Only)

The secondary latency timer register does not apply to PCI Express. It must be read-only and return all zeros when read.

14.4.3.5 PCI Express I/O Base Register (RC Mode Only)

Note that this device does not support inbound I/O transactions. The I/O base register is shown in [Figure 14-27](#).




Figure 14-27. PCI Express I/O Base Register

[Table 14-26](#) describes the I/O base register fields.

Table 14-26. PCI Express I/O Base Register Field Descriptions

| Bits | Name | Description |
|------|---------------------|---|
| 7–4 | I/O Start Address | Specifies bits 15:12 of the I/O space start address |
| 3–0 | Address Decode Type | Specifies the number of I/O address bits. 0x00 16-bit I/O address decode 0x01 32-bit I/O address decode All other settings reserved. |

14.4.3.6 PCI Express I/O Limit Register (RC Mode Only)

Note that this device does not support inbound I/O transactions. The I/O limit register is shown in [Figure 14-28](#).




Figure 14-28. PCI Express I/O Limit Register

Table 14-27 describes the I/O limit register fields.

Table 14-27. PCI Express I/O Limit Register Field Descriptions

| Bits | Name | Description |
|------|---------------------|---|
| 7–4 | I/O Limit Address | Specifies bits 15:12 of the I/O space ending address |
| 3–0 | Address Decode Type | Specifies the number of I/O address bits. 0x00 16-bit I/O address decode 0x01 32-bit I/O address decode All other settings reserved. |

14.4.3.7 PCI Express Secondary Status Register (RC Mode Only)

The PCI Express secondary status register is shown in Figure 14-29. Note that the errors in this register can be masked by corresponding bits in the secondary status interrupt mask register (PEX_SS_INTR_MASK) and that by default all the errors are masked. See Section 14.4.8.3, “Secondary Status Interrupt Mask Register (PEX_SS_INTR_MASK) (RC Mode Only),” for more information.

Offset 0x01E

Access: Mixed




Figure 14-29. PCI Express Secondary Status Register

Table 14-28 describes the PCI Express secondary status register fields.

Table 14-28. PCI Express Secondary Status Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 15 | DPE | Detected parity error. This bit is set when the secondary side receives a poisoned TLP regardless of the state of the parity error response bit. |
| 14 | SSE | Signaled system error. This bit is set when a device sends a ERR_FATAL or ERR_NONFATAL message if the SERR enable bit in the command register is set to enable reporting. |
| 13 | RMA | Received master abort. This bit is set when the secondary side receives an unsupported request (UR) completion. |
| 12 | RTA | Received target abort. This bit is set when the secondary side receives a completer abort (CA) completion. |
| 11 | STA | Signaled target abort. This bit is set when the secondary side issues a CA completion. |
| 10–9 | — | Reserved. |
| 8 | MDPE | Master data parity error. This bit is set when the parity error response bit is set and the secondary side requestor receives a poisoned completion or poisons a write request. If the parity error response bit is cleared, this bit is never set. |
| 7–0 | — | Reserved |

14.4.3.8 PCI Express Memory Base Register (RC Mode Only)

The memory base register is shown in [Figure 14-30](#).



Figure 14-30. PCI Express Memory Base Register

[Table 14-29](#) describes the memory base register fields.

Table 14-29. PCI Express Memory Base Register Field Descriptions

| Bits | Name | Description |
|------|-------------|--|
| 15–4 | Memory Base | Specifies bits 31:20 of the non-prefetchable memory space start address. Typically used for specifying memory-mapped I/O space. Note: Inbound posted transactions hitting into the mem base/limit range are ignored; inbound non-posted transactions hitting into the mem base/limit range results in an unsupported request response. |
| 3–0 | — | Reserved |

14.4.3.9 PCI Express Memory Limit Register (RC Mode Only)

The memory limit register is shown in [Figure 14-31](#).




Figure 14-31. PCI Express Memory Limit Register

[Table 14-30](#) describes the memory base register fields.

Table 14-30. PCI Express Memory Limit Register Field Descriptions

| Bits | Name | Description |
|------|--------------|--|
| 15–4 | Memory Limit | Specifies bits 31:20 of the non-prefetchable memory space ending address. Typically used for specifying memory-mapped I/O space. Note: Inbound posted transactions hitting into the mem base/limit range are ignored; inbound non-posted transactions hitting into the mem base/limit range result in an unsupported request response. |
| 3–0 | — | Reserved |

14.4.3.10 PCI Express Prefetchable Memory Base Register (RC Mode Only)

The prefetchable memory base register is shown in [Figure 14-32](#).



Figure 14-32. PCI Express Prefetchable Memory Base Register

[Table 14-31](#) describes the prefetchable memory base register fields.

Table 14-31. PCI Express Prefetchable Memory Base Register Field Descriptions

| Bits | Name | Description |
|------|---------------------|---|
| 15–4 | PF Memory Base | Specifies bits 31:20 of the prefetchable memory space start address. |
| 3–0 | Address Decode Type | Number of prefetchable memory address bits. 0x00 32-bit memory address decode 0x01 64-bit memory address decode All other settings reserved. |

14.4.3.11 PCI Express Prefetchable Memory Limit Register (RC Mode Only)

The PCI Express prefetchable memory limit register is shown in [Figure 14-33](#).




Figure 14-33. PCI Express Prefetchable Memory Limit Register

[Table 14-32](#) describes the prefetchable memory limit register fields.

Table 14-32. PCI Express Prefetchable Memory Limit Register Field Descriptions

| Bits | Name | Description |
|------|---------------------|---|
| 15–4 | PF Memory Limit | Specifies bits 31:20 of the prefetchable memory space ending address. |
| 3–0 | Address Decode Type | Specifies the number of prefetchable memory address bits. 0x00 32-bit memory address decode 0x01 64-bit memory address decode All other settings reserved. |

14.4.3.12 PCI Express Prefetchable Base Upper 32-Bit Register (RC Mode Only)

The PCI Express prefetchable memory base upper 32-bit register is shown in [Figure 14-34](#).




Figure 14-34. PCI Express Prefetchable Base Upper 32-Bit Register

[Table 14-33](#) describes the PCI Express prefetchable memory base upper 32-bit register fields.

Table 14-33. PCI Express Prefetchable Base Upper 32-Bit Register Field Descriptions

| Bits | Name | Description |
|------|-----------------------|--|
| 31–0 | PF Base Upper 32 Bits | Specifies bits 64:32 of the prefetchable memory space start address when the address decode type field in the prefetchable memory base register is 0x01. |

14.4.3.13 PCI Express Prefetchable Limit Upper 32-Bit Register (RC Mode Only)

The PCI Express prefetchable memory base upper 32-bit register is shown in [Figure 14-35](#).




Figure 14-35. PCI Express Prefetchable Limit Upper 32-Bit Register

[Table 14-34](#) describes the PCI Express prefetchable memory limit upper 32-bit register fields.

Table 14-34. PCI Express Prefetchable Limit Upper 32-Bit Register Field Descriptions

| Bits | Name | Description |
|------|------------------------|--|
| 31–0 | PF Limit Upper 32 Bits | Specifies bits 64–32 of the prefetchable memory space ending address when the address decode type field in the prefetchable memory limit register is 0x01. |

14.4.3.14 PCI Express I/O Base Upper 16-Bit Register (RC Mode Only)

Note that this device does not support inbound I/O transactions. The I/O base upper 16-bit register is shown in [Figure 14-36](#).




Figure 14-36. PCI Express I/O Base Upper 16-Bit Register

Table 14-35 describes the I/O base upper 16 bits register fields.

Table 14-35. PCI Express I/O Base Upper 16-Bit Register Field Descriptions

| Bits | Name | Description |
|------|------------------------|--|
| 15–0 | I/O Base Upper 16 Bits | Specifies bits 31–16 of the I/O space start address when the address decode type field in the I/O base register is 0x01. |

14.4.3.15 PCI Express I/O Limit Upper 16-Bit Register (RC Mode Only)

Note that this device does not support inbound I/O transactions. The I/O limit upper 16-bit register is shown in Figure 14-37.




Figure 14-37. PCI Express I/O Limit Upper 16-Bit Register

Table 14-36 describes the I/O limit upper 16-bit register fields.

Table 14-36. PCI Express I/O Limit Upper 16-Bit Register Field Descriptions

| Bits | Name | Description |
|------|-------------------------|--|
| 15–0 | I/O Limit Upper 16 Bits | Specifies bits 31–16 of the I/O space ending address when the address decode type field in the I/O limit register is 0x01. |

14.4.3.16 PCI Express Capabilities Pointer Register

The PCI Express capabilities pointer identifies additional functionality supported by the device.




Figure 14-38. PCI Express Capabilities Pointer Register

Table 14-37. PCI Express Capabilities Pointer Register Field Descriptions

| Bits | Name | Description |
|------|----------------------|--|
| 7–0 | Capabilities Pointer | Provides the offset (0x44) for additional PCI-compatible registers above the common 64-byte header. Refer to Section 14.4.4, “PCI Compatible Device-Specific Configuration Space Registers.” |

14.4.3.17 PCI Express Interrupt Line Register

The PCI Express interrupt line register is used by device drivers and OS software to communicate interrupt line routing information. Values in this register are programmed by system software and are system-specific.




Figure 14-39. PCI Express Interrupt Line Register

Table 14-38. PCI Express Interrupt Line Register Field Descriptions

| Bits | Name | Description |
|------|----------------|--|
| 7–0 | Interrupt Line | Communicates interrupt line routing information. |

14.4.3.18 PCI Express Interrupt Pin Register

The interrupt pin register identifies the legacy interrupt (INTx) messages of the device (or function).




Figure 14-40. PCI Express Interrupt Pin Register

Table 14-39. PCI Express Interrupt Pin Register Field Descriptions

| Bits | Name | Description |
|------|---------------|---|
| 7–0 | Interrupt Pin | Legacy INTx message used by this device. 0x01 INTA only is supported by this device. |

14.4.3.19 PCI Express Bridge Control Register (RC Mode Only)

The PCI Express bridge control register is shown in [Figure 14-41](#).




Figure 14-41. PCI Express Bridge Control Register

Table 14-40 describes the PCI Express bridge control register fields.

Table 14-40. PCI Express Bridge Control Register Field Descriptions

| Bits | Name | Description |
|------|----------|---|
| 15–7 | — | Reserved |
| 6 | Scnd_RST | Secondary bus reset |
| 5–4 | — | Reserved |
| 3 | VGA_EN | VGA enable |
| 2 | ISA_EN | ISA enable |
| 1 | SERR_EN | SERR enable. Controls the propagation of ERR_COR, ERR_NONFATAL, and ERR_FATAL responses received on the secondary side. |
| 0 | PER | Parity error response. |

14.4.4 PCI Compatible Device-Specific Configuration Space Registers

The PCI-compatible device-specific configuration space is a PCI-compatible configuration space from 0x040 to 0x0FF (just above the 64-byte PCI-compatible configuration header).

| Reserved | Address Offset (Hex) |
|--|----------------------|
| <div style="border: 1px solid black; padding: 10px; margin: 5px;"> PCI-Compatible Configuration Header (See Section 14.4.1, "Common PCI-Compatible Configuration Header Registers," for more information.) </div> | 000 03F |
| | 040 |
| Power Mgmt Capabilities | 044 |
| Next Pointer (0x4C) | 044 |
| Power Mgmt Capability ID | 044 |
| Data | 048 |
| Power Management Status & Control | 048 |
| PCI Express Capabilities | 04C |
| Next Pointer (0x70 — EP mode) (NULL — RC mode) | 04C |
| PCI Express Capability ID | 04C |
| Device Capabilities | 050 |
| Device Status | 054 |
| Device Control | 054 |
| Link Capabilities | 058 |
| Link Status | 05C |
| Link Control | 05C |
| Slot Capabilities | 060 |
| Slot Status | 064 |
| Slot Control | 064 |
| Root Control (RC mode only) | 068 |
| Root Status | 06C |
| MSI Message Control | 070 |
| Next Pointer (NULL) | 070 |
| MSI Message Capability ID | 070 |
| MSI Message Address | 074 |
| MSI Upper Message Address | 078 |
| MSI Message Data | 07C |
| | 080 |
| | 0FF |

Figure 14-42. PCI-Compatible Device-Specific Configuration Space

14.4.4.1 PCI Express Power Management Capability ID Register

The PCI Express power management capability ID register is shown in [Figure 14-43](#).




Figure 14-43. PCI Express Power Management Capability ID Register

[Table 14-41](#) describes the PCI Express power management capability ID register fields.

Table 14-41. PCI Express Power Management Capability ID Register Field Descriptions

| Bits | Name | Description |
|------|--------------------------------|-------------------------|
| 7–0 | Power Management Capability ID | Power Management = 0x01 |

14.4.4.2 PCI Express Power Management Next Capabilities Pointer Register

The PCI Express power management next capabilities pointer register is shown in [Figure 14-44](#).




Figure 14-44. PCI Express Power Management Next Capabilities Pointer

[Table 14-42](#) describes the PCI Express power management next capabilities pointer register fields.

Table 14-42. PCI Express Power Management Next Capabilities Pointer Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7–0 | — | Points to the PCI Express Capability Registers |

Table 14-44. PCI Express Power Management Status and Control Register Field Descriptions (continued)

| Bits | Name | Description |
|------|-------------|---|
| 8 | PME_EN | PME Enable |
| 7-2 | — | Reserved |
| 1-0 | Power State | Power state. Indicates the current power state of the function. 00 D0 01 D1 02 D2 03 D3 |

14.4.4.5 PCI Express Power Management Data Register

The PCI Express power management data register is shown in [Figure 14-47](#).




Figure 14-47. PCI Express Power Management Data Register

[Table 14-45](#) describes the PCI Express power management data register fields.

Table 14-45. PCI Express Power Management Data Register Field Descriptions

| Bits | Name | Description |
|------|------|-------------|
| 7-0 | Data | — |

14.4.4.6 PCI Express Capability ID Register

The PCI Express capability ID register is shown in [Figure 14-48](#).




Figure 14-48. PCI Express Capability ID Register

[Table 14-46](#) describes the PCI Express capability ID register fields.

Table 14-49. PCI Express Device Capabilities Register Field Descriptions (continued)

| Bits | Name | Description |
|------|-----------------|---|
| 4–3 | PHAN_FCT | Phantom functions supported |
| 2–0 | MAX_PL_SIZE_SUP | Maximum payload size supported. 000 = 128 bytes |

14.4.4.10 PCI Express Device Control Register

The PCI Express device control register is shown in [Figure 14-52](#).

Offset 0x054

Access: Read/write

| | 15 | 14 | 12 | 11 | 10 | 9 | 8 | 7 | 5 | 4 | 3 | 2 | 1 | 0 |
|-------|----|---------------|-----|-----|-----|-----|-----|------------------|---|----|-----|-----|------|-----|
| R | — | MAX_READ_SIZE | | NSE | APE | PFE | ETE | MAX_PAYLOAD_SIZE | | RO | URR | FER | NFER | CER |
| W | — | MAX_READ_SIZE | | NSE | APE | PFE | ETE | MAX_PAYLOAD_SIZE | | RO | URR | FER | NFER | CER |
| Reset | 0 | 0 | 1 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 |

Figure 14-52. PCI Express Device Control Register

[Table 14-50](#) describes the PCI Express device control register fields.

Table 14-50. PCI Express Device Control Register Field Descriptions

| Bits | Name | Description |
|-------|------------------|-------------------------------|
| 15 | — | Reserved |
| 14–12 | MAX_READ_SIZE | Maximum read request size |
| 11 | NSE | No snoop enable |
| 10 | APE | AUX power PM enable |
| 9 | PFE | Phantom functions enable |
| 8 | ETE | Extended tag field enable |
| 7–5 | MAX_PAYLOAD_SIZE | Maximum payload size |
| 4 | RO | Relaxed ordering |
| 3 | URR | Unsupported request reporting |
| 2 | FER | Fatal error reporting |
| 1 | NFER | Non-fatal error reporting |
| 0 | CER | Correctable error reporting |

14.4.4.11 PCI Express Device Status Register

The PCI Express device status register is shown in [Figure 14-53](#).



Figure 14-53. PCI Express Device Status Register

[Table 14-51](#) describes the PCI Express device status register fields.

Table 14-51. PCI Express Device Status Register Field Descriptions

| Bits | Name | Description |
|------|------|------------------------------|
| 15–6 | — | Reserved |
| 5 | TP | Transactions pending |
| 4 | APD | AUX power detected |
| 3 | URD | Unsupported request detected |
| 2 | FED | Fatal error detected |
| 1 | NFED | Non-fatal error detected |
| 0 | CED | Correctable error detected |

14.4.4.12 PCI Express Link Capabilities Register

The PCI Express link capabilities register is shown in [Figure 14-54](#). Note that for End Point mode, some of these fields can indirectly be set using the PCI Express Link Capabilities Update Register (PEX_LINKCAP_UPDATE). See [Section 14.4.6.10, “PCI Express Link Capabilities Update Register \(PEX_LINKCAP_UPDATE\)”](#) for more details.




Figure 14-54. PCI Express Link Capabilities Register

[Table 14-52](#) describes the PCI Express link capabilities register fields.

Table 14-52. PCI Express Link Capabilities Register Field Descriptions


| Bits | Name | Description |
|-------|-------------|---|
| 31–24 | Port Number | — |
| 23–18 | — | Reserved |
| 17–15 | L1_EX_LAT | L1 exit latency. 0b111 indicates more than 64 microseconds |
| 14–12 | L0s_EX_LAT | L0s exit latency. 0b011 indicates 256 ns to less than 512 ns |
| 11–10 | ASPM | Active state power management (ASPM) Support, L0s Entry Supported |
| 9–4 | MAX_LINK_W | Maximum link width 0b000001 x1 |
| 3–0 | MAX_LINK_SP | Maximum link speed, 0b0001 indicates 2.5 Gb/s |

14.4.4.13 PCI Express Link Control Register

The PCI Express link control register is shown in [Figure 14-55](#).

Offset 0x05C

Access: Read/Write


Figure 14-55. PCI Express Link Control Register

[Table 14-53](#) describes the PCI Express link control register fields.

Table 14-53. PCI Express Link Control Register Field Descriptions

| Bits | Name | Description |
|------|----------|--|
| 15–8 | — | Reserved |
| 7 | EXT_SYNC | Extended synch |
| 6 | CCC | Common clock configuration |
| 5 | RL | Retrain link |
| 4 | LD | Link disable |
| 3 | RCB | Read completion boundary |
| 2 | — | Reserved |
| 1–0 | ASPM_CTL | Active state power management (ASPM) control |

Table 14-55. PCI Express Slot Capabilities Register Field Descriptions


| Bits | Name | Description |
|-------|----------------------|--|
| 31–19 | Physical Slot Number | This field indicates the physical slot number attached to this Port. |
| 18–17 | — | Reserved |
| 16–15 | SPLS | Slot power limit scale. |
| 14–17 | SPLV | Slot power limit value |
| 6 | HPD | Hot plug capable |
| 5 | HPS | Hot plug surprise |
| 4 | PIP | Power indicator present |
| 3 | AIP | Attention indicator present |
| 2 | MRLSP | MRL sensor present |
| 1 | PCP | Power controller present |
| 0 | ABP | Attention button present |

14.4.4.16 PCI Express Slot Control Register

The PCI Express slot control register is shown in [Figure 14-58](#).

Offset 0x064

Access: Read/Write


Figure 14-58. PCI Express Slot Control Register

[Table 14-56](#) describes the PCI Express slot control register fields.

Table 14-56. PCI Express Slot Control Register Field Descriptions

| Bits | Name | Description |
|-------|--------|------------------------------------|
| 15–11 | — | Reserved |
| 10 | PCC | Power controller control |
| 9–8 | PIC | Power indicator control |
| 7–6 | AIC | Attention indicator control |
| 5 | HPIE | Hot plug interrupt enable |
| 4 | CCIE | Command completed interrupt enable |
| 3 | PDCE | Presence detect changed enable |
| 2 | MRLSCE | MRL sensor changed enable |

Table 14-56. PCI Express Slot Control Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---------------------------------|
| 1 | PFDE | Power fault detected enable |
| 0 | ABPE | Attention button pressed enable |

14.4.4.17 PCI Express Slot Status Register

The PCI Express slot status register is shown in [Figure 14-59](#).

Offset 0x066

Access: Mixed




Figure 14-59. PCI Express Slot Status Register

[Table 14-57](#) describes the PCI Express slot status register fields.

Table 14-57. PCI Express Slot Status Register Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 15–7 | — | Reserved |
| 6 | PDS | Presence detect state. Indicates whether a card is present in the slot. 0 Slot empty 1 Card is present |
| 5 | MRLSS | MRL sensor state 0 MRL closed 1 MRL open |
| 4 | CC | Command completed |
| 3 | PDC | Presence detect changed |
| 2 | MRLSC | MRL sensor changed |
| 1 | PFD | Power fault detected |
| 0 | ABP | Attention button pressed |

14.4.4.18 PCI Express Root Control Register (RC Mode Only)

The PCI Express root control register is shown in [Figure 14-60](#).

Offset 0x068

Access: Read/Write




Figure 14-60. PCI Express Root Control Register

Table 14-58 describes the PCI Express root control register fields.

Table 14-58. PCI Express Root Control Register Field Descriptions

| Bits | Name | Description |
|------|--------|---|
| 15–4 | — | Reserved |
| 3 | PMEIE | PME interrupt enable. |
| 2 | SEFEE | System error on fatal error enable. |
| 1 | SENFEE | System error on non-fatal error enable. |
| 0 | SECEE | System error on correctable error enable. |

14.4.4.19 PCI Express Root Status Register (RC Mode Only)

The PCI Express root status register is shown in Figure 14-61.




Figure 14-61. PCI Express Root Status Register

Table 14-59 describes the PCI Express root status register fields.

Table 14-59. PCI Express Root Status Register Field Descriptions

| Bits | Name | Description |
|-------|------------------|------------------|
| 31–18 | — | Reserved |
| 17 | PMEP | PME pending |
| 16 | PMES | PME status |
| 15–0 | PME Requestor ID | PME requestor ID |

14.4.4.20 PCI Express MSI Message Capability ID Register (EP Mode Only)

The PCI Express MSI message capability ID register is shown in Figure 14-62.




Figure 14-62. PCI Express Capability ID Register

Table 14-60 describes the PCI Express MSI message capability ID register fields.

Table 14-60. PCI Express Capability ID Register Field Descriptions

| Bits | Name | Description |
|------|---------------------------|--------------------|
| 7-0 | MSI Message Capability ID | MSI Message = 0x05 |

NOTE

The value of the Next Pointer register at offset 0x071 is 0x00 (NULL), as this is the last capability of the list.

14.4.4.21 PCI Express MSI Message Control Register (EP Mode Only)

The PCI Express MSI message control register is shown in Figure 14-63.




Figure 14-63. PCI Express MSI Message Control Register

Table 14-61 describes the PCI Express MSI message control register fields.

Table 14-61. PCI Express MSI Message Control Register Field Descriptions

| Bits | Name | Description |
|------|------|--------------------------|
| 15-8 | — | Reserved |
| 7 | 64AC | 64-bit address capable |
| 6-4 | MME | Multiple message enable |
| 3-1 | MMC | Multiple message capable |
| 0 | MSIE | MSI enable |

14.4.4.22 PCI Express MSI Message Address Register (EP Mode Only)

The PCI Express MSI message address register is shown in Figure 14-64.




Figure 14-64. PCI Express MSI Message Address Register

Table 14-62 describes the PCI Express MSI message address register fields.

Table 14-62. PCI Express MSI Message Address Register Field Descriptions

| Bits | Name | Description |
|------|-----------------|----------------------------------|
| 31–2 | Message Address | System-specified message address |
| 1–0 | — | Reserved |

14.4.4.23 PCI Express MSI Message Upper Address Register (EP Mode Only)

The PCI Express MSI message upper address register is shown in Figure 14-65.




Figure 14-65. PCI Express MSI Message Upper Address Register

Table 14-63 describes the PCI Express MSI message upper address register fields.

Table 14-63. PCI Express MSI Message Upper Address Register Field Descriptions

| Bits | Name | Description |
|------|-----------------------|--|
| 31–0 | Message Upper Address | System-specified message upper address |

14.4.4.24 PCI Express MSI Message Data Register (EP Mode Only)

The PCI Express MSI message data register is shown in Figure 14-66.




Figure 14-66. PCI Express MSI Message Data Register

Table 14-64 describes the PCI Express MSI message data register fields.

Table 14-64. PCI Express MSI Message Data Register Field Descriptions

| Bits | Name | Description |
|------|--------------|--------------------------|
| 15–0 | Message Data | System-specified message |

14.4.5 PCI Express Extended Configuration Space

| Reserved | Address Offset (Hex) | | |
|--|--|--|-----|
| PCI Compatible Configuration Header (See Section 14.4.1 , "Common PCI-Compatible Configuration Header Registers," for more information.) | 000 03F | | |
| PCI-Compatible Device-Specific Configuration Space (See Section 14.4.4 , "PCI Compatible Device-Specific Configuration Space Registers," for more information.) | 040 0FF | | |
| <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; text-align: center;">Next Capability Offset (NULL¹)/Capability Version</td> <td style="width: 50%; text-align: center;">Advanced Error Reporting Capability ID</td> </tr> </table> | Next Capability Offset (NULL ¹)/Capability Version | Advanced Error Reporting Capability ID | 100 |
| Next Capability Offset (NULL ¹)/Capability Version | Advanced Error Reporting Capability ID | | |
| Uncorrectable Error Status | 104 | | |
| Uncorrectable Error Mask | 108 | | |
| Uncorrectable Error Severity | 10C | | |
| Correctable Error Status | 110 | | |
| Correctable Error Mask | 114 | | |
| Advanced Error Capabilities and Control | 118 | | |
| Header Log | 11C 120 124 128 | | |
| Root Error Command | 12C | | |
| Root Error Status | 130 | | |
| <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; text-align: center;">Error Source ID</td> <td style="width: 50%; text-align: center;">Correctable Error Source ID</td> </tr> </table> | Error Source ID | Correctable Error Source ID | 134 |
| Error Source ID | Correctable Error Source ID | | |
| <div style="border: 1px solid black; height: 40px; width: 100%;"></div> | 138 3FF | | |
| PCI Express Controller Internal CSRs | 400 5A3 5A4 FFF | | |

Figure 14-67. PCI Express Extended Configuration Space

¹ Even though the default value of this field is not NULL, it should be considered so by the software.

14.4.5.1 PCI Express Advanced Error Reporting Capability ID Register

The PCI Express advanced error reporting capability ID register is shown in [Figure 14-68](#).




Figure 14-68. PCI Express Advanced Error Reporting Capability ID Register

Table 14-65. PCI Express Advanced Error Reporting Capability ID Register Field Descriptions

| Bits | Name | Description |
|-------|-------------------------|---|
| 13–20 | Next Capability Pointer | Note: Even though the default value of this field is not NULL, it should be considered so by the software. |
| 19–16 | Capability Version | — |
| 15–0 | Capability ID | Advanced error reporting capability. |

14.4.5.2 PCI Express Uncorrectable Error Status Register

The PCI Express uncorrectable error status register is shown in [Figure 14-69](#). When a particular bit of this status register is set, it indicates that the error has occurred.




Figure 14-69. PCI Express Uncorrectable Error Status Register

Table 14-66. PCI Express Uncorrectable Error Status Register Field Descriptions

| Bits | Name | Description |
|-------|-------|----------------------------------|
| 31–21 | — | Reserved |
| 20 | URE | Unsupported request error status |
| 19 | ECRCE | ECRC error status |

Table 14-66. PCI Express Uncorrectable Error Status Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|------------------------------------|
| 18 | MTLP | Malformed TLP status |
| 17 | R XO | Receiver overflow status |
| 16 | UC | Unexpected completion status |
| 15 | CA | Completer abort status |
| 14 | CTO | Completion timeout status |
| 13 | FCPE | Flow control protocol error status |
| 12 | PTLP | Poisoned TLP status |
| 11–5 | — | Reserved |
| 4 | DLPE | Data link protocol error status |
| 3–1 | — | Reserved |
| 0 | TE | Training error status |

14.4.5.3 PCI Express Uncorrectable Error Mask Register

The PCI Express uncorrectable error mask register is shown in [Figure 14-70](#). An error is masked if the corresponding mask bit in this register is set.




Figure 14-70. PCI Express Uncorrectable Error Mask Register

Table 14-67. PCI Express Uncorrectable Error Mask Register Field Descriptions

| Bits | Name | Description |
|-------|---------|--------------------------------|
| 31–21 | — | Reserved |
| 20 | UREM | Unsupported request error mask |
| 19 | ECRCCEM | ECRC error mask |
| 18 | MTLPM | Malformed TLP mask |
| 17 | R XOM | Receiver overflow mask |
| 16 | UCM | Unexpected completion mask |

Table 14-71. PCI Express Advanced Error Capabilities and Control Register Field Descriptions (continued)

| Bits | Name | Description |
|------|---------------------|---|
| 7 | ECRCCC | ECRC checking capable. Status bit indicates if this capability has been enabled. 0 ECRC checking capability is disabled. 1 ECRC checking capability is enabled. |
| 6 | ECRCGE | ECRC generation enable. Set this bit to enable ECRC generation. |
| 5 | ECRCGC | ECRC generation capable. Status bit indicates if this capability has been enabled. 0 ECRC generation capability is disabled. 1 ECRC generation capability is enabled. |
| 4–0 | First Error Pointer | First error pointer. Identifies the bit position of the first error reported in uncorrectable error status register, Section 14.4.5.2, “PCI Express Uncorrectable Error Status Register.” |

14.4.5.8 PCI Express Header Log Register

The PCI Express header log register is shown in [Figure 14-75](#) and its fields are described in [Table 14-72](#).

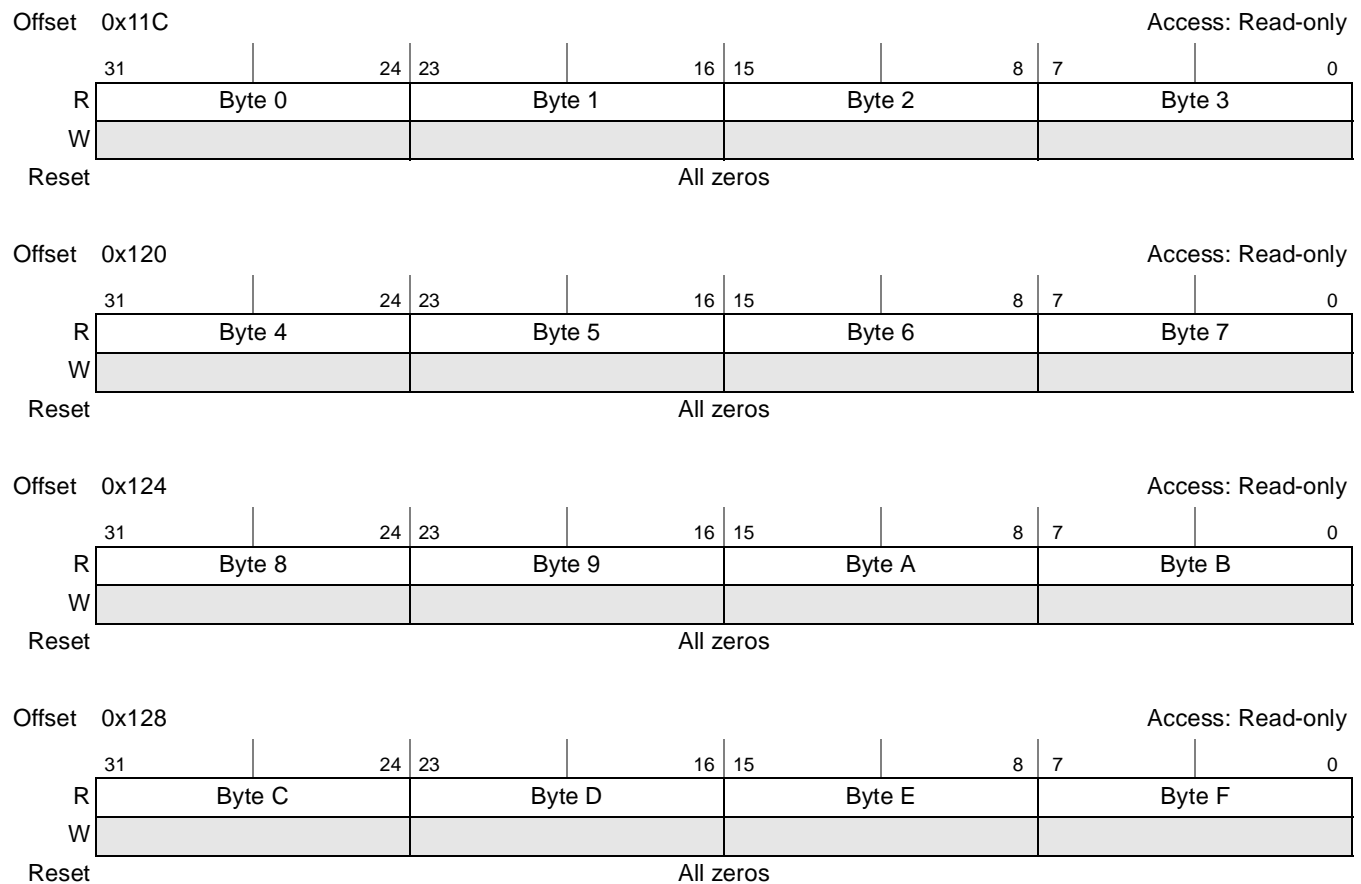



Figure 14-75. PCI Express Header Log Register

Table 14-72. PCI Express Header Log Register Field Descriptions

| Bits | Name | Description |
|-------|------------|--------------------------------------|
| 127–0 | Header Log | Header of TLP associated with error. |

14.4.5.9 PCI Express Root Error Command Register

The PCI Express root error command register is shown in [Figure 14-76](#).


Figure 14-76. PCI Express Root Error Command Register


The fields of PCI Express root error command register are described in [Table 14-73](#).

Table 14-73. PCI Express Root Error Command Register Field Descriptions

| Bits | Name | Description |
|------|-------|------------------------------------|
| 31–3 | — | Reserved |
| 2 | FERE | Fatal error reporting enable. |
| 1 | NFERE | Non-fatal error reporting enable |
| 0 | CERE | Correctable error reporting enable |

14.4.5.10 PCI Express Root Error Status Register

The PCI Express root error status register is shown in [Figure 14-77](#).


Figure 14-77. PCI Express Root Error Command Register

The fields of PCI Express root error status register are described in [Table 14-74](#).

Table 14-74. PCI Express Root Error Command Register Field Descriptions

| Bits | Name | Description |
|-------|--------|--|
| 31–27 | AEIMN | Advanced error interrupt message number. |
| 26–7 | — | Reserved |
| 6 | FEMR | Fatal error messages received. |
| 5 | NFEMR | Non-fatal error messages received. |
| 4 | FUF | First uncorrectable fatal. |
| 3 | MEFNFR | Multiple ERR_FATAL/NONFATAL received. |
| 2 | EFNFR | ERR_FATAL/NONFATAL received. |
| 1 | MECR | Multiple ERR_COR received. |
| 0 | ECR | ERR_COR received. |

14.4.5.11 PCI Express Error Source Identification Register

The Error Source Identification register shown in [Figure 14-78](#) identifies the source (Requestor ID) of first correctable and uncorrectable (Non-fatal/Fatal) errors reported in the Root Error Status register. This register is relevant for RC only.




Figure 14-78. PCI Express Error Source Identification Register

The fields of the Error Source Identification register are described in [Table 14-75](#).

Table 14-75. PCI Express Error Source Identification Register Field Descriptions

| Bits | Name | Description |
|-------|------------------------------|--|
| 31–16 | ERR_FATAL/NONFATAL Source ID | ERR_FATAL/NONFATAL Source Identification. Loaded with the Requestor ID indicated in the received ERR_FATAL or ERR_NONFATAL Message when the ERR_FATAL/NONFATAL Received register is not already set. |
| 15–0 | ERR_COR Source ID | Correctable Error Source Identification. Loaded with the Requestor ID indicated in the received ERR_COR Message when the ERR_COR Received register is not already set. |

14.4.6 PCI Express Controller Internal Control and Status Registers (CSRs)

14.4.6.1 PCI Express LTSSM State Status Register (PEX_LTSSM_STAT)

PEX_LTSSM_STAT, shown in Figure 14-79, provides details on link training status. This register is useful for debugging link training failures.




Figure 14-79. PCI Express LTSSM State Status Register (PEX_LTSSM_STAT)

The fields of the PEX_LTSSM_STAT are described in Table 14-76.

Table 14-76. PEX_LTSSM_STAT Field Descriptions

| Bits | Name | Description |
|------|-------------|---|
| 31–7 | — | Reserved |
| 6–0 | Status code | Status code. See Table 14-77 for encodings. |

Table 14-77 provides the encodings for the status code field of the PEX_LTSSM_STAT register.

Table 14-77. PEX_LTSSM_STAT Status Codes

| Status Code (Hex) | LTSSM State Description | Status Code (Hex) | LTSSM State Description |
|-------------------|-------------------------------------|-------------------|-------------------------|
| 00 | Detect quiet | 27 | TX L0s FTS; RX L0s FTS |
| 01 | Detect active (0) | 28 | L0 to L1 (0) |
| 02 | Detect active (1) | 29 | L0 to L1 (1) |
| 03 | Detect active (2) | 2A | L1 entry |
| 04 | Polling active (0) | 2B | L1 idle (0) |
| 05 | Polling active (1) | 2C | L1 idle (1) |
| 06 | Polling config (0) | 2D | L0 to L2 (0) |
| 07 | Polling config (1) | 2E | L0 to L2 (1) |
| 08 | Polling compliance | 2F | L2 entry |
| 09 | Configuration link width start (0) | 30 | L2 idle (0) |
| 0A | Configuration link width start (1) | 31 | L2 idle (1) |
| 0B | Configuration link width accept (0) | 32 | Recovery lock (0) |
| 0C | Configuration link width accept (1) | 33 | Recovery lock (1) |
| 0D | Configuration lane number wait (0) | 34 | Recovery lock (2) |
| 0E | Configuration lane number wait (1) | 35 | Recovery cfg (0) |
| 0F | Configuration lane number wait (2) | 36 | Recovery cfg (1) |
| 10 | Configuration lane number wait (3) | 37 | Recovery idle (0) |

Table 14-77. PEX_LTSSM_STAT Status Codes (continued)

| Status Code (Hex) | LTSSM State Description | Status Code (Hex) | LTSSM State Description |
|-------------------|--|-------------------|-------------------------------------|
| 11 | Configuration lane number accept | 38 | Recovery idle (1) |
| 12 | Configuration complete (0) | 39 | Recovery to configuration |
| 13 | Configuration complete (1) | 3A | Recovery cfg to configuration |
| 14 | Configuration idle (0) | 3F | L0 no training |
| 15 | Configuration idle (1) | 7F | Detect quiet EI |
| 16 | L0 | 49 | Configuration link width start—RC |
| 17 | TX L0; RX L0s entry | 4A | Configuration link width accept—RC |
| 18 | TX L0; RX L0s idle | 4B | Configuration lane number wait—RC |
| 19 | TX L0; RX L0s fast training sequence (FTS) | 4C | Configuration lane number accept—RC |
| 1A | TX L0s entry (0); RX L0 | 60 | Loopback slave active (0) |
| 1B | TX L0s entry (0); RX L0s idle | 61 | Loopback slave active (1) |
| 1C | TX L0s entry (0); RX L0s FTS | 62 | Loopback slave exit |
| 1D | TX L0s entry (1); RX L0 | 68 | Hot reset (0) |
| 1E | TX L0s entry (1); RX L0s idle | 69 | Hot reset (1) |
| 1F | TX L0s entry (1); RX L0s FTS | 6A | Hot reset (0)—RC |
| 20 | TX L0s idle; RX L0 | 6B | Hot reset (1)—RC |
| 21 | TX L0s idle; RX L0s entry | 75 | Disabled (0) |
| 22 | TX L0s idle; RX L0s idle | 71 | Disabled (1) |
| 23 | TX L0s idle; RX L0s FTS | 72 | Disabled (2) |
| 24 | TX L0s FTS; RX L0 | 73 | Disabled (3) |
| 25 | TX L0s FTS; RX L0s entry | 74 | Disabled (4) |
| 26 | TX L0s FTS; RX L0s idle | 78 | L0 to L1/L2—RC |

14.4.6.2 PCI Express N_FTS Control Register (PEX_NFTS_CTRL)

The PCI Express N_FTS Control Register, shown in [Figure 14-80](#), is used to set the N_FTS value that is advertised by the PCI Express controller during link training. It is preferable to set it before the PCI Express core internal reset is negated. If this value is changed after the link is up, the new value will take effect during the next link training. The N_FTS value is should be set according to the L0s exit latency of the Rx link of PHY. At a given time, either N_FTS or N_FTS_COM value is used based on the setting of common clock configuration bit in the configuration space.

received for some more time due to the link and processing latency involved before downstream device processes the NAK msg and stops the request. Meanwhile, these request DLLPs should not be considered as a new L1 entry request and responded with Ack DLLP. For more details refer to the PCI Express Base Specification Rev 1.0a errata, C7. ASPM & PCI-PM L1.

The fields of the PCI Express ASPM Request Timer Register are described in [Table 14-83](#).




Figure 14-85. PCI Express ASPM Request Timer Register

Table 14-83. PCI Express ASPM Request Timer Register Field Descriptions

| Bits | Name | Description |
|-------|-----------|---|
| 31–13 | — | Reserved |
| 12–0 | ASPML1TMR | ASPM L1 request timer value. This is the time-out interval after sending NAK message, before a new ASPM L1 request from a downstream device is treated as a new ASPM L1 entry request. For example, if the upstream device rejects an ASPM L1 entry request from a downstream device with ASPM NAK message, the next ASPM L1 entry request from downstream device will be entertained only after this timeout interval or only after the Rx link of the downstream port enters L0s state. This value is specified in terms of system clock cycles (CSB clock). The default value is equivalent to 9.5 μs in a 166 MHz system clock. This value can be calculated as [Time in microseconds * SYSTEM_CLK in MHz]. For example 9.5[μs]*166[MHz] = 1577 (0x629). This register is used only in RC mode. |

14.4.6.8 PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE)

The PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE) shown in [Figure 14-86](#) is used to configure Subsystem Vendor ID and Subsystem ID fields of configuration header (offset 0x2C) for End Point devices. This register has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration.



Figure 14-86. PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE)

Table 14-84. PEX_SSVID_UPDATE Field Descriptions

| Bits | Name | Description |
|-------|-------|---------------------|
| 31–16 | SSID | Subsystem ID |
| 15–0 | SSVID | Subsystem Vendor ID |

14.4.6.9 PCI Express Device Capabilities Update Register (PEX_DEVCAP_UPDATE)

The PCI Express Device Capabilities Update Register shown in [Figure 14-87](#) is used to set the values to the PCI Express Device Capabilities Register in the PCI Express configuration header (offset 0x50). It can be used when the device is configured as an End Point to make the correct device information available to the upstream device. This register has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration.




Figure 14-87. PCI Express Device Capabilities Update Register

The fields of the PCI Express Device Capabilities Update Register are described in [Table 14-85](#).

Table 14-85. PCI Express Device Capabilities Update Register Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 31–15 | — | Reserved |
| 14 | PIP | Power Indicator Present |
| 13 | AIP | Attention Indicator Present |
| 12 | APB | Attention Button Present |
| 11–9 | L1AL | Endpoint L1 Acceptable Latency. This field indicates the acceptable latency that an Endpoint can withstand due to the transition from L1 state to the L0 state. Defined encodings are: 000b Less than 1us 001b 1 μs to less than 2 μs 010b 2 μs to less than 4 μs 011b 4 μs to less than 8 μs 100b 8 μs to less than 16 μs 101b 16 μs to less than 32 μs 110b 32 μs-64 is 111b More than 64 μs |
| 8–6 | LOSAL | Endpoint L0s Acceptable Latency. This field indicates the acceptable total latency that an Endpoint can withstand due to the transition from L0s state to the L0 state. Defined encodings are: 000b Less than 64 ns 001b 64 ns to less than 128 ns 010b 128 ns to less than 256 ns 011b 256 ns to less than 512 ns 100b 512 ns to less than 1 μs 101b 1 μs to less than 2 μs 110b 2 μs-4 μs 111b More than 4 μs |
| 5 | — | Reserved (Extended Tag Field Supported). Must be set to 0b. |
| 4–3 | — | Reserved (Phantom Functions Supported). Must be set to 00b. |
| 2–0 | MPLS | Max Payload Size Supported. Must be set to 000b (128bytes) |

Table 14-86. PCI Express Link Capabilities Update Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 7–6 | ASPM | ASPM Support. Indicates the level of ASPM supported on the given PCI Express Link. Defined encodings are: 00b Reserved 01b L0s Entry Supported 10b Reserved 11b Reserved (L0s and L1 not supported by this device) |
| 5–0 | MLW | Maximum Link Width of the given PCI Express Link. Defined encodings are: 000001b x1 Other: Reserved |

14.4.6.11 PCI Express Slot Capabilities Update Register (PEX_SLCAP_UPDATE)

The PCI Express Slot Capabilities Update Register shown in [Figure 14-89](#) is used to set the values to the PCI Express Slot Capabilities Register in the PCI Express configuration header (offset 0x60). It can be used when the device is configured as an End Point to make the correct slot information available to the upstream device. This register has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration.




Figure 14-89. PCI Express Slot Capabilities Update Register

Table 14-87. PCI Express Slot Capabilities Update Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–30 | — | Reserved |
| 29–17 | PSN | Physical Slot Number. Physical Slot number attached to the port. |
| 16–15 | SPLS | Slot Power Limit Scale. Specifies the scale used for the Slot Power Limit Value. Range of Values: 00 1.0x 01 0.1x 10 0.01x 11 0.001x |
| 14–7 | SPLV | Slot power limit value. In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot. Power limit (in Watts) calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. |
| 6 | HPC | Hot plug capable |
| 5 | HPS | Hot plug surprise |
| 4 | PIP | Power indicator present |
| 3 | AIP | Attention indicator present |
| 2 | MRL | MRL sensor present |

Table 14-87. PCI Express Slot Capabilities Update Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--------------------------|
| 1 | PCP | Power controller present |
| 0 | ABP | Attention button present |

14.4.6.12 PCI Express Configuration Ready Register

The PCI Express configuration ready register, shown in [Figure 14-90](#), indicates configuration complete status to the transaction layer. The transaction layer handles configuration requests from external hosts only after the CFG_READY bit is set. All the configuration requests received from external hosts before the CFG_READY bit is set are completed with configuration request retry status (CRS). To ensure that the external host reads the correct capability advertisements during enumeration, the CFG_READY bit in this register should be set only after all relevant configuration registers are programmed.


Figure 14-90. PCI Express Configuration Ready Register (PEX_CFG_READY)

The fields of the PCI Express configuration ready register are described in [Table 14-88](#).

Table 14-88. PEX_CFG_READY Field Descriptions

| Bits | Name | Description |
|------|-----------|---|
| 31-1 | — | Reserved |
| 0 | CFG_READY | Configuration ready 1 The transaction layer accepts inbound configuration requests. 0 The transaction layer responds to all inbound configuration requests with retry (CRS) |

14.4.7 PCI Express BAR Configuration Registers (EP Mode)

The registers described in this section configure the size and the prefetchable attributes of the base address registers (BAR) in the PCI Express configuration space. The local host should program these attributes before setting the config-ready bit in the PCI Express configuration space so that the RC host reads the correct information during enumeration. These registers are used only in EP mode.

14.4.7.1 PCI Express BAR Enable Register (PEX_BAR_ENABLE)

PEX_BAR_ENABLE, shown in [Figure 14-91](#), is used to enable BARs. It supports a maximum of four BARs for the controller. To enable a given BAR, the bit corresponding to the BAR has to be set.

The fields of the PEX_BAR_SIZEL register are described in [Table 14-90](#).

Table 14-90. PEX_BAR_SIZEL Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–12 | MASK | Mask. Sets the mask for the BAR, and any bit with a value of zero is masked. When the RC does a configuration write to the BAR during the enumeration sequence, bits that are masked cannot be modified and remain zeros. All ones and zeros in this register must be consecutive. The actual size is according to the location of the least significant bit in the MASK[31:12] field, which is set. If MASK[31:m] is all ones, the size is 2^m bytes. If MASK[31:12] is all zeros, the window size is 4 Gigabytes. For example: 1111...1111 - 2^{12} , 4 Kilobytes window. 1111...1110 - 2^{13} , 8 Kilobytes window. ... 1100...0000 - 2^{30} , 1 Gigabytes window. 1000...0000 - 2^{31} , 2Gigabytes window. 0000...0000 - 4 Gigabytes window (for 64 bit BAR only). |
| 11–0 | — | Reserved. Must be zeros |

14.4.7.3 PCI Express BAR Select Configuration Register (PEX_BAR_SEL)

(PEX_BAR_SEL, shown in [Figure 14-92](#), is used to select the specific BAR for which the size is being configured by the PEX_BAR_SIZEL register. This register should be programmed before the PEX_BAR_SIZEL register is accessed, and it is used only in EP mode.




Figure 14-93. PCI Express BAR Select Configuration Register (PEX_BAR_SEL)

The fields of the PEX_BAR_SEL register are described in [Table 14-91](#).

Table 14-91. PEX_BAR_SEL Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 31–2 | — | Reserved. Must be zeros. |
| 1–0 | SEL | Select. Selects the specific BAR size to be programmed by the PEX_BAR_SIZEL and PEX_BAR_SIZEH registers. 00 PEX_BAR_SIZEL points to BAR0 in address 0x010 (Window 0, 32 bit address). 01 PEX_BAR_SIZEL points to BAR1 in address 0x014 (Window 1, 32 bit address). 10 PEX_BAR_SIZEL points to BAR2 in address 0x018 (low portion of window 2, 64 bit address). 11 PEX_BAR_SIZEL points to BAR4 in address 0x020 (low portion of window 3, 64 bit address). |

14.4.7.4 PCI Express BAR Prefetch Configuration Register (PEX_BAR_PF)

PEX_BAR_PF, shown in [Figure 14-94](#), sets the Prefetchable field in the base address registers (BAR) of the PCI Express configuration space. The local host should program this register before setting the

- Read DMA engine control
- Mailbox
- Message signaled interrupts (MSI) generation
- Events monitoring

NOTE

The registers described in this section use little-endian byte ordering. Software running on the local processor in big-endian mode must byte-swap the data. No byte swapping occurs when the registers are accessed from the PCI bus.

14.5.2 Global Registers

14.5.2.1 PCI Express CSB Bridge Control Register (PEX_CSB_CTRL)

PEX_CSB_CTRL, shown in Figure 14-98, controls the operation of the PCI Express to CSB bridge.



Figure 14-98. PCI Express CSB Bridge Control Register (PEX_CSB_CTRL)

Table 14-96 defines the bit fields of PEX_CSB_CTRL.

Table 14-96. PEX_CSB_CTRL Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–10 | — | Reserved |
| 9–8 | DSAD | Depth of descriptors array. Indicates the number of descriptors that should be placed at a contiguous addresses block. The PCI Express controller DMA uses this information to fetch each block of descriptors by a burst transaction. The implicit address of the next descriptor is the next memory location. The last descriptor in the contiguous block contains the explicit address pointer of the next set of descriptors. See Section 14.7.4, “Descriptor-Based DMA,” for detailed description. Note that for most usages this field should be programmed to 00. 00 1—Single fetch descriptor chain mode. Each descriptor explicitly contains the address of the next descriptor. 01 2—Burst fetch descriptor chain mode. The PCI Express controller will fetch two contiguous descriptors in a burst. 10 4—Burst fetch descriptor chain mode. The PCI Express controller will fetch four contiguous descriptors in a burst. 11 Reserved |

Table 14-96. PEX_CSB_CTRL Register Field Descriptions (continued)

| Bits | Name | Description |
|------|--------|--|
| 7–4 | — | Reserved |
| 3 | RDMAE | Read DMA enable. Must be set to enable the read DMA operation. |
| 2 | WDMAE | Write DMA enable. Must be set to enable the write DMA operation. |
| 1 | IBPIOE | Inbound PIO enable. Must be set to enable the PCI Express Inbound PIO operation. |
| 0 | OBPIOE | Outbound PIO enable. Must be set to enable the PCI Express outbound PIO operation. |

14.5.2.2 PCI Express DMA Descriptor Timer Register (PEX_DMA_DSTMR)

PEX_DMA_DSTMR, shown in [Figure 14-99](#), contains the timer value the DMA engine should wait for before reading the next descriptor once it encounters a descriptor that is not ready. The timer should be programmed to allow sufficient number of clocks before the DMA tries to fetch the descriptors again.




Figure 14-99. PCI Express DMA Descriptor Timer Register (PEX_DMA_DSTMR)

[Table 14-97](#) defines the bit field for PEX_DMA_DSTMR.

Table 14-97. PEX_DMA_DSTMR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 31–0 | DSRT | Descriptor ready timer. Represents the number of CSB bridge clocks that the DMA engine should wait before checking whether the next descriptor is ready when it encounters invalid descriptor. |

14.5.2.3 PCI Express CSB Bridge Status Register (PEX_CSB_STAT)

PEX_CSB_STAT, shown in [Figure 14-100](#), maintains the activity status of the DMA and PIO transactions. When a transaction is initiated by the PCI Express DMA or PIO, the corresponding pending bit is set until a response is received.




Figure 14-100. PCI Express CSB Bridge Status Register (PEX_CSB_STAT)

Table 14-98 defines the bit field for PEX_CS_B_STAT.

Table 14-98. PEX_CS_B_STAT Register Field Descriptions

| Bits | Name | Description |
|-------|---------|--|
| 31–25 | — | Reserved |
| 24 | RDMARP | Read DMA read transaction pending. Indicates whether a response is pending from the PCI Express bus to a transfer by the read DMA engine. 0 No response pending 1 Response is pending |
| 23–17 | — | Reserved |
| 16 | WDMARP | Write DMA read transaction pending. Indicates whether a response is pending from the CSB bus to a transfer from the write DMA engine. 0 No response pending 1 Response is pending |
| 15–9 | — | Reserved |
| 8 | IBPIORP | PCI Express inbound PIO read transaction pending. Indicates whether a response is pending from the CSB bus for an inbound transfer from the PCI Express bus 0 No response pending 1 Response is pending |
| 7–1 | — | Reserved |
| 0 | OBPIORP | PCI Express outbound PIO read transaction pending. Indicates whether a response is pending from the PCI Express bus for a transfer initiated on the CSB bus. 0 No response pending 1 Response is pending |

14.5.3 PCI Express Outbound PIO Registers

The registers discussed in this section control PIO outbound transactions initiated by a CSB master.

14.5.3.1 PCI Express Outbound PIO Control Register (PEX_CS_B_OBCTRL)

PEX_CS_B_OBCTRL, shown in Figure 14-101, controls the PCI Express Outbound PIO operations.




Figure 14-101. PCI Express Outbound PIO Control Register (PEX_CS_B_OBCTRL)

Table 14-100. PEX_CS_B_OBSTAT Register Field Descriptions (continued)

| Bit | Name | Description |
|-----|-------|---|
| 1 | BMPER | Bridge mapping error. Hardware sets this bit to indicate that an outbound PIO operation could not be completed successfully because a CSB bridge address mapping error has occurred. |
| 0 | BENER | Bridge enable error. Hardware sets this bit to indicate that the an outbound PIO operation could not be completed successfully because the CSB bridge outbound PIO operation was not enabled. |

14.5.4 PCI Express Inbound PIO Registers

The registers discussed in this section control PIO inbound transactions initiated by a PCI Express device.

14.5.4.1 PCI Express Inbound PIO Control Register (PEX_CS_B_IBCTRL)

PEX_CS_B_IBCTRL, shown in [Figure 14-103](#), controls the PCI Express inbound PIO operations.


Figure 14-103. PCI Express Inbound PIO Control Register (PEX_CS_B_IBCTRL)


[Table 14-101](#) defines the bit fields for PEX_CS_B_IBCTRL.

Table 14-101. PEX_CS_B_IBCTRL Register Field Descriptions

| Bit | Name | Description |
|------|------|---|
| 31–1 | — | Reserved |
| 0 | PIOE | PIO enable. Must be set to enable an inbound PIO transaction. This field controls the general enable of the PCI Express CSB bridge inbound PIO operation. |

14.5.4.2 PCI Express Inbound PIO Status Register (PEX_CS_B_IBSTAT)

PEX_CS_B_IBSTAT, shown in [Figure 14-104](#), maintains the status of PCI Express Inbound PIO operations through the PCI Express CSB bridge.


Figure 14-104. PCI Express Inbound PIO Status Register (PEX_CS_B_IBSTAT)

14.5.5.2 PCI Express Write DMA First Address Register (PEX_WDMA_ADDR)

PEX_WDMA_ADDR, shown in Figure 14-106, contains the CSB local memory address of the first descriptor. Note that the content is byte swapped from a CSB native address.




Figure 14-106. PCI Express Write DMA First Address Register (PEX_WDMA_ADDR)

Table 14-104 defines the bit fields of PEX_WDMA_ADDR.

Table 14-104. PEX_WDMA_ADDR Register Field Descriptions

| Bit | Name | Description |
|------|------|---|
| 31–0 | FDSA | First descriptor address. Indicates the address of the first descriptor on the CSB local memory (byte-swapped). |

14.5.5.3 PCI Express Write DMA Status Register (PEX_WDMA_STAT)

PEX_WDMA_STAT, shown in Figure 14-107, maintains the status of write DMA operations.




Figure 14-107. PCI Express Write DMA Status Register (PEX_WDMA_STAT)

Table 14-105 defines the bit fields of PEX_WDMA_STAT.

Table 14-105. PEX_WDMA_STAT Register Field Descriptions

| Bit | Name | Description |
|------|-------|---|
| 31–7 | — | Reserved |
| 6 | DAFER | DMA data fetch error. Hardware sets this bit to indicate an error during the data fetch operation. |
| 5 | BRER | Bridge error. Hardware sets this bit to indicate that DMA operation cannot complete successfully because of a CSB bridge error. |
| 4 | DSUER | Descriptor update error. Hardware sets this bit to indicate an error during descriptor update operation. |
| 3 | — | Reserved |
| 2 | DSFER | Descriptor fetch error. This bit is set by hardware to indicate that a descriptor read from the CSB has terminated with an error. |

Table 14-105. PEX_WDMA_STAT Register Field Descriptions (continued)

| Bit | Name | Description |
|-----|-------|---|
| 1 | DSCPL | Descriptor DMA transfer completed. Hardware sets this bit after completing the transaction for the descriptor. |
| 0 | CHCPL | DMA chain transfer completed. Hardware sets this bit after completing the transaction in all the descriptors that are currently programmed. This bit is set when DMA operation is complete and the DMA controller encounters a NULL descriptor. Note: When hardware sets this bit it is not guaranteed that the transferred data has fully reached its final destination. Software should guarantee this another way. |

14.5.5.4 PCI Express Read DMA Control Register (PEX_RDMA_CTRL)

PEX_RDMA_CTRL, shown in [Figure 14-108](#), controls the RDMA operations.




Figure 14-108. PCI Express Read DMA Control Register (PEX_RDMA_CTRL)

[Table 14-106](#) defines the bit fields of PEX_RDMA_CTRL.

Table 14-106. PEX_RDMA_CTRL Register Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 31–11 | — | Reserved |
| 10 | SNOOP | Snoop for write and read transactions to the descriptor. Controls the snooping of the e300 core on the CSB bus to a transaction initiated by the RDMA. 0 Snoop disabled. 1 Snoop enabled. |
| 9 | RLXO | Relaxed ordering for PCI Express. Indicates the relaxed ordering bit to be used for all PCI Express transactions initiated by the DMA controller. |
| 8–2 | — | Reserved |
| 1 | SUS | DMA suspend. Software sets this bit to suspend the DMA controller. |
| 0 | START | DMA start. Software can set this bit to indicate that a descriptor is ready and the DMA controller can start transmission. Hardware resets this bit when a descriptor fetch cycle is initiated. |

14.5.5.5 PCI Express Read DMA First Address Register (PEX_RDMA_ADDR)

PEX_RDMA_ADDR, shown in [Figure 14-109](#), contains the local memory address of the first descriptor. Note that the content is byte swapped from a CSB native address.




Figure 14-109. PCI Express Read DMA First Address Register (PEX_RDMA_ADDR)

[Table 14-107](#) defines the bit fields of PEX_RDMA_ADDR.

Table 14-107. PEX_RDMA_ADDR Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 31–0 | FDSA | First descriptor address. Indicates the address of the first descriptor on the CSB local memory (byte swapped). |

14.5.5.6 PCI Express Read DMA Status Register (PEX_RDMA_STAT)

PEX_RDMA_STAT, shown in [Figure 14-110](#), maintains the status of read DMA operations.




Figure 14-110. PCI Express Read DMA Status Register (PEX_RDMA_STAT)

[Table 14-108](#) defines the bit fields of PEX_RDMA_STAT.

Table 14-108. PEX_RDMA_STAT Register Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 31–8 | — | Reserved |
| 7 | CPLER | PCI Express completion error. Hardware sets this bit to indicate that DMA operation cannot complete successfully because of a PCI Express error. |
| 6 | DAFER | DMA data write error. The hardware sets this bit to indicate an error during data write operation. |
| 5 | BRER | Bridge error. Hardware sets this bit to indicate that DMA operation cannot complete successfully because of a CSB bridge error. |
| 4 | DSUER | Descriptor update error. Hardware sets this bit to indicate an error during descriptor update operation. |
| 3 | — | Reserved |
| 2 | DSFER | Descriptor fetch error. This bit is set by hardware to indicate that a descriptor read from the CSB has terminated with an error. |

Table 14-108. PEX_RDMA_STAT Register Field Descriptions (continued)

| Bits | Name | Description |
|------|-------|---|
| 1 | DSCPL | Descriptor DMA transfer completed. Hardware sets this bit after completing the transaction for the descriptor. |
| 0 | CHCPL | DMA chain transfer completed. Hardware sets this bit after completing the transaction in all the descriptors that are currently programmed. This bit is set when DMA operation is complete and the DMA controller encounters a NULL descriptor. Note: When hardware sets this bit it is not guaranteed that the transferred data has fully reached its final destination. Software should guarantee this another way. |

14.5.6 Mailbox Registers

14.5.6.1 PCI Express Outbound Mailbox Control Register (PEX_OMBCR)

PEX_OMBCR, shown in [Figure 14-111](#), controls the generation of an outbound interrupt from the CSB local host to the PCI Express and indicates that the local host has programmed the data mailbox register, and it is ready to be read. Setting the ready bit generates an interrupt to the PCI Express host if enabled. The PCI Express host should clear the ready bit after reading the data mailbox register.




Figure 14-111. PCI Express Outbound Mailbox Control Register (PEX_OMBCR)

Table 14-109 defines the bit fields of PEX_OMBCR.

Table 14-109. PEX_OMBCR Register Field Descriptions

| Bits | Name | Description |
|------|-------|---|
| 31–1 | — | Reserved |
| 0 | READY | Outbound mailbox ready. If set, indicates that mailbox has valid data to be read by the PCI Express host and generates an interrupt if enabled. |

14.5.6.2 PCI Express Outbound Mailbox Data Register (PEX_OMBDR)

PEX_OMBDR, shown in Figure 14-112, contains the data to be read by the PCI Express host when it receives an interrupt.



Figure 14-112. MPCl Express Outbound Mailbox Data Register (PEX_OMBDR)

Table 14-110 defines the bit fields of PEX_OMBDR.

Table 14-110. PEX_OMBDR Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 31–0 | MBD | Mailbox Data. Contains the data to be read by the PCI Express host upon receiving an interrupt. |

14.5.6.3 PCI Express Inbound Mailbox Control Register (PEX_IMBCR)

PEX_IMBCR, shown in Figure 14-113, controls the generation of an interrupt to the local host indicating that the PCI Express host has programmed the data mailbox register, and it is ready to be read. The CSB host clears the bit after reading out the mailbox register. Note that setting the ready bit generates an interrupt to the CSB host if enabled. The CSB host should clear the ready bit after reading the data mailbox register.




Figure 14-113. PCI Express Inbound Mailbox Control Register (PEX_IMBCR)

Table 14-111 defines the bit fields of PEX_IMBCR.

Table 14-111. PEX_IMBCR Register Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 31–1 | — | Reserved |
| 0 | READY | Inbound mailbox ready. If set, indicates that mailbox has valid data to be read by the CSB local host and generates an interrupt if enabled. |

14.5.6.4 PCI Express Inbound Mailbox Data Register (PEX_IMBDR)

PEX_IMBDR, shown in [Figure 14-114](#), contains the data to be read by the local CSB host.



Figure 14-114. PCI Express Inbound Mailbox Data Register (PEX_IMBDR)

[Table 14-112](#) defines the bit fields of PEX_IMBDR.

Table 14-112. PEX_IMBDR Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 31–0 | MBD | Mailbox data. Contains the data to be read by the CSB local host upon receiving an interrupt. |

14.5.7 PCI Express Host Interrupt Registers

This section describes the registers for generating interrupts to the PCI Express host. It consists of interrupt status registers and enable registers. Interrupts are generated only if the corresponding enable bit is set. The device supports generation of INTx and MSI interrupts. Using these registers to generate interrupts to the PCI Express host is applicable for only PCI Express EP applications.

14.5.7.1 PCI Express Host Interrupt Enable Register (PEX_HIER)

PEX_HIER, shown in Figure 14-115, enables the generation of interrupts to the PCI Express host at various events during the CBS bridge, PIO and DMA operation.




Figure 14-115. PCI Express Host Interrupt Enable Register (PEX_HIER)

Table 14-113 defines the bit fields of PEX_HIER.

Table 14-113. PEX_HIER Register Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 31–12 | — | Reserved |
| 11 | CBRIE | CSB bridge reset interrupt enable. If set, enables the generation of an interrupt upon a CSB bridge reset. |
| 10 | OMBIE | Outbound mailbox ready interrupt enable. If set, enables the generation of interrupt when the outbound mailbox control register ready bit (PEX_OMBCR[READY]) is set. |
| 9 | RDAIE | RDMA transfer aborted interrupt enable. If set, enables the generation of interrupt for every Read DMA transaction aborted. |
| 8 | RDSIE | RDMA descriptor transfer completed interrupt enable. If set, enables the generation of an interrupt when a read DMA transaction corresponding to a descriptor successfully completes. |
| 7 | RDCIE | RDMA chain descriptor transfer completed interrupt enable. If set, enables the generation of interrupt when a read DMA transaction for an end-of-descriptor successfully completes. |
| 6 | WDAIE | WDMA transfer aborted interrupt enable. If set, enables the generation of interrupt for every DMA transaction aborted. |
| 5 | WDSIE | WDMA descriptor transfer completed interrupt enable. If set, enables the generation of interrupt when DMA transactions corresponding to a descriptor successfully complete. |
| 4 | WDCIE | WDMA chain descriptor transfer completed interrupt enable. If set, enables the generation of interrupt when DMA transactions for end-of-descriptor successfully complete. |
| 3 | IPAIE | Inbound PIO transaction aborted interrupt enable. If set, enables the generation of an interrupt for every inbound PCI Express PIO transaction aborted. |
| 2 | IPCIE | Inbound PIO transaction completed interrupt enable. If set, enables the generation of an interrupt for every inbound PCI Express PIO transaction that successfully completes. |
| 1 | OPAIE | Outbound PIO transaction abort interrupt enable. If set, enables the generation of an interrupt for every outbound PIO transaction aborted. |
| 0 | OPCIE | Outbound PIO transaction completed interrupt enable. If set, enables the generation of an interrupt for every outbound PIO transaction that successfully completes. |

14.5.7.2 PCI Express Host Interrupt Status Register (PEX_HISR)

PEX_HISR, shown in [Figure 14-116](#), maintains the status for interrupts issued to the PCI Express host.

Offset 0xBA4

Access: w1c




Figure 14-116. PCI Express Host Interrupt Status Register (PEX_HISR)

[Table 14-114](#) defines the bit fields of PEX_HISR.

Table 14-114. PEX_HISR Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–12 | — | Reserved |
| 11 | CBR | CSB bridge reset. Hardware sets this bit upon a CSB bridge reset. |
| 10 | OMB | Outbound mailbox ready. Hardware sets this bit when the outbound mailbox control register ready bit (PEX_OMBCR[READY]) is set, indicating that the outbound mailbox is loaded with data to be read by the PCI-Express host. |
| 9 | RDA | RDMA transfer aborted. Hardware sets this bit when a read DMA transaction aborts. |
| 8 | RDS | RDMA descriptor transfer completed. Hardware sets this bit when a read DMA transaction corresponding to a descriptor successfully completes. |
| 7 | RDC | RDMA chain descriptor transfer completed. Hardware sets this bit when a read DMA transaction for the last descriptor in a chain descriptor successfully completes. |
| 6 | WDA | WDMA transfer aborted. Hardware sets this bit when a write DMA transaction aborts. |
| 5 | WDS | WDMA descriptor transfer completed. Hardware sets this bit when a write DMA transaction corresponding to a descriptor successfully completes. |
| 4 | WDC | WDMA chain descriptor transfer completed. Hardware sets this bit when a write DMA transaction for the last descriptor in a chain descriptor successfully completes. |
| 3 | IPA | Inbound PIO transaction aborted. Hardware sets this bit when an inbound PCI Express PIO transaction aborts. |
| 2 | IPC | Inbound PIO transaction completed. Hardware sets this bit when an inbound PCI Express PIO transaction successfully completes. |
| 1 | OPA | Outbound PIO transaction aborted. Hardware sets this bit when an outbound PIO transaction aborts. |
| 0 | OPC | Outbound PIO transaction completed. Hardware sets this bit when an outbound PIO transaction successfully completes. |

14.5.7.3 PCI Express Host Outbound PIO Interrupt Vector Register (PEX_HOPIVR)

PEX_HOPIVR, shown in [Figure 14-117](#), contains the interrupt vector for MSI interrupt generation upon an outbound PIO event. This vector is sent by the MSI interrupts to the PCI Express host if the interrupt is enabled.




Figure 14-117. PCI Express Host Outbound PIO Interrupt Vector Register (PEX_HOPIVR)

[Table 14-115](#) defines the bit fields of PEX_HOPIVR.

Table 14-115. PEX_HOPIVR Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 31–5 | — | Reserved |
| 4–0 | IVEC | Interrupt Vector. Contains the vector value for MSI. |

14.5.7.4 PCI Express Host Inbound PIO Interrupt Vector Register (PEX_HIPIVR)

PEX_HIPIVR, shown in [Figure 14-118](#), contains the interrupt vector for MSI interrupt generation upon an inbound PIO event. This vectors will be sent by the MSI interrupts to the PCI Express host if the interrupt is enabled.




Figure 14-118. PCI Express Host Inbound PIO Interrupt Vector Register (PEX_HIPIVR)

[Table 14-116](#) defines the bit fields of PEX_HIPIVR.

Table 14-116. PEX_HIPIVR Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 31–5 | — | Reserved |
| 4–0 | IVEC | Interrupt vector. Contains the vector value for MSI. |

14.5.7.5 PCI Express Host Write DMA Interrupt Vector Register (PEX_HWDIVR)

PEX_HWDIVR, shown in [Figure 14-119](#), contains the interrupt vector for an MSI interrupt issued to the PCI Express host upon WDMA events.




Figure 14-119. PCI Express Host Write DMA Interrupt Vector Register (PEX_HWDIVR)

[Table 14-117](#) defines the bit fields of PEX_HWDIVR.

Table 14-117. PEX_HWDIVR Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 31–5 | — | Reserved |
| 4–0 | IVEC | Interrupt vector. Contains the vector value for MSI. |

14.5.7.6 PCI Express Host Read DMA Interrupt Vector Register (PEX_HRDIVR)

PEX_HRDIVR, shown in [Figure 14-120](#), contains the RDMA interrupt vector for MSI interrupt issued to the PCI Express host upon RDMA events.




Figure 14-120. PCI Express Host Read DMA Interrupt Vector Register (PEX_HRDIVR)

[Table 14-118](#) defines the bit fields of PEX_HRDIVR.

Table 14-118. PEX_HRDIVR Register Field Descriptions

| Bit | Name | Description |
|------|------|--|
| 31–5 | — | Reserved |
| 4–0 | IVEC | Interrupt Vector. Contains the vector value for MSI. |

14.5.7.7 PCI Express Host Miscellaneous Interrupt Vector Register (PEX_HMIVR)

PEX_HMIVR, shown in [Figure 14-121](#), contains the interrupt vector for mailbox message for MSI interrupts issued to the PCI Express host. This includes CSB bridge reset and mailbox message.




Figure 14-121. PCI Express Host Miscellaneous Interrupt Vector Register (PEX_HMIVR)

[Table 14-119](#) defines the bit fields of PEX_HMIVR.

Table 14-119. PEX_HMIVR Register Field Descriptions

| Bit | Name | Description |
|------|------|--|
| 31–5 | — | Reserved |
| 4–0 | IVEC | Interrupt Vector. Contains the vector value for MSI upon miscellaneous events. |

14.5.8 CSB System Interrupt Registers

This section describes the registers for generating interrupts to the CSB system host (through the IPIC). It consists of interrupt status registers and enable registers. Interrupts are generated only if the corresponding enable bit is set, upon PIO, DMA and Miscellaneous events. The user has the flexibility of combining all or partial interrupt signals to generate interrupts to the CSB system host.

14.5.8.1 CSB System PIO Interrupt Enable Register (PEX_CSPIER)

PEX_CSPIER, shown in [Figure 14-122](#), controls the generation of interrupt to the CSB processor at various events during an CSB or PCI Express PIO operation.




Figure 14-122. CSB System PIO Interrupt Enable Register (PEX_CSPIER)

[Table 14-120](#) defines the bit fields of PEX_CSPIER.

Table 14-120. PEX_CSPIER Register Field Descriptions

| Bit | Name | Description |
|-------|-------|---|
| 31–25 | — | Reserved |
| 24 | IPAIE | Inbound PIO transaction aborted interrupt enable. If set, enables the generation of an interrupt for every inbound PCI Express PIO transaction aborted. |

Table 14-120. PEX_CSPIER Register Field Descriptions (continued)

| Bit | Name | Description |
|-------|-------|---|
| 23–17 | — | Reserved |
| 16 | IPCIE | Inbound PIO transaction completed interrupt enable. If set, enables the generation of an interrupt for every inbound PCI Express PIO transaction that successfully completes. |
| 15–9 | — | Reserved |
| 8 | OPAIE | Outbound PIO transaction abort interrupt enable. If set, enables the generation of an interrupt for every outbound PIO transaction aborted. |
| 7–1 | — | Reserved |
| 0 | OPCIE | Outbound PIO transaction completed interrupt enable. If set, enables the generation of an interrupt for every outbound PIO transaction that successfully completes. |

14.5.8.2 CSB System Write DMA Interrupt Enable Register (PEX_CSWDIER)

PEX_CSWDIER, shown in [Figure 14-123](#), controls the generation of interrupt to the CSB processor at various events during a WDMA operation.




Figure 14-123. CSB System Write DMA Interrupt Enable Register (PEX_CSWDIER)

[Table 14-121](#) defines the bit fields of PEX_CSWDIER.

Table 14-121. PEX_CSWDIER Register Field Descriptions

| Bit | Name | Description |
|-------|-------|---|
| 31–17 | — | Reserved |
| 16 | WDAIE | WDMA transfer aborted interrupt enable. If set, enables the generation of interrupt for every DMA transaction aborted. |
| 15–9 | — | Reserved |
| 8 | WDSIE | WDMA descriptor transfer completed interrupt enable. If set, enables the generation of interrupt when DMA transactions corresponding to a descriptor complete successfully. |
| 7–1 | — | Reserved |
| 0 | WDCIE | WDMA chain descriptor transfer completed interrupt enable. If set, enables the generation of interrupt when DMA transactions for end-of-descriptor complete successfully. |

14.5.8.3 CSB System Read DMA Interrupt Enable Register (PEX_CSRDIER)

PEX_CSRDIER, shown in [Figure 14-124](#), controls the generation of interrupt to the CSB system host at various events during a RDMA operation.




Figure 14-124. CSB System Read DMA Interrupt Enable Register (PEX_CSRDIER)

[Table 14-122](#) defines the bit fields of PEX_CSRDIER.

Table 14-122. PEX_CSRDIER Register Field Descriptions

| Bit | Name | Description |
|-------|-------|---|
| 31–17 | — | Reserved |
| 16 | RDAIE | RDMA transfer aborted interrupt enable. If set, enables the generation of interrupt for every Read DMA transaction aborted. |
| 15–9 | — | Reserved. |
| 8 | RDSIE | RDMA descriptor transfer completed interrupt enable. If set, enables the generation of an interrupt when a Read DMA transactions corresponding to a descriptor complete successfully. |
| 7–1 | — | Reserved. |
| 0 | RDCIE | RDMA chain descriptor transfer completed interrupt enable. If set, enables the generation of interrupt when a Read DMA transactions for end-of-descriptor complete successfully. |

14.5.8.4 CSB System Miscellaneous Interrupt Enable Register (PEX_CSMIER)

PEX_CSMIER, shown in [Figure 14-125](#), controls the generation of interrupt to the CSB processor at various events.




Figure 14-125. CSB System Miscellaneous Interrupt Enable Register (PEX_CSMIER)

Table 14-123 defines the bit fields of PEX_CSMIER.

Table 14-123. PEX_CSMIER Register Field Descriptions

| Bit | Name | Description |
|-------|---------|---|
| 31–30 | — | Reserved |
| 29 | ERRDIE | Error detected interrupt enable. If set, enables the generation of an interrupt when a PCI Express event occurs. The PCI Express event is reported by the Secondary status register (PCI Express Secondary Status Register) at address 0x901E. Note that a secondary mask exist by the PCI Express PCI Interrupt Mask Register (PEX_SS_INTR_MASK) at address 0x95A0. Valid only for RC. |
| 28–21 | — | Reserved |
| 20 | PSCHIE | Power state change interrupt enable. If set, enables the generation of an interrupt when a device power state change occurs. |
| 19–18 | — | Reserved |
| 17 | PMETOIE | PME to Ack timeout event interrupt enable. If set, enables the generation of an interrupt when a PME to Ack timeout occurs. This bit is valid only in RC mode. |
| 16 | — | Reserved |
| 15 | ROFIE | Receive overflow error interrupt enable. If set, enables the generation of an interrupt when PCI Express reports receive overflow error. |
| 14 | ECRCIE | ECRC error interrupt enable. If set, enables the generation of an interrupt when a TLP is received by PCI Express and it fails ECRC check. |
| 13 | PTLPIE | Poisoned TLP interrupt enable. If set, enables the generation of an interrupt when a poisoned TLP is received by PCI Express. |
| 12 | MFTLPIE | Malformed TLP interrupt enable. If set, enables the generation of an interrupt when a malformed TLP is received by PCI Express |
| 11 | CTOIE | Completion timeout interrupt enable. If set, enables the generation of an interrupt when PCI Express completion timeout occurs. |
| 10 | CAIE | Completer abort interrupt enable. If set, enables the generation of an interrupt when PCI Express completion Abort is received. |
| 9 | URIE | Unsupported request interrupt enable. If set, enables the generation of an interrupt when an unsupported request is received. |
| 8 | INTDIE | PCI Express INTD interrupt enable. If set, enables the generation of an interrupt when an INTD interrupt is received on the PCI Express link. Valid for RC applications only. |
| 7 | INTCIE | PCI Express INTC interrupt enable. If set, enables the generation of an interrupt when an INTC interrupt is received on the PCI Express link. Valid for RC applications only. |
| 6 | INTBIE | PCI Express INTB interrupt enable. If set, enables the generation of an interrupt when an INTB interrupt is received on the PCI Express link. Valid for RC applications only. |
| 5 | INTAIE | PCI Express INTA interrupt enable. If set, enables the generation of an interrupt when an INTA interrupt is received on the PCI Express link. Valid for RC applications only. |
| 4–2 | — | Reserved |
| 1 | RSTIE | PCI Express reset interrupt enable. If set, enables the generation of an interrupt when the PCI Express is reset, for example, in the case of link down. |
| 0 | IMBIE | Inbound mailbox ready interrupt enable. If set, enables the generation of interrupt whenever the inbound mailbox control register ready bit (PEX_IMBCR[READY]) is set. |

Table 14-125 defines the bit fields of PEX_CSWDISR.

Table 14-125. PEX_CSWDISR Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–17 | — | Reserved |
| 16 | WDA | WDMA transfer aborted. Indicates that a write DMA transaction was aborted. |
| 15–9 | — | Reserved |
| 8 | WDS | WDMA descriptor transfer completed. Indicates that a write DMA transaction corresponding to a descriptor successfully completed. |
| 7–1 | — | Reserved |
| 0 | WDC | WDMA chain descriptor transfer completed. Indicates that a write DMA transaction corresponding to the last descriptor in a chain successfully completed. |

14.5.8.7 CSB System Read DMA Interrupt Status Register (PEX_CSRDISR)

PEX_CSRDISR, shown in Figure 14-128, maintains the status for interrupts issued to the CSB host related to the RDMA operation.




Figure 14-128. CSB System Read DMA Interrupt Status Register (PEX_CSRDISR)

Table 14-126 defines the bit fields of PEX_CSRDISR.

Table 14-126. PEX_CSRDISR Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–17 | — | Reserved |
| 16 | RDA | RDMA transfer aborted. Indicates that a Read DMA transaction was aborted. |
| 15–9 | — | Reserved |
| 8 | RDS | RDMA descriptor transfer completed. Indicates that a read DMA transaction corresponding to a descriptor successfully completed. |
| 7–1 | — | Reserved |
| 0 | RDC | RDMA chain descriptor transfer completed. Indicates that a read DMA transaction corresponding to the last descriptor in a chain successfully completed. |

14.5.8.8 CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR)

PEX_CSMISR, shown in Figure 14-129, maintains the status for interrupt issued to the CSB.




Figure 14-129. CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR)

Table 14-127 defines the bit fields of PEX_CSMISR.

Table 14-127. PEX_CSMISR Register Field Descriptions

| Bits | Name | Description |
|-------|--------|---|
| 31–30 | — | Reserved |
| 29 | ERRD | Error detected. Indicates that a PCI Express event occurred. The PCI Express event is reported by the secondary status register (PCI Express secondary status register) at address 0x901E. Note that there is a secondary mask, the PCI Express interrupt mask register (PEX_SS_INTR_MASK), at address 0x95A0. Valid only for RC. This bit must be cleared after the interrupt is serviced and after the associated status registers in the PCI Express controller causing the interrupt are cleared. |
| 28–21 | — | Reserved |
| 20 | PSCH | Power state change. Indicates that a device power state change occurred. Change in Device power state (D state) for function-0. D-state can transition between the supported values of D0, D1, D2 and D3-hot. PM software changes D-state with a configuration write to PMCSR register in PM capability of the PCI Express. The new D-state is available in the corresponding field of the PMCSR register. This bit must be cleared after the interrupt is serviced. |
| 19 | EXL2L3 | Exited L2/L3 state. Indication that the L2/L3 ready state has been exited and the current link state is L0. Traffic can be re-started on the link. This bit is set when the link switches from the L2/L3 ready to L0 state in response to an Exit L2 command. After issuing this command, the user must wait until the exited EXL2L3 bit is set before initiating traffic. This bit is valid only in RC mode. Setting this bit causes an interrupt to be sent to CSB side. The bit must be cleared after the interrupt is serviced. |
| 18 | ENL2L3 | Entered L2/L3 state. Indication to the Power manager that it is safe to switch off power to the downstream device 100nsec after this bit is set. It is set when the PCI Express controller enters L2/L3 ready state. This bit is valid only in RC mode. Setting this bit causes an interrupt to be sent to the CSB side. The bit must be cleared after the interrupt is serviced. |
| 17 | PMETMO | PME Turn Off Ack timeout event. Indication to power manager software that it is safe to switch off power to the downstream device. It is set when the PCI Express controller detects that the timeout interval for receiving a PME_To_Ack message from the downstream device has expired. This bit is valid only in RC mode. This bit must be cleared after the interrupt is serviced. |

Table 14-127. PEX_CSMISR Register Field Descriptions (continued)

| Bits | Name | Description |
|------|--------|--|
| 16 | RPMETO | Received PME Turn off message. Notifies that main power to the device is to be removed. After this notification is received, Uplink must not try to transmit TLPs or initiate PME requests because the power may be switched off. After this message is received, the user should indicate the readiness to lose power by setting the Initiate L2/L3 entry bit. This bit is valid only in EP mode. Setting this bit causes an interrupt to be sent to the CSB side. The bit must be cleared after the interrupt is serviced. |
| 15 | ROF | Receive overflow error. Indicates that the PCI Express reported a receive overflow error. |
| 14 | ECRC | ECRC error. Indicates that a TLP received by the PCI Express failed the ECRC check. |
| 13 | PTLP | Poisoned TLP. Indicates that a poisoned TLP was received by the PCI Express. |
| 12 | MFTLP | Malformed TLP. Indicates that a malformed TLP was received by the PCI Express |
| 11 | CTO | Completion timeout. Indicates that a PCI Express completion timeout occurred. |
| 10 | CA | Completer abort. Indicates that a PCI Express completion Abort was received. |
| 9 | UR | Unsupported request. Indicates that an unsupported request was received. |
| 8 | INTD | PCI Express INTD. Indicates that an INTD interrupt was received on the PCI Express link. Valid for RC applications only. |
| 7 | INTC | PCI Express INTC. Indicates that an INTC interrupt was received on the PCI Express link. Valid for RC applications only. |
| 6 | INTB | PCI Express INTB. Indicates that an INTB interrupt was received on the PCI Express link. Valid for RC applications only. |
| 5 | INTA | PCI Express INTA. Indicates that an INTA interrupt was received on the PCI Express link. Valid for RC applications only. |
| 4–2 | — | Reserved |
| 1 | RST | PCI Express reset. Indicates that a PCI Express reset or link down occurred. |
| 0 | IMB | Inbound mailbox ready. Indicates that the inbound mailbox control register ready bit (PEX_IMBCR[READY]) was set and the CSB host can read the mailbox data. |

14.5.9 PCI Express Power Management Registers

14.5.9.1 PCI Express Power Management Control Register (PEX_PM_CTRL)

This PCI Express PM Control Register shown in [Figure 14-130](#), is used to control the link power management by the PCI Express controller.



Figure 14-130. PCI Express PM Control Register (PEX_PM_CTRL)

[Table 14-128](#) defines the bit fields of the PEX_PM_CTRL.

Table 14-128. PEX_PM_CTRL Register Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 31–4 | — | Reserved. Must be zeros. |
| 3 | PMESS | PME Status Set. Set the PME Status bit in the PCI Express Power Management Status and Control Register of the configuration space (Offset 0x048). In EP mode, this also causes the PCI Express controller to send PM_PME message. PME message transmission is not supported in RC mode, but the PME status bit can still be set. PM PME message can be used to request for a device power state change. The message will be transmitted only if PME is enabled and if PME Turn off message has not been received by the endpoint. This field is Write only. Read always returns zero. |
| 2 | EXL2 | Exit L2. This bit is valid only in RC mode and instructs the PCI Express controller to exit from L2/L3 ready state and move to L0 active state so that traffic can be re-started on the PCI-Express link. This bit can be set under the following conditions: the downstream device was shut down by the power manager (through L2/L3 protocol) and later has its power restored, whereas the upstream device (CI Express controller as RC) was in L2/L3 ready state (with power and clocks available). This field is Write only. Read always returns zero. |
| 1 | PMETO | Send PME Turn off message. This bit is valid only in RC mode and instructs the PCI Express controller to send PME_Turn_Off message to downstream devices. After setting this bit, the user must not try to transmit TLPs or initiate PME requests as the power may be switched off. This field is Write only. Read always returns zero. |
| 0 | IL2L3 | Initiate L2/L3 entry. This bit is valid only in EP mode and instructs the PCI Express controller to transition to L2/L3 ready state. This bit has to be asserted only after preparing for power removal. After setting this bit, the user must not try to transmit TLPs or initiate PME requests as the power may be switched off. This field is Write only. Read always returns zero. |

14.5.9.2 PCI Express Slot Control Misc Register

The PCI Express slot control misc register, shown in [Figure 14-131](#), is used to control the PCI Express slot capability function.



Figure 14-131. PCI Express Slot Control Misc Register

Table 14-129 describes the PCI Express slot control misc register fields.

Table 14-129. PCI Express Slot Control Misc Register Field Descriptions

| Bits | Name | Description |
|------|-------------------|---|
| 0 | Slot implemented | When set, this bit indicates that the PCI Express link associated with this port is connected to a slot. |
| 1 | Slot clock config | This bit indicates that the component uses the same physical reference clock that the platform provides on the connector. If the device uses an independent clock regardless of the presence of a reference on the connector, this bit must be cleared. |
| 2–31 | — | Reserved |

14.5.10 PCI Express Outbound Address Mapping Registers

The registers discussed in this section control the outbound transactions attributes and address mapping from the CSB domain to the PCI Express domain. These registers are used in both RC and EP modes, and serve both PIO and DMA transactions.

14.5.10.1 PCI Express Outbound Window Attributes Register *n* (PEX_OWAR0–PEX_OWAR3)

PEX_OWAR0–PEX_OWAR3, shown in Figure 14-132, sets the attributes for the respective address window defined by the base and translation address registers for mapping of addresses related to CSB outbound transactions to PCI Express addresses.

Offset 0xCA0, 0xCB0, 0xCC0, 0xCD0

Access: Read/Write




Figure 14-132. PCI Express Outbound Window Attributes Register *n* (PEX_OWAR0–PEX_OWAR3)

Table 14-130 defines the bit fields of the PEX_OWAR0–PEX_OWAR3.

Table 14-130. PEX_OWAR0–PEX_OWAR3 Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–12 | SIZE | CSB window size. Indicates the size of window in bytes. The actual size is a concatenation of the SIZE field as most significant bits and 12 zeros as least significant bits {SIZE[31:12], 0x000}. |
| 11–8 | — | Reserved. Must be zeros. |
| 7–5 | TC | Traffic class. Indicates the traffic class of the packet. Applicable only if user wants to send traffic using multiple TC but single VC. |

Table 14-130. PEX_OWAR0–PEX_OWAR3 Register Field Descriptions (continued)


| Bits | Name | Description |
|------|------|---|
| 4 | NSNP | No snoop enable. When this bit and the PCI Express device control register [Enable No Snoop] bit are set, the No Snoop bit for the packet is enabled. This attribute is not applicable and must be cleared for configuration requests, I/O requests, and memory requests that are Message Signaled Interrupts. 0 PCI Express TLP snoop enabled 1 PCI Express TLP snoop disabled |
| 3 | RLXO | Relax ordering enable. When this bit and the PCI Express device control register [Enable Relaxed] bit are set, this bit enables the relaxed ordering bit for the packet. This applies only to memory transactions. |
| 2–1 | TYPE | Window type. Indicates the type to which CSB transactions to the window address are mapped. 00 CFG 01 I/O 10 Memory 11 Reserved |
| 0 | EN | Enable. Must be set to enable this window. |

14.5.10.2 PCI Express Outbound Window Base Address Register *n* (PEX_OWBAR0–PEX_OWBAR3)

PEX_OWBAR0–PEX_OWBAR3, shown in [Figure 14-133](#), contains the base address of the CSB window for mapping to a PCI Express address. Note that the CSB base address must be aligned to 1 Kbyte. Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an outbound window, or where an outbound translation window points back into an inbound window, are not allowed.

Offset 0xCA4, 0xCB4, 0xCC4, 0xCD4

Access: Read/Write


Figure 14-133. PCI Express Outbound Window Base Address Register *n* (PEX_OWBAR0–PEX_OWBAR3)

[Table 14-131](#) defines the bit fields of the PEX_OWBAR0–PEX_OWBAR3.

Table 14-131. PEX_OWBAR*n* Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–12 | BA | Base address. The CSB window base address. Represents the CSB-based address for the window. The actual address is a concatenation of the BAR field as most significant bits and 12 zeros as least significant bits {BA[31:12], 0x000}. |
| 11–0 | — | Reserved. Must be zeros. |

14.5.10.3 PCI Express Outbound Window Translation Address Register Low *n* (PEX_OWTARL0–PEX_OWTARL3)

PEX_OWTARL0–PEX_OWTARL3, shown in [Figure 14-134](#), contain the lower base address of the PCI Express domain corresponding to this window. When this window is enabled and a CSB-based

transaction hits its base address register, the address is translated to a PCI Express address, according to the PEX_OWTARL n and PEX_OWTARH n registers.




Figure 14-134. PCI Express Outbound Window Translation Address Register Low n (PEX_OWTARL0–PEX_OWTARL3)

Table 14-132 defines the bit fields of PEX_OWTARL n .

Table 14-132. PEX_OWTARL n Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–12 | TAL | Translation address low. The lower portion of the PCI Express address base. The actual address is a concatenation of the TA field as most significant bits and 12 zeros as least significant bits {TAL[31:12], 0x000}. The complete 64 bits address on the PCI Express bus is built of {PEX_OWTARH[TAH], PEX_OWTARL[TAL], 0x000}. |
| 11–0 | — | Reserved. |

14.5.10.4 PCI Express Outbound Window Translation Address Register High n (PEX_OWTARH0–PEX_OWTARH3)

PEX_OWTARH0–PEX_OWTARH3, shown in Figure 14-135, contains the higher base address of the PCI Express domain corresponding to this window. When this window is enabled and a CSB based transaction hits its base address register, the address is translated to a PCI Express address according to the PEX_OWTARL n and PEX_OWTARH n registers. This register should be used in 64-bit addressing. Otherwise it should contain all zeros.




Figure 14-135. PCI Express Outbound Window Translation Address Register High n (PEX_OWTARH0–PEX_OWTARH3)

Table 14-133 defines the bit fields of PEX_OWTARH n .

Table 14-133. PEX_OWTARH n Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 31–0 | TAH | Translation address high. Higher portion of the PCI Express address base ([63:32]). The complete 64-bit address on the PCI Express bus is built of {PEX_OWTARH[TAH], PEX_OWTARL[TAL], 0x000}. |

14.5.11 PCI Express EP Inbound Address Translation Registers

The following registers are used as the base address for the CSB domain to translate the address of an inbound transaction from the PCI Express domain. They are valid only in End Point (EP) mode and operate in conjunction with the base address registers in the PCI Express configuration space.

When a PCI Express inbound transaction hits a valid address window defined by the PCI Express configuration space base address registers and the respective translation register is enabled, the incoming address is translated to a CSB domain address and the transaction is forwarded to the CSB.

The correspondence between the PCI Express configuration space base address registers and the respective translation address registers is shown in [Table 14-134](#).

Table 14-134. EP Inbound Base and Translation Address Registers Correspondence

| Base Address Register (Configuration Space) | BAR Name | Type | TAR Address | TAR Name |
|---|----------|--|-------------|--------------|
| 0x010 | BAR0 | Window 0, 32-bit address | 0xDE0 | PEX_EPIWTAR0 |
| 0x014 | BAR1 | Window 1, 32-bit address | 0xDE4 | PEX_EPIWTAR1 |
| 0x018 | BAR2 | Window 2, 64-bit address, low portion | 0xDE8 | PEX_EPIWTAR2 |
| 0x01C | BAR3 | Window 2, 64-bit address, high portion | | |
| 0x020 | BAR4 | Window 3, 64-bit address, low portion | 0xDEC | PEX_EPIWTAR3 |
| 0x024 | BAR5 | Window 3, 64-bit address, high portion | | |

14.5.11.1 PCI Express EP Inbound Window Translation Address Register *n* (PEX_EPIWTAR0–PEX_EPIWTAR3)

PEX_EPIWTAR0–PEX_EPIWTAR3, shown in [Figure 14-136](#), contain the CSB address to be mapped for a PCI Express inbound transaction hitting the respective BAR window. Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an outbound window, or where an outbound translation window points back into an inbound window, are not allowed.




Figure 14-136. PCI Express EP Inbound Window Translation Address Register *n* (PEX_EPIWTAR0–PEX_EPIWTAR3)

Table 14-135 defines the bit fields of PEX_EPIWTAR n .

Table 14-135. PEX_EPIWTAR n Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–12 | TA | Translation address. Contains the CSB base address to be mapped for a PCI Express inbound transaction hitting the respective BAR window. The actual address is a concatenation of the TA field as most significant bits and 12 zeros as least significant bits {TA[31:12], 0x000}. |
| 11–1 | — | Reserved |
| 0 | EN | Enable. If set, indicates that the address mapping window is enabled. |

14.5.12 PCI Express RC Inbound Address Mapping Registers

The registers discussed in this section control the inbound transactions attributes and address mapping from the PCI Express bus to the CSB domain. These registers are used only in RC mode, and they serve inbound transactions. When a PCI Express inbound transaction hits a valid address window defined by these registers and the respective translation register is enabled, the incoming address is translated to a CSB domain address and the transaction is forwarded to the CSB.

14.5.12.1 PCI Express RC Inbound Window Attributes Register n (PEX_RCIWAR0 – PEX_RCIWAR3)

PEX_RCIWAR0–PEX_RCIWAR3, shown in Figure 14-137, controls the mapping of a PCI Express inbound PIO transaction to a CSB transaction. This register is valid only in RC mode.

Offset 0xE60, 0xE70, 0xE80, 0xE90

Access: Read/Write




Figure 14-137. PCI Express RC Inbound Window Attributes Register n (PEX_RCIWAR0–PEX_RCIWAR3)

Table 14-136 defines the bit fields of PEX_RCIWAR n .

Table 14-136. PEX_RCIWAR n Register Field Descriptions


| Bits | Name | Description |
|-------|------|--|
| 31–28 | — | Reserved |
| 27–12 | SIZE | PCI Express window size. Indicates the size of the window in bytes. The actual size is a concatenation of the SIZE field as most significant bits and 12 zeros as least significant bits {SIZE[27:12], 0x000}. |
| 11–5 | — | Reserved. Must be zeros. |

Table 14-136. PEX_RCIWAR n Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 4 | NSNP | No snoop. If the no-snoop override enable bit in this register (NSOV) is set, this bit defines the snooping behavior on the CSB domain for the inbound packet hitting this window. This applies only to memory transactions. 0 CSB snoop enabled 1 CSB snoop disabled |
| 3 | NSOV | No-snoop override. If set, the No-Snoop attribute in the packet is overridden by the No Snoop bit in this register (NSNP). Otherwise, the No-Snoop bit in the inbound packet defines the snooping behavior. 0 Snoop behavior is defined by the inbound packet 1 Snoop behavior is defined by the NSNP field |
| 2–1 | TYPE | Type. Indicates the type of the window to which the PCI Express transactions are mapped. 00 Reserved 01 Reserved 10 Prefetchable memory. Inbound read transactions are optimized for CSB performance. The address and size of the actual memory read transaction may differ from those of the original PCI Express packet, aligning the first and last segments of the data read from the memory to cache line boundaries. 11 Non-prefetchable memory. Inbound read transactions from the PCI Express bus access the exact address and size of the memory. This mode is not optimized for CSB performance. |
| 0 | EN | Enable. Must be set to enable this window. |

14.5.12.2 PCI Express RC Inbound Window Translation Address Register n (PEX_RCIWTAR0–PEX_RCIWTAR3)

PEX_RCIWTAR0–PEX_RCIWTAR3, shown in [Figure 14-138](#), contain the CSB address to be mapped for a PCI Express inbound transaction hitting the respective base address register of this window. These registers are valid only in RC mode. Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an outbound window, or where an outbound translation window points back into an inbound window, are not allowed.


Figure 14-138. PCI Express RC Inbound Window Translation Address Register n (PEX_RCIWTAR0–PEX_RCIWTAR3)

[Table 14-137](#) defines the bit fields of PEX_RCIWTAR n .

Table 14-137. PEX_RCIWTAR n Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–12 | TA | Translation address. Contains the CSB base address to be mapped for a PCI Express inbound transaction hitting the respective base address register of this window. The actual address is a concatenation of the TA field as most significant bits and 12 zeros as least significant bits {TA[31:12], 0x000}. |
| 11–0 | — | Reserved. Must be zeros. |

14.5.12.3 PCI Express RC Inbound Window Base Address Register Low *n* (PEX_RCIWBARL0–PEX_RCIWBARL3)

PEX_RCIWBARL0–PEX_RCIWBARL3, shown in Figure 14-139, contains the lower portion base address of the PCI Express domain corresponding to this window.




Figure 14-139. PCI Express RC Inbound Window Base Address Register Low *n* (PEX_RCIWBARL0–PEX_RCIWBARL3)

Table 14-138 defines the bit fields of PEX_RCIWBARLn.

Table 14-138. PEX_RCIWBARLn Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–12 | BAL | Base address low. Lower portion of the PCI Express address base. Represents the PCI Express-based address for the window. The actual address is a concatenation of the BAL field as most significant bits and 12 zeros as least significant bits {BAL[31:12], 0x000}. |
| 11–0 | — | Reserved. Must be zeros. |

14.5.12.4 PCI Express RC Inbound Window Base Address Register High *n* (PEX_RCIWBARH0–PEX_RCIWBARH3)

PEX_RCIWBARH0–PEX_RCIWBARH3, shown in Figure 14-140, contain the higher-portion base address of the PCI Express domain corresponding to this window. This register should be used in 64-bit addressing. Otherwise, it should contain all zeros.



Figure 14-140. CI Express RC Inbound Window Base Address Register High *n* (PEX_RCIWBARH0–PEX_RCIWBARH3)

Table 14-139 defines the bit fields of PEX_RCIWBARHn.

Table 14-139. PEX_RCIWBARHn Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 31–0 | BAH | Base address high. Higher portion of the PCI Express address base ([63:32]). The complete 64-bit address on the PCI Express bus is built of {PEX_RCIWBARH[BAH], PEX_RCIWBARL[BAL], 0b0000000000}. |

14.6 Functional Description

The PCI Express protocol relies on a requestor/completer relationship in which one device requests that a target device perform an action, and the target device completes the task and responds. Usually, the requests and responses occur through a network of links, but to the requestor and to the completer, the intermediate components are transparent.




Figure 14-141. Requestor/Completer Relationship

Each PCI Express device is divided into two halves, transmit (Tx) and receive (Rx), and each of these halves is further divided into three layers—transaction, data link, and physical—as shown in [Figure 14-142](#).

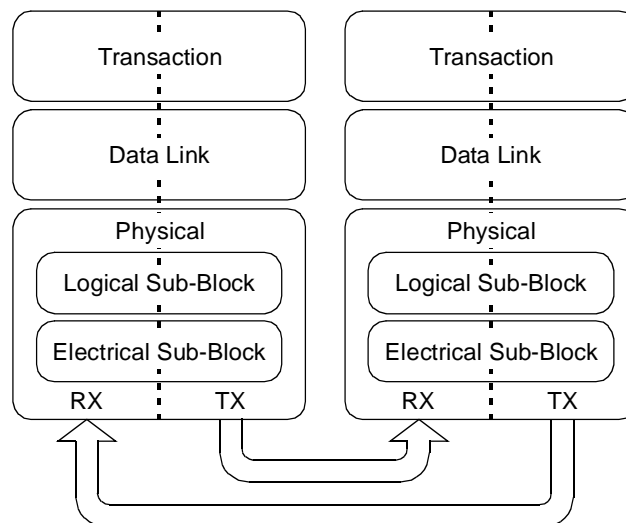


Figure 14-142. PCI Express High-Level Layering

Packets are formed in the transaction layer (TL) and data link layer (DLL), and each subsequent layer adds the necessary encoding and framing. As packets are received, they are decoded and processed by the same layers but in reverse order, so they may be processed by the layer or by the device application software.

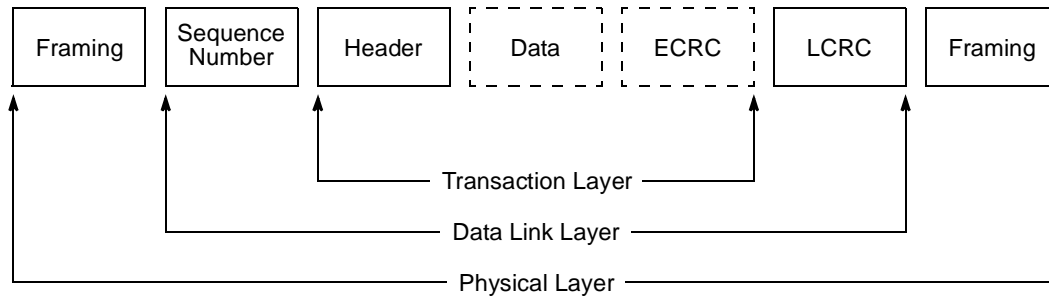


Figure 14-143. PCI Express Packet Flow

14.6.1 Architecture

14.6.1.1 Address Translation Windows (ATMUs)

The device includes four general purpose inbound and outbound address translation windows which are used to map between the PCI Express domain and the CSB domain (referred as local address space). The outbound windows are common for both Root Complex and End Point modes, and their configuration registers reside in the CSB bridge space. For Inbound windows there is a different set of configuration registers between Root Complex and End Point modes. EP inbound base address registers (BARs) reside in the standard configuration header space of the PCI Express core, while RC inbound base address registers (BARs) reside in the CSB bridge space. Since the CSB domain supports only 32 bit addressing, outbound base address registers and inbound translation registers are 32 bit only. [Table 14-140](#) specifies the available combinations between base address registers and translation address registers.

Table 14-140. Address Translation Window Combinations

| Window Number | Type | BAR Name | BAR Address | TAR Name | TAR Address |
|---|------------------------------|------------|-------------|--------------|-------------|
| Outbound Windows—Common for Root Complex and End Point Modes | | | | | |
| 0 | 64-bit address, low portion | PEX_OWBAR0 | 0x9CA4 | PEX_OWTARL0 | 0x9CA8 |
| | 64-bit address, high portion | | | PEX_OWTARH0 | 0x9CAC |
| 1 | 64-bit address, low portion | PEX_OWBAR1 | 0x9CB4 | PEX_OWTARL1 | 0x9CB8 |
| | 64-bit address, high portion | | | PEX_OWTARH1 | 0x9CBC |
| 2 | 64-bit address, low portion | PEX_OWBAR2 | 0x9CC4 | PEX_OWTARL2 | 0x9CC8 |
| | 64-bit address, high portion | | | PEX_OWTARH2 | 0x9CCC |
| 3 | 64-bit address, low portion | PEX_OWBAR3 | 0x9CD4 | PEX_OWTARL3 | 0x9CD8 |
| | 64-bit address, high portion | | | PEX_OWTARH3 | 0x9CDC |
| Inbound Windows—End Point Mode | | | | | |
| 0 | 32-bit address | BAR0 | 0x010 | PEX_EPIWTAR0 | 0xDE0 |
| 1 | 32-bit address | BAR1 | 0x014 | PEX_EPIWTAR1 | 0xDE4 |

Table 14-140. Address Translation Window Combinations

| Window Number | Type | BAR Name | BAR Address | TAR Name | TAR Address |
|--|------------------------------|---------------|-------------|--------------|-------------|
| 2 | 64-bit address, low portion | BAR2 | 0x018 | PEX_EPIWTAR2 | 0xDE8 |
| | 64-bit address, high portion | BAR3 | 0x01C | | |
| 3 | 64-bit address, low portion | BAR4 | 0x020 | PEX_EPIWTAR3 | 0xDEC |
| | 64-bit address, high portion | BAR5 | 0x024 | | |
| Inbound Windows—Root Complex Mode | | | | | |
| 0 | 64-bit address, low portion | PEX_RCIWBARL0 | 0x9E68 | PEX_RCIWTAR0 | 0x9E64 |
| | 64-bit address, high portion | PEX_RCIWBARH0 | 0x9E6C | | |
| 1 | 64-bit address, low portion | PEX_RCIWBARL1 | 0x9E78 | PEX_RCIWTAR1 | 0x9E74 |
| | 64-bit address, high portion | PEX_RCIWBARH1 | 0x9E7C | | |
| 2 | 64-bit address, low portion | PEX_RCIWBARL2 | 0x9E88 | PEX_RCIWTAR2 | 0x9E84 |
| | 64-bit address, high portion | PEX_RCIWBARH2 | 0x9E8C | | |
| 3 | 64-bit address, low portion | PEX_RCIWBARL3 | 0x9E98 | PEX_RCIWTAR3 | 0x9E94 |
| | 64-bit address, high portion | PEX_RCIWBARH3 | 0x9E9C | | |

The attributes of the outbound windows are controlled by the PEX_OWAR n registers and the attributes of the RC inbound windows are controlled by the PEX_RCIWAR n registers, residing in the CSB bridge address space. The attributes of the EP inbound windows are controlled by both the translation registers, PEX_EPIWTAR n , in the CSB bridge space, and by the PCI Express BAR configuration registers residing in the PCI Express core address space.

NOTE

For outbound transactions, both the PCI Express DMA and CSB Masters share the same set of windows.

14.6.1.2 Byte Ordering

Whenever data must cross a bridge between two busses, the byte ordering of data on the source and destination buses must be considered. The internal platform bus of this device is inherently big endian and the PCI Express bus interface is inherently little endian.

There are two methods to handle ordering of data as it crosses a bridge—address invariance and data invariance. Address invariance preserves the addressing of bytes within a scalar data element, but not the relative significance of the bytes within that scalar. Conversely, data invariance preserves the relative significance of bytes within a scalar, but not the addressing of the individual bytes that make up a scalar.

This device uses address invariance as its byte ordering policy.

As stated above, address invariance preserves the byte address of each byte on an I/O interface as it is placed in memory or moved into a register. This policy can have the effect of reversing the significance order of bytes (most significant to least significant and vice versa), but it has the benefit of preserving the

format of general data structures. Provided that software is aware of the endianness and format of the data structure, it can correctly interpret the data on either side of the bridge.

Figure 14-144 shows the transfer of a 4-byte scalar, 0x4142_4344, from a big endian source across an address invariant bridge to a little endian destination.




Figure 14-144. Address Invariant Byte Ordering—4 bytes Outbound

Note that although the significance of the bytes within the scalar have changed, the address of the individual bytes that make up the scalar have not changed. As long as software is aware that the source of the data used a big endian format, the data can be interpreted correctly.

Figure 14-145 shows data flowing the other way, from a little endian source to a big endian destination.




Figure 14-145. Address Invariant Byte Ordering—4 bytes Inbound

Figure 14-146 shows an outbound transfer of an 8-byte scalar, 0x5455_1617_CDCE_2728, using address invariance.




Figure 14-146. Address Invariant Byte Ordering—8 bytes Outbound

Figure 14-147 shows an inbound transfer of a 2-byte scalar, 0x5837, using address invariance.




Figure 14-147. Address Invariant Byte Ordering—2 bytes Inbound

Note that in all of these examples, the original addresses of the individual bytes within the scalars (as created by the source) have been preserved.

14.6.1.2.1 Byte Order for Configuration Transactions

All internal memory-mapped registers in the CCSR space use big endian byte ordering. However, the PCI Express specification defines PCI Express configuration registers as little endian. All accesses to the PCI Express configuration port, PEX_CONFIG_DATA, including the those targeting the internal PCI Express configuration registers, use the address invariance policy as shown in Figure 14-148. Therefore, software must access PEX_CONFIG_DATA with little-endian formatted data—either using the `lwbrx/stwbrx` instructions or by manipulating the data before writing to and after reading from PEX_CONFIG_DATA.




Figure 14-148. PEX_CONFIG_DATA Byte Ordering

14.6.1.3 Transaction Ordering Rule

In general, transactions are serviced in the order that they were received. However, transactions can be reordered as they are sent due to stalled conditions such as the internal buffer full condition. Below are the ordering rules for sending the next outstanding request.

- A posted request can and will bypass all other transactions except another posted request.
- Completion can and will bypass non-posted and posted requests only if the relaxed ordering (RO) bit of the PCI Express packet's header is set.
- A non-posted request cannot bypass a posted or non-posted request but can bypass completion if the relaxed ordering bit is set.

Note that it is possible for one outbound configuration transaction to bypass another outbound configuration transaction due to CRS status and the ability for hardware to retry the transaction.

14.6.1.4 Memory Space Addressing

A PCI Express memory transaction can address a 32- or 64-bit memory space. The Fmt[0] field in the PCI Express header for a 32-bit address packet is 0 while a 64-bit address packet has a 1 indication. A memory read transaction has settings of 00000 for the Type[4:0] field in the PCI Express header and 0 for Fmt[1]. A memory write transaction has the Type[4:0] field in the PCI Express header as 00000 and Fmt[1] as 1. As an initiator, the controller is capable of sending a 32- or 64-bit memory packet depending on the window translation address. Any transaction from the CSB that has a translated address greater than 4G after going through the translation window is sent as a 64-bit memory packet. Otherwise, a 32-bit memory packet is sent. As a target device, the controller can decode a 32- or 64-bit memory packet using two 32-bit inbound windows and two 64-bit inbound windows. All inbound addresses are translated to the CSB address, which is 32 bits wide.

14.6.1.5 I/O Space Addressing

The PCI Express controller does not support I/O transactions as a target. As an initiator, the controller can send I/O transactions in RC mode by programming one of the outbound translation window's attributes to send I/O transactions. All I/O transactions access only a 32-bit address I/O space. An I/O read transaction has the Type[4:0] field in the PCI Express header as 00010 and the Fmt[1] as 0. An I/O write transaction has the Type[4:0] field in the PCI Express header as 00010 and the Fmt[1] as 1.

14.6.1.6 Configuration Space Access

To access the PCI Express controller's internal configuration header by MPC8315E itself, the only mechanism supported is the direct access via the CSB, since the whole internal configuration space is memory-mapped. This is true regardless the PCI Express controller is configured as RC or EP.

If the PCI Express controller is configured as RC,

- Inbound configuration transaction is not supported.
- Outbound configuration transaction to access downstream PCI Express devices is supported. The only mechanism can be used to initiate either Type 0 or Type 1 configuration cycle is via outbound ATMU windows.

If the PCI Express controller is configured as EP,

- Outbound configuration transaction is not supported. In other words, the PCI Express EP controller does not generate configuration transactions in EP mode.
- Inbound configuration transaction to access the PCI Express EP controller's configuration space is supported.

14.6.1.6.1 Outbound ATMU Configuration Transaction Generation (RC)

In RC mode, the PCI Express controller can generate both Type 0 and Type 1 configuration cycles to access downstream PCI Express devices via the outbound ATMU windows mechanism.

As RC the PCI Express controller configuration access mechanism utilizes a memory-mapped address space to access device configuration registers. To achieve this, software can program the TYPE field of the PEX_OWARn register to 0x0 in one of the outbound ATMU windows to perform a configuration access. The other bit fields of the PEX_OWARn register should be programmed as below:

- TC = 0x0;
- NSNP = 0;
- RLXO = 0;
- EN = 1;

The SIZE field of the PEX_OWARn register should be set to a value to correspond with the BA (base address) field of the base address register (PEX_OWBARn), normally based on the Bus Number(s) of the downstream PCI Express device(s) to be accessed. The base address registers, PEX_OWBARn, set the CSB address window for the configuration transactions. The translation address registers, PEX_OWTARLn, can be used to define the translated PCI Express address of the CSB-based configuration transaction.

Once the PCI Express outbound window attributes register (PEX_OWARn), base address register (PEX_OWBARn) and translation address register (PEX_OWTARLn) are fully defined, a CSB-based memory transaction hitting the defined base address register will be converted to an external PCI Express configuration transaction cycle appeared on the downstream link of the PCI Express RC controller. In this case, the CSB memory address determines the configuration register accessed and the memory data returns the contents of the addressed register. Software must only issue 4-byte or less access to the ATMU configuration window and the access cannot cross a 4-byte boundary.

The mapping from the CSB local address to PCI Express configuration space is defined by the Outbound ATMU as in a memory transaction. The actual PCI Express configuration header will be formatted from the CSB address and the ATMU translation, defined by PEX_OWTARLn.

The formatted address defines the configuration transactions parameter, as shown in [Table 14-141](#). Note that there is no byte swapping for the address itself, although the programming of the registers content do require byte swapping. The table also translates between the bit ordering commonly used by the Power PC terminology (0:31) and the bit ordering used by the PCI Express terminology (31:0).

Table 14-141. Configuration Address Mapping

| CSB Address Bits Numbering | PCI Express Address Bits Numbering | PCI Express Configuration Space |
|----------------------------|------------------------------------|---------------------------------|
| 0:7 | 31:24 | Bus number |
| 8:12 | 23:19 | Device number |
| 13:15 | 18:16 | Function number |
| 16:19 | 15:12 | Reserved |
| 20:23 | 11:8 | Extended register number |
| 24:29 | 7:2 | Register number |

Table 14-141. Configuration Address Mapping (continued)

| CSB Address Bits Numbering | PCI Express Address Bits Numbering | PCI Express Configuration Space |
|----------------------------|------------------------------------|---------------------------------|
| 30:31 | 1:0 | Reserved |

Note: It is the user’s responsibility to set the reserved bit fields to zero (especially bits 15–12).

Note: The configuration cycle generation mechanism does not differentiate from internal or external configuration cycle. This means that any transaction which hits a configuration window will be passed to the PCI Express link with a relevant transaction type.

NOTE

The location of the configuration fields described in the table above are slightly different than the definition of a “flat memory addressing method” in the PCI Express specification, and are directly aligned to the fields in the configuration transaction header. The user software should take this difference into considerations.

The PCI Express RC controller initiates the Type 0 or Type 1 configuration cycle on its downstream link based on the following rules:

- If the bus number of the CSB-initiated transaction equals the secondary bus number from Type 1 header of RC’s configuration space and the device number is 0, a Type 0 configuration cycle will be sent to the link.
- If the bus number of the CSB-initiated transaction does not equal the RC’s primary bus number, and does not equal the secondary bus number (from RC’s Type 1 header), and is less than or equal to the subordinate bus number (from RC’s Type 1 header), a Type 1 configuration cycle will be sent to the link. Note that according to PCI and PCI Express base specifications, the relationship where the Secondary Bus Number ≤ Subordinate Bus Number must be ensured when configuring the two bus numbers within RC’s Type 1 header.
- For all other cases, the PCI Express RC controller will issue a configuration cycle on the link whenever an outbound configuration window is hit on the CSB side, no matter what the parameters are. It is the software driver’s responsibility to block transactions with unsupported bus, device, and function numbers and return “1”s for such reads. If the bus number in the CSB-initiated transaction equals the primary bus number of RC hitting the outbound configuration window, software error will occur which must be handled by the driver.

The following is an example showing how to configure the related registers of one of the MPC8315E outbound windows for configuration transaction generation purpose. To simplify the illustration, this example has the following assumption:

- The Boot ROM location is configured to locate within 0x0000_0000 to 0x007F_FFFF.
- The overall PCI Express system has a total of 16 buses to be configured.

Sixteen buses mean that the bus number can range from 0x00 to 0x0F. In general, if there are 2^n bus numbers to be configured, n number of address bits are needed to represent the variation of bus number ranging from 0 to $2^n - 1$. For this example, four address bits from CSB[4:7] are required to represent the bus number variation and therefore become the most significant four bits within the total 28 bit offset (CSB[4:31], including reserved bits) of the outbound window to be configured. The CSB[0:3] in this case

will not participate in the bus number mapping process as shown in [Table 14-141](#), where all eight bits of CSB[0:7] are mapped to PCI Express address bits [31:24] to represent the possible of 256 bus numbers within a very large system. In other words, in the example, CSB[0:3] become the four most significant bits of the base address of the outbound window to be configured and will be translated to a new “address” in the PCI Express space as defined by the corresponding PEX_OWTARLn.

Based on the above information, the most significant four bits of the base address of the outbound window came from CSB[0:3] must be unique within the total 4 Gbyte local CSB memory space. This essentially defines a window with size of 256 Mbytes, since the lower 28 bits are offset. For this example, assume a 256 Mbyte outbound window is therefore defined between 0x5000_0000 and 0x5FFF_FFFF in local CSB space. This yields the PEX_OWBARn’s BA[31:12] to be 0x5000_0, with the actual base address of the outbound window as 0x5000_0000. The SIZE[31:12] field of the PEX_OWARn register is 0x1000_0 to reflect the actual size of this outbound window as 0x1000_0000 or 256 Mbytes.

Note that the final configuration transaction to be generated is based on the information gathered from two areas: the most significant four bits from the defined TAL (translation address low) bit field of PEX_OWTARLn and the lower order bits from the CSB[4:31] offsets directly mapped to PCI Express address bits [27:0]. As mention in the note section of [Table 14-141](#), software should ensure all the reserved bits within CSB[4:31] are filled with zeros.

Since TAL[31:24] (total of eight bits) of PEX_OWTARLn could be used for mapping the bus number bit field for a general configuration transaction, while in this example the bus number to be configured ranges from 0x00 to 0x0F, the most significant four bits (TA[31:28]) are not needed for the bus number mapping and therefore must be configured as 0x0. The TA[27:12] bit field of PEX_OWTARLn is left as all zeros. This yields the PEX_OWTARLn’s TAL[31:12] to be 0x0000_0, with the actual translation address of the outbound window as 0x0000_0000. Since the size of this window is 256 Mbytes, the upper limit of the translation address is then locate at 0x0FFF_FFFF.

With the outbound window configured as above, if software intends to scan the PCI Express bus and attempts to probe the first bus immediately underneath the PCI Express RC, the software will need to issue a Configuration Transaction to read Register Number 0 (Vendor ID Register) at Bus Number/Device Number/Function Number of 1/0/0, assuming the RC’s secondary bus number has been configured as 0x1 along with subordinate bus number being configured with a big number like 0xFF initially. As long as the LAW is configured to ensure that the local address space between 0x5000_0000 and 0x5FFF_FFFF is configured for PCI Express, a CSB-based memory read transaction to local address 0x5100_0000 initiated by software will be translated to a transaction hitting PCI Express RC controller with PCI Express address of 0x0100_0000. Upon receiving this CSB-based transaction, the RC controller checks the Type attribute of this outbound window and realizes the transaction is of a configuration type, instead of directly using the translated address as in usual memory transaction, the RC controller starts compose a configuration transaction with header information from this translation address based on the mapping defined in [Table 14-141](#). The decode of the mapping process yields a type 0 configuration transaction to be generated on the PCI Express link with Bus Number=1, Device Number=0, Function Number=0, and Register Number=0.

Similarly, if the software intends to read the above EP’s configuration space offset 0x440, it can generate a CSB-based memory transaction to address 0x5100_0440. Once the transaction hitting the above configured outbound window, the ATMU translates it into a transaction with PCI Express address 0x0100_0440. This RC controller, once receiving the transaction, will issue a Type 0 configuration

transaction on the PCI Express link with Bus Number=1, Device Number=0, Function Number=0, Extended Register Number=0x4, and Register Number=0x40.

NOTE

In the example above the translation address register (PEX_OWTARLn) is set once. It is also possible to use a dynamic approach of updating the translation address register before every configuration access. In this method the PEX_OWBARn and the PEX_OWARn for the configuration window can be set for a relatively small address range, but the software needs to adjust the translation address register (PEX_OWTARn) for the configuration window to the desired parameters prior to issuing the configuration transaction.

The programming of the ATMU registers must guarantee that there is no overlap between address bits defined by the base address and size of the window, and address bits defined by the translation address. In other words, the bits in the lower portion of the PEX_OWTARLn covered by the base address must be zero.

14.6.1.6.2 EP Configuration Register Access

When the PCI Express controller is configured as an EP device it responds to remote host generated configuration cycles. This is indicated by decoding the configuration command along with type 0 access in the packet. A remote host can access up to 4096 bytes of the PCI Express configuration area. While in EP mode, the PCI Express controller does not support generating configuration accesses as a master. There is no configuration mechanism supported in EP mode using the ATMU window. If the outbound ATMU window is configured to issue a configuration transaction, all posted transactions hitting this window are ignored and all non-posted transactions will get a response with an error and can lead to unexpected results.

14.6.1.7 Messages

14.6.1.7.1 Inbound Messages

The following tables lists the messages and the actions that take place depending on whether RC or EP mode is configured. The actual events are logged in the PCI Express Root Error Status Register and in the CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR). See [Section 14.4.5.10, “PCI Express Root Error Status Register](#), [Section 14.5.8.4, “CSB System Miscellaneous Interrupt Enable Register \(PEX_CSMIER\)](#) and [Section 14.5.8.8, “CSB System Miscellaneous Interrupt Status Register \(PEX_CSMISR\)](#) for further details.

[Table 14-142](#) lists the messages and the actions that take place in RC mode.

Table 14-142. PCI Express RC Inbound Message Handling

| Name | Code[7:0] | Routing[2:0] | Action |
|---------------------------|-----------|--------------|--|
| Assert_INTA | 0010 0000 | 100 | Set INTA event |
| Assert_INTB | 0010 0001 | 100 | Set INTB event |
| Assert_INTC | 0010 0010 | 100 | Set INTC event |
| Assert_INTD | 0010 0011 | 100 | Set INTD event |
| Deassert_INTA | 0010 0100 | 100 | Clear INTA event |
| Deassert_INTB | 0010 0101 | 100 | Clear INTB event |
| Deassert_INTC | 0010 0110 | 100 | Clear INTC event |
| Deassert_INTD | 0010 0111 | 100 | Clear INTD event |
| PM_Active_State_Nak | 0001 0100 | 100 | No action taken |
| PM_PME | 0001 1000 | 000 | Set PM_PME event |
| PME_Turn_Off | 0001 1001 | 011 | No action taken |
| PME_TO_Ack | 0001 1011 | 101 | Log entered_I23_state in PME and message detect register and generate interrupt to IPIC if enabled |
| ERR_COR | 0011 0000 | 000 | Set correctable error event |
| ERR_NONFATAL | 0011 0001 | 000 | Set non-fatal error event |
| ERR_FATAL | 0011 0011 | 000 | Set fatal error event |
| Unlock | 0000 0000 | 000 | No action taken |
| Set_Slot_Power_Limit | 0101 0000 | 100 | No action taken |
| Vendor_Defined Type 0 | 0111 1110 | | No action taken |
| Vendor_Defined Type 1 | 0111 1111 | | No action taken |
| Attention_Indicator_On | 0100 0001 | 100 | No action taken |
| Attention_Indicator_Blink | 0100 0011 | 100 | No action taken |
| Attention_Indicator_Off | 0100 0000 | 100 | No action taken |
| Power_Indicator_On | 0100 0101 | 100 | No action taken |
| Power_Indicator_Blink | 0100 0111 | 100 | No action taken |
| Power_Indicator_Off | 0100 0100 | 100 | No action taken |
| Attention_Button_Pressed | 0100 1000 | 100 | Log in PME and message detect register if enabled. Send interrupt if enabled. |

Table 14-143 lists the messages and the actions that take place in EP mode.

Table 14-143. PCI Express EP Inbound Message Handling

| Name | Code[7:0] | Routing[2:0] | Action |
|---------------------------|-----------|--------------|--|
| Assert_INTA | 0010 0000 | 100 | No action taken |
| Assert_INTB | 0010 0001 | 100 | No action taken |
| Assert_INTC | 0010 0010 | 100 | No action taken |
| Assert_INTD | 0010 0011 | 100 | No action taken |
| Deassert_INTA | 0010 0100 | 100 | No action taken |
| Deassert_INTB | 0010 0101 | 100 | No action taken |
| Deassert_INTC | 0010 0110 | 100 | No action taken |
| Deassert_INTD | 0010 0111 | 100 | No action taken |
| PM_Active_State_Nak | 0001 0100 | 100 | No action taken |
| PM_PME | 0001 1000 | 000 | No action taken |
| PME_Turn_Off | 0001 1001 | 011 | 1. Log in PME and message detect register if enabled. Send interrupt if enabled. |
| PM_TO_Ack | 0001 1011 | 101 | No action taken |
| ERR_COR | 0011 0000 | 000 | No action taken |
| ERR_NONFATAL | 0011 0001 | 000 | No action taken |
| ERR_FATAL | 0011 0011 | 000 | No action taken |
| Unlock | 0000 0000 | 000 | No action taken |
| Set_Slot_Power_Limit | 0101 0000 | 100 | No action taken |
| Vendor_Defined Type 0 | 0111 1110 | — | No action taken |
| Vendor_Defined Type 1 | 0111 1111 | — | No action taken |
| Attention_Indicator_On | 0100 0001 | 100 | Log in PME and message detect register if enabled. Send interrupt if enabled. |
| Attention_Indicator_Blink | 0100 0011 | 100 | Log in PME and message detect register if enabled. Send interrupt if enabled. |
| Attention_Indicator_Off | 0100 0000 | 100 | Log in PME and message detect register if enabled. Send interrupt if enabled. |
| Power_Indicator_On | 0100 0101 | 100 | Log in PME and message detect register if enabled. Send interrupt if enabled. |
| Power_Indicator_Blink | 0100 0111 | 100 | Log in PME and message detect register if enabled. Send interrupt if enabled. |
| Power_Indicator_Off | 0100 0100 | 100 | Log in PME and message detect register if enabled. Send interrupt if enabled. |
| Attention_Button_Pressed | 0100 1000 | 100 | No action taken |

14.6.2 Interrupts

Both INTx and Message Signaled Interrupt (MSI) generation and handling are supported; however there are subtle differences depending on whether the device is configured as an RC or EP device.

14.6.2.1 EP Interrupt Generation

Both hardware INTx messages generation and Hardware MSI generation are supported, for the interrupt events described in [Section 14.5.7, “PCI Express Host Interrupt Registers.”](#) The INTx message and MSI mechanism are mutually exclusive. Only one of these two mechanisms should be enabled and used at a given time.

14.6.2.1.1 Hardware INTx Message Generation

Hardware INTx message is generated when any interrupt event occurs and the corresponding interrupt is enabled in PEX_HIER register, if the interrupt disable bit is cleared in the PCI Express EP’s command register (see [Section 14.4.1.3, “PCI Express Command Register”](#)) and the MSI interrupt mechanism is not enabled.

Only INTA message is supported by this device.

Note that MSI interrupt mechanism will be used as long as it is enabled, regardless the setting of the interrupt disable bit in the PCI Express EP’s command register.

14.6.2.1.2 Software INTx Message Generation

Software INTx message generation is not supported.

14.6.2.1.3 Hardware MSI Generation

Host software must set up the MSI capability registers to enable MSI mode and put the correct MSI address and data values into the MSI capability registers prior to setting up various interrupt event enable bits in the PEX_HIER register to enable the generation of the correct MSI cycle to RC.

Note that the value being programmed by the host software into the MSI Data register of EP’s Type 0 configuration space is a 16-bit base message data pattern, which is referred to as “base vector [15:0]” in the description below:

- If only one MSI message is desired, the EP software can directly use this “base vector” as the interrupt vector for MSI interrupt generation. In such case, there is no need to program any of the “vector registers” among PEX_HOPIVR, PEX_HIPIVR, PEX_HWDIVR, PEX_HRDIVR and PEX_HMIVR. The PEX_HIER register setting determines which event can trigger this single MSI interrupt message to RC. Note that multiple interrupt events are allowed to share the same MSI interrupt vector.
- If multiple MSI messages are desired, the Multiple Message Capable bit field of EP’s MSI Message Control Register can be used to indicate how many MSI messages (in the power of two, up to 32 messages allowed per EP) the EP wants to use. During configuration stage, after examining the above desired value, the Host software will allocate the actual number of MSI messages for an EP to use by configuring the Multiple Message Enable bit field in the same register, in addition to

programming the “base vector” in EP’s MSI Data Register. Once this is accomplished, the EP software can program the IVEC bit field of each individual “vector register” based on the number of MSI messages allocated by host. The IVEC value must be unique for the same EP and start from 0x00. At last, the EP software can set the corresponding interrupt event enable bits in the PEX_HIER to enable the Hardware MSI generation. The actual MSI data value for a given interrupt event used by the EP in its MSI message to RC is formed by the concatenation of the “base vector [15:5]” and the IVEC [4:0] value of the corresponding interrupt event.

As an example, if the value of the Multiple Message Enable bit field allocated by host software is 010b (4 MSI messages allocated) and the “base vector” being programmed by host software in EP’s MSI Data Register is 0x55A0 (lower-16bits little endian), the actual MSI data values of all possible MSI messages can be used by the EP are 0x0000_55A0, 0x0000_55A1, 0x0000_55A2, and 0x0000_55A3. The EP software only needs to program the four possible lower-order bits (0x00, 0x01, 0x02, and 0x03) as unique IVEC values in its “vector registers” among PEX_HOPIVR, PEX_HIPIVR, PEX_HWDIVR, PEX_HRDIVR and PEX_HMIVR. If both OPAIE and OPCIE bit fields are enabled in the PEX_HIER register, assuming PEX_HOPIVR register’s IVEC bit field is programmed as 0x03h, when any one of these two interrupt events occurs, the EP will use 0x0000_55A3 as the actual MSI data value in its MSI message sent to RC. Once receiving such MSI message, the device driver running at RC is responsible to issue a read to EP’s PCI Express Interrupt Status Register (PEX_HISR) to find out exactly which of the two interrupt events caused the interrupt.

14.6.2.1.4 Software MSI Generation

Host software needs to set up the MSI capability registers to enable MSI mode and put the correct values for the MSI address and data register. Next, local software must read the MSI address in the MSI capability register and configure the outbound ATMU window to map the translated address to the MSI address. Software determines the number of allocated messages in the MSI capability register and allocates the appropriate data values to use. A write to the MSI ATMU window with the appropriate data value generates the MSI transaction to the RC.

14.6.2.2 RC Handling of INTx Message and MSI Interrupt

14.6.2.2.1 INTx Message Handling

MSIs are the preferred interrupt signaling mechanism for the PCI Express. However, in RC mode, the PCI Express controller supports the INTx virtual-wire interrupt signaling mechanism (as described in the *PCI Express Base Specification*). When the controller receives an inbound INTx asserted message, it sets the appropriate INTA, INTB, INTC, or INTD bit in the CSB system miscellaneous interrupt status register (PEX_CSMISR). An interrupt is generated to the local host if the interrupts are enabled in the CSB system miscellaneous interrupt enable register (PEX_CSMIER) and the IPIC.

14.6.2.2.2 MSI Handling

An MSI interrupt cycle must hit into the IMMRBAR window (window 0) with the address offset that points to MIISR register in the IPIC. Note that the host software must configure the EP’s MSI capability register so that an MSI cycle generated from the device is routed to the correct MIISR register in the IPIC and for the appropriate interrupt to be generated to the core.

14.6.2.3 Initial Credit Advertisement

To prevent overflowing of the link partner's receiver buffers and for compliance with ordering rules, the transmitter cannot send transactions unless it has enough credits to send. Each device maintains a flow control (FC) credit pool. The FC information is conveyed between two links by DLLPs during link training (initial credit advertisement). The transaction layer performs the FC accounting functions. It functions as the FC gate. One unit of FC is 4 DWs (16 bytes) of data.

Table 14-144. Initial Credit Advertisement

| Credit Type | Initial Credit Advertisement |
|---|-------------------------------|
| PH (memory write, message write) | 4 |
| PD (memory write, message write) | $(256 \div 16) \times 4 = 64$ |
| NPH (memory read, I/O read, cfg read) | 8 |
| NPD (I/O write, cfg write) | 2 |
| CPLH (memory read completion, I/O R/W completion, cfg R/W completion) | Infinite |
| CPLD (memory read completion, I/O read completion, cfg read completion) | Infinite |

14.6.3 Mailbox

The Mailbox mechanism is useful when the device is in EP mode, and enables exchanging information between the remote PCI Express root complex and the local host using interrupts and data storage registers. There are sets of register for both inbound and outbound mailbox messages and interrupts.

14.6.3.1 Outbound Mailbox

The EP local host uses the outbound mailbox messages to signal the remote RC device across the PCI Express link. The local host, for example the e300 core, stores the required message in the PCI Express outbound mailbox data register (PEX_OMBDR) and then initiates an interrupt (MSI) to the remote PCI Express device. When the remote PCI Express RC device receives the interrupt it can perform a memory read from the mailbox data register and read the message. The following steps are required in order to use the outbound mailbox mechanism.

1. The RC should set the MSIE bit of the PCI Express MSI message control register (address 0x72 of the configuration space) to enable the generation of MSI.
2. The EP local host should set the OMBIE bit of the PCI Express host interrupt enable register (PEX_HIER) to enable interrupt generation to the PCI Express (MSI) at the event of setting the READY bit of the PCI Express outbound mailbox control register (PEX_OMBCR).
3. The EP local host should program the IVEC field of the PCI Express host miscellaneous interrupt vector register (PEX_HMIVR) with an appropriate vector value. This value, along with the value programmed in the EP's PCI Express MSI message control register's MME field determines the MSI message data sent to the RC. For example, if the MME field has a value of N, then the lower N bits of the MSI message data are replaced with the lower N bits of the PEX_HMIVR register.
4. The EP local host should program the MBD field of the PCI Express outbound mailbox data register (PEX_OMBDR) with the message to be read by the PCI Express remote device (user defined).

5. The EP local host should set the READY bit of the PCI Express outbound mailbox control register (PEX_OMBCR). This will generate an interrupt (MSI) to the PCI Express root complex.
6. The PCI Express RC, after receiving the MSI will perform a memory read to the EP's PCI Express outbound mailbox data register (PEX_OMBDR) and get the message content.
7. The PCI Express RC should perform a memory write to clear the READY bit of the EP's PCI Express outbound mailbox control register (PEX_OMBCR).
8. The EP can repeat steps 3–5 with an appropriate MSI and message data after verifying that the PEX_OMBCR[READY] is cleared.

14.6.3.2 Inbound Mailbox

The remote PCI Express RC device uses the inbound mailbox messages to signal the EP local host across the PCI Express link. The RC performs a memory write and stores the required message in the PCI Express inbound mailbox data register (PEX_IMBDR) and then initiates an interrupt to the local host by performing a memory write and setting the READY bit of the PCI Express inbound mailbox control register (PEX_IMBCR). When the local host detects the interrupt it can read the message from the mailbox data register. The following steps are required in order to use the outbound mailbox mechanism.

1. The EP local host should enable PCI Express interrupts by programming the integrated programmable interrupt controller (IPIC).
2. The EP local host should set the IMBIE bit of the CSB system miscellaneous interrupt enable register (PEX_CSMIER), to allow interrupt at the event of mailbox ready.
3. The PCI Express RC should perform a memory write and store the required message in the EP's PCI Express Inbound Mailbox Data Register (PEX_IMBDR).
4. The PCI Express RC should perform a memory write and set the READY bit of the EP's PCI Express inbound mailbox control register (PEX_IMBCR). This will issue an interrupt to the local host.
5. The local host can read the message content from the PCI Express inbound mailbox data register (PEX_IMBDR).
6. The local host should clear PEX_IMBCR[READY] and all the interrupt event and status bits in the relevant registers.
7. The RC can repeat steps 3–4 after verifying that the PEX_IMBCR[READY] is cleared.

14.6.4 Power Management


All device power states are supported except D3cold. In addition, all link power states are supported except the L2 and L3 states. Only L0s ASPM (active state power management) mode is supported if enabled by configuring the link control register bits 1–0 in configuration space. Note that there is no power saving in the controller when the device is put into a non-D0 state. The only power saving is the I/O drivers when the controller is put into a non-L0 link state.

Table 14-145. Power Management State Supported

| Component D-State | Permissible Interconnect State | Action |
|-------------------|--------------------------------|---|
| D0 | L0, L0s | In full operation. |
| D1 | L1 | All outbound traffic is stopped. All inbound traffic is thrown away. The only exceptions are PME messages and configuration transactions. If the device is in RC mode, a PM_Turn_Off message can be sent through the PCI Express power management command register. |
| D2 | L1 | All outbound traffic is stopped. All inbound traffic is thrown away. The only exceptions are PME messages and configuration transactions. If the device is in RC mode, a PM_Turn_Off message can be sent through the PCI Express power management command register. |
| D3hot | L1, L2/L3 ready | All outbound traffic is stopped. All inbound traffic is thrown away. The only exceptions are PME messages and configuration transactions. If the device is in RC mode, a PM_Turn_Off message can be sent through the PCI Express power management command register. Note that if a transition of D3 → D0 occurs, a reset is performed to the controller configuration space. In addition, link training restarts. |
| D3cold | Not supported | Not supported. |

14.6.4.1 L2/L3 Ready Link State

The L2/L3 Ready link state is entered after the EP device is put into a D3hot state followed by a PME_Turn_Off/PME_TO_Ack message handshake protocol. Exiting this state requires a POR or a detection of a beacon or a WAKE# signal from the EP device. The PCI Express controller as an EP device does not support the generation of beacon, so the device can alternatively use one of the GPIO signals as an enable to an external tristate buffer to generate the WAKE# signal if the device needs to wake up from an L2/L3 Ready state. As an RC device, the WAKE# signal from the EP device can be connected to one of the external interrupt input pins to service the WAKE# request if needed.


Figure 14-149. Example on How to Generate WAKE#

14.6.5 Hot Reset

When a hot reset condition occurs, the controller (in both RC and EP mode) initiates a cleanup of all outstanding transactions and goes into suspend mode. The RC controller then needs to be reset (issuing CBRST in PECR1/PECR2) to bring it back to the idle state followed by a reprogramming of all the CSB bridge registers (offset 0x800–0xFFC, such as the ATMU windows). The EP controller also needs to be

reset, followed by a reprogramming of the entire configuration space and CSB bridge registers (offset 0x800–0xFFC, such as the ATMU windows). All configuration register bits that are non-sticky are reset. Link training takes place subsequently. The device is permitted to generate a hot reset condition on the bus when it is configured as an RC device by setting the secondary bus reset bit in the bridge control register in the configuration space. In EP mode, the device is not permitted to generate a hot reset condition; it can only detect a hot reset condition and initiate the cleanup procedure appropriately.

14.6.6 Initialization Sequence

The following sequence must be followed. Note the specific stages for RC or EP modes. Upon chip reset, the default SerDes reference clock frequency is assumed to be 100 MHz. Steps 3 and 4 are required only if the reference clock frequency need to be changed to 125 MHz.

1. The device performs its power-on reset sequence. The PCI Express controller and the SerDes PHY are held in reset (controlled by memory-mapped registers).
2. Program the system configuration registers of the device, including some PCI Express controller related options (local memory windows, clock ratio, and so on) See the system configuration registers in [Section 5.2.3.1, “Local Access Register Memory Map,”](#) and the clock configuration registers in [Section 4.5.2, “Clock Configuration Registers.”](#)
3. Set the reference clock in the SRDSCR4 register. See [Section 16.3.5, “SerDes Control Register 4 \(SRDSCR4\).”](#)
4. Start the SerDes reset sequence by setting RST field (bit 0) in the SerDes reset control register (SRDSRSTCTL). See [Section 16.3.6, “SerDes Reset Control Register \(SRDSRSTCTL\).”](#)
5. Poll RDONE field (bit 1) in the SerDes reset control register (SRDSRSTCTL).
6. After RDONE is set, wait at least 1 ms.
7. Program the PCI Express control registers 1 and 2. See [Section 5.3.2.10, “PCI Express Control Registers \(PECR1 and PECR2\).”](#) Set the DEV_TYPE field, to select between EP or RC, take the PCI Express controller out of reset by setting bits [0:2] and optionally select the parameters in the PRI_DATA, PRI_DES, and PRI_PIO fields.
8. Configure the PCI Express core and CSB bridge control registers and address mapping windows to the desired values (that is, inbound/outbound windows).
9. Poll the Status Code field from the LTSSM State Status Register (see [Section 14.4.6.1, “PCI Express LTSSM State Status Register \(PEX_LTSSM_STAT\)”](#)) to determine when link negotiation is done and link is up (that is, Status Code = 16 link is up).
10. For EP mode only: After system configuration is done, set the CFG_READY bit in the configuration ready register. See [Section 14.4.6.12, “PCI Express Configuration Ready Register.”](#)
11. For RC mode only: Set the bus master and memory space fields in the PCI Express command register (in the PCI Express configuration space) to allow inbound and outbound transactions. See [Section 14.4.1.3, “PCI Express Command Register.”](#)

NOTE

Only in RC mode the local host should perform this programming. When the device is in EP mode, it is expected that the remote PCI Express RC will do this operation by a configuration access.

12. The device is ready to generate or accept PCI Express transactions according to its mode.

14.7 DMA Functional Operation

Software uses the DMA engine to initiate memory transfers from the CSB subsystem to the PCI Express system and vice versa without using programmed input/output (PIO), where data is transferred by sending control data through the CPU. The DMA engine provides control registers to enable the transfers. The PCI Express CSB bridge supports two separate engines for read and write DMA operations, referred as RDMA and WDMA, respectively. The DMA engines use a descriptor-based programming model.

NOTE

In general, the DMA descriptors use little-endian byte ordering. Software running on the local processor in big-endian mode must byte-swap the data. No byte swapping occurs when the registers are accessed from the PCI Express link.

14.7.1 DMA Descriptor Format

The DMA descriptor, shown in [Figure 14-150](#), is five DW wide and consists of control and status fields. The software is responsible to prepare the descriptors at the local memory and program the relevant DMA control registers. The hardware updates the status register on completion of a DMA data transfer for a descriptor and a chain of descriptors, and report errors.

Offset *Anywhere in local memory* :

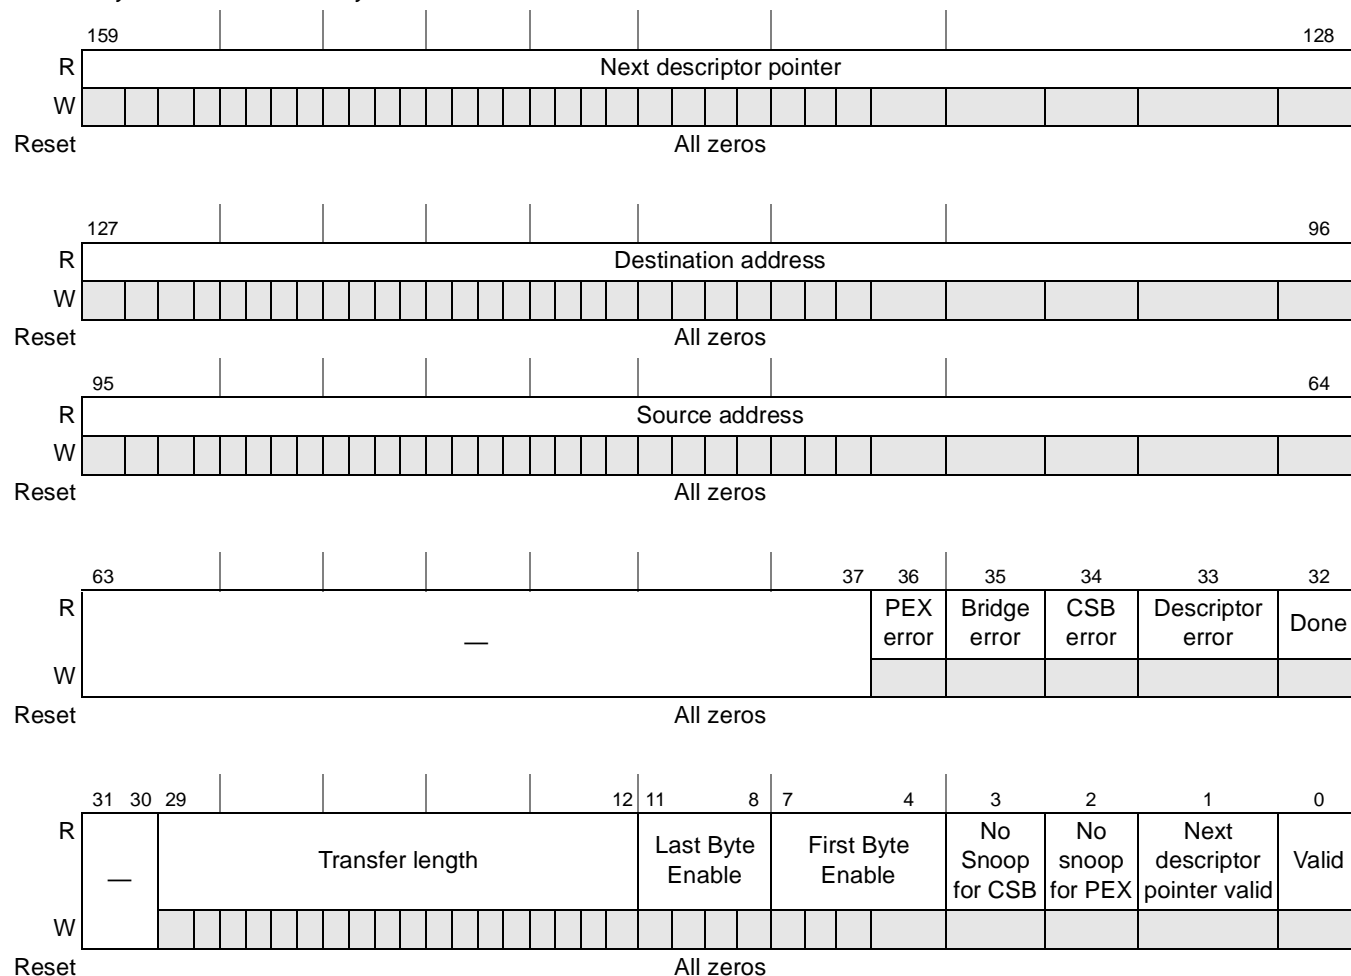


Figure 14-150. DMA Descriptor Format

Table 14-146 defines the bit fields of a DMA descriptor.

Table 14-146. DMA Descriptor Bit Field Descriptions

| Bit | Width | Attribute | Description |
|---------|-------|-----------|--|
| 159–128 | — | Control | Next descriptor pointer. Indicates the location of the next descriptor. Valid only if next_dp is set. |
| 127–96 | 32 | Control | Destination address. Software programs this field to indicate the destination address. For a write DMA, the destination address is a CSB address that is mapped to a PCI Express memory address using the outbound window address translation. For a read DMA, the destination address is the CSB address. |
| 95–64 | 32 | Control | Source address. Software programs this field to indicate the source address. For a write DMA, the source address is the CSB address. For a read DMA, the source address is a CSB address that is mapped to a PCI Express memory address using the outbound window address translation. |
| 63–37 | 27 | — | Reserved |
| 36 | — | Status | PCI Express error. Hardware sets this bit to indicate that a DMA data transfer corresponding to descriptor cannot complete due to a PCI Express error. |

Table 14-146. DMA Descriptor Bit Field Descriptions (continued)

| Bit | Width | Attribute | Description |
|-------|-------|-----------|--|
| 35 | 1 | Status | Bridge error. Hardware sets this bit to indicate that DMA data transfer corresponding to descriptor cannot complete due to Bridge error. |
| 34 | 1 | Status | CSB error. Hardware sets this bit to indicate that complete data cannot be fetched due to an CSB Error. |
| 33 | 1 | Status | Descriptor error. Hardware sets this bit to indicate that the next descriptor cannot be fetched due to an CSB error. |
| 32 | 1 | Status | Done. Hardware sets this bit after completing the transaction. |
| 31–30 | 2 | — | Reserved |
| 29–12 | 18 | Control | Transfer length. Software programs this field to indicate the length of transfers in DW or data payload size. A value of zero means that no data is transferred. |
| 11–8 | 4 | Control | Last byte enable indicates the byte enables of the last DW to be transmitted by DMA. |
| 7–4 | 4 | Control | First byte enable. Indicates the byte enables of first DW to be transmitted by DMA. |
| 3 | 1 | Control | No snoop for CSB transactions. 0 The memory transaction is broadcast on the CSB as non-global (that is, not snooped). 1 The memory transaction is broadcast on the CSB as global (that is, snooped) |
| 2 | 1 | Control | No snoop for PCI Express transactions. Indicates the no snoop value to be used in TLP header for PCI Express transactions. |
| 1 | 1 | Control | Next descriptor pointer—valid. Software sets this bit to indicate that a descriptor has a valid next descriptor pointer next_dp. When this bit is cleared, the address of the next descriptor is implicitly specified. This bit must be cleared for block descriptor-based memory. |
| 0 | 1 | Control | Valid. Software sets this bit to indicate that a descriptor is valid and has the information related to data transfer. |

14.7.2 Write DMA

The PCI Express or CSB software can program the write DMA engines to send data from the CSB system to the PCI Express. After the control registers are programmed, the write DMA engine issues a CSB read request through the DMA read master. To improve system performance, the DMA request is segmented according to a natural aligned CSB address boundary of maximum transfer size (32 bytes).

The entire address space accessed by the DMA controller is prefetchable, so the CSB read request for DMA operation always reads all the bytes. All segments use the same ID and are posted without waiting for the previous read response.

When a response to a CSB read request is received, the segments are packed into the PCI Express memory write request according to the PCI Express MPS. All data requested by the DMA controller is processed as a single data stream as described in this section. For example, when a 256-byte request starting from address 0 is segmented to eight 32-byte CSB read requests, then it is segmented to two 128-byte PCI Express write requests (assuming a 128-byte MPS).

If all data can be packed into one PCI Express write request, no segmentation is performed. Otherwise, the first segment starts from the start address and ends at the MPS address boundary. Byte enables are set according to DMA control register settings. All other segments except the last one start from the MPS

address boundary and end at the next boundary. That is, the size of a write request is equal to MPS. All bytes are enabled. Remaining data, if applicable from the last MPS address boundary to the PCI Express end address, is packed into the last segment. Byte enables are set according to DMA control register settings. When all segments get an CSB response with a status of OKAY, the DMA data transfer has completed successfully.

If any data within an CSB read response is aborted by the CSB slave (SLVERR response), all the data received before that point is packed and sent. Any remaining data returned from the CSB slave is discarded, and an error is logged.

If any segment gets a response of DECERR, all data received before that point is packed and sent. Any remaining data returned from the CSB slave is discarded, and an error is logged.

The CSB read master can issue multiple pending CSB read requests if the requests belong to a single DMA request. When all segments are successfully received (that is, they get an CSB response with a status of OKAY), the CSB read master starts processing the next DMA request.

14.7.3 Read DMA

The read DMA engines can be programmed by the PCI Express or CSB software to send data from the PCI Express system to the CSB system. After the control registers are programmed, the read DMA engine issues a PCI Express read request. The DMA request is segmented according to the PCI Express MRRS natural aligned address boundary. If all data can be requested in one read request, no segmentation is performed. Otherwise, the first segment starts from the start address and ends at the first MRRS boundary. All other segments except the last one start from the MRRS address boundary and end at the next boundary. That is, the length of the read request is equal to MRRS. Any remaining data, if applicable from the last MRRS address boundary to the end address, is requested in the last segment.

The entire address space accessed by the DMA controller is considered as prefetchable, so a PCI Express read request for DMA operation always enables all byte enables. All segments have a unique tag, so multiple read requests can be pending at any given time.

To improve system performance, when a PCI Express completion comes back, it is segmented according to the CSB maximum size natural aligned address boundary. If all data can be packed into one CSB write request, no segmentation is performed, and the write is performed according to the DMA control register settings. Otherwise, the first segment starts from the start address and ends at the first CSB address boundary. All other segments except the last one start from the CSB address boundary and end at the next boundary. That is, the size of the write request is equal to the CSB maximum size packet. All bytes within each data phase are enabled.

Any remaining data, if applicable from the last CSB address boundary to the PCI Express end address, is packed into the last segment.

Each completion is segmented as a single data stream according to these rules. For example, when a 256-byte request starting from address 0 is converted to two 128-byte PCI Express read requests (assuming a 128-byte MPS/MRRS), the bridge issues two read requests if a tag and a completion buffer are available. When any completion comes back, it is segmented as described in this section. No scatter gather is performed. Because PCI Express can return completions in any order, the bridge may not issue an CSB write request in address order.

When all segments get a response with a PCI Express completion that has a status of successful, the DMA data transfer has completed successfully. If any segment gets a response with a status other than successful, or a completion timeout occurs, the DMA data transfer completes with an error status after all requests complete either normally or abnormally. The data returned for prior requests is still processed and sent to the CSB accordingly.

14.7.4 Descriptor-Based DMA

Descriptor-based DMA operation has a specific format in which the host can store information about a data transfer, such as the source address for the data to be transferred, the destination address, the data transfer size, location of the next descriptor, and so on. The host can program a series of descriptors and store them in host local memory. The host also programs the DMA control register to indicate the location of the first descriptor. The DMA control registers are part of the bridge device-specific registers. After programming the control register and descriptors, the host is free to continue with its other functions. The DMA engine is responsible for fetching the descriptor from host memory and moving data from/to host memory.

Software can organize the descriptors in two different formats that are only for reference and do not affect the hardware functionality and requirements:

- Chain descriptor
- Block descriptor

14.7.4.1 Chain Descriptor

Chain descriptors form an n -way chain in which each descriptor implicitly or explicitly contains the address of the next descriptor. This enables the host to use memory efficiently to store the descriptors even when contiguous memory locations are not available. When the host needs to initiate another transfer, it adds another descriptor in its memory and modifies the pointer of the last descriptor in the chain to the location of the new descriptor. [Figure 14-151](#) illustrates the chain descriptor organization in host memory.




Figure 14-151. n -Way Chain Descriptor Organization in Host Memory

The number of ways, n , is software configurable. In this mode, n descriptors are written in contiguous memory locations. The implicit address of the next descriptor is the next memory location. The last descriptor in the contiguous block contains the explicit address pointer of the next set of n descriptors. Address 0x0 will never be part of the chain since it should close the chain. Non-contiguous valid descriptors are not supported. If the valid bit of a descriptor in the chain is not set, all of the succeeding descriptors should also have the valid bit as zero.

The software need to follow the following sequence on receiving a chain transfer done interrupt:

1. Software receives an interrupt for chain transfer done.

2. Software polls the main memory and waits for descriptor done bit to be set (bit 32 in the descriptor). This guarantees that the final data has been written into memory for read DMA. For write DMA, it guarantees that the final data has been sent to PCI Express controller. If a read is sent to the same location, PCI Express controller guarantees that the outbound read will not bypass the outbound write.

The host can reuse the descriptor memory after the DMA/bridge logic processes it. The exact handshake between the hardware and software is described later in this chapter. The DMA registers are described in [Section 14.5.5, “DMA Registers.”](#)

When $n > 1$, the hardware uses this knowledge to prefetch descriptors in advance, thereby reducing possible holes in transmission. This might be useful for applications that request several small DMA transfer requests, such as an Ethernet traffic. For applications that request large data transfers, the effect on performance due to descriptor fetching is not significant.

14.7.4.2 Block Descriptors

Block descriptors are a special case of chain descriptor in which all the descriptors are part of a single array. A portion of host memory is reserved to store the descriptors. The starting address of this block is indicated in the DMA control register. This block of memory serves as a circular buffer. Software writes the first descriptor to the first address in the descriptor block address space. Subsequent descriptors are written in the consecutive locations until the last location in the descriptor space. After that, the next descriptor is again written into the first location. All descriptors except the one written into the last descriptor location in the block have an implicit address that points to the immediate next descriptor location in the block. The last descriptor contains an explicit address pointer to the first descriptor location of the block. Software must ensure that a descriptor is processed by hardware before overwriting it. [Figure 14-152](#) illustrates the organization of the block descriptors in memory.




Figure 14-152. Block Descriptor Organization in Host Memory

The DMA can fetch prefetch a configurable number of descriptors in a single shot. This number depends on the chain organization (n) as well as the number of descriptor registers in the DMA engine.

14.7.4.3 Descriptor Format

For more information on descriptor format, refer to [Section 14.7.1, “DMA Descriptor Format.”](#) The fetched descriptors are stored in the bridge configuration space registers. Each time the DMA controller fetches a new set of descriptors, the register is updated to indicate the value/fields of the descriptors. After completion of a transaction specified by a descriptor, the status fields are updated in the descriptor registers in the configuration space. Also, the descriptor status is written back into host local memory and an interrupt is generated if enabled. When the last of n descriptors is processed, the next set of descriptors is fetched from memory.

14.7.4.4 Software-Hardware Handshake

Hardware and software communicate through descriptors, control registers, and interrupts. The control register programmed by software indicates to hardware the location of the first descriptor. The descriptors programmed by software in host memory provide information to the hardware about the impending data transfer. After the initial descriptor location and other DMA parameters are programmed in the DMA control registers, software sets the ‘start’ bit in the control register to trigger the DMA controller to initiate the operation. When it detects that the ‘start’ bit is set, the DMA controller uses these programmed DMA parameters to initiate a descriptor fetch and the corresponding data transfer, and then it clears the “start” bit.

The DMA controller executes the transfer according to the instructions given in the descriptor and then updates both the descriptor and DMA status register to indicate the status of the transfer operation. Descriptor status is updated in host local memory. An interrupt is also generated to the host, if enabled.

At any time, software can parse the done bit in the descriptor chain located in host local memory to determine the point to which the DMA controller has executed and also if the descriptors were successfully completed. The status register also gives details about transfer status, including details about errors if any occurred.

If the DMA encounters an unprogrammed descriptor (ready = 0) in the n descriptor array that it fetched, it first executes any remaining prefetched valid descriptors and then sends an interrupt to the host, if enabled. The transfers can resume in one of two ways:

- The DMA controller automatically fetches the same set of descriptors again after the time indicated in PEX_DMA_DSTMR[DSRT] (see [Section 14.5.2.2, “PCI Express DMA Descriptor Timer Register \(PEX_DMA_DSTMR\)”](#)). If the descriptor is ready now, it continues execution or else checks again later.
- Software can force immediate resuming of the transfers by disabling and then re-enabling the DMA.

When the DMA controller completes the data transfer for the last descriptor in the chain (null descriptor), it updates the descriptor and status register. An interrupt is generated, if enabled. The DMA engine moves into the IDLE state until software re-triggers it by setting the ‘start’ bit in the control register.



Chapter 15

SATA Controller

15.1 Overview

NOTE

The MPC8314E and MPC8314 do not support dual SATA controllers.

The serial ATA controller is a high-performance SATA solution incorporating some of the latest SATA-IO extensions. The SATA may also be referred to as a host bus adapter (HBA). The SATA controller is designed to operate in a system that supports command queuing and, in particular, a switching scheme based on a frame information structure (FIS) using port multipliers.

FIS-based switching requires the SATA controller to maintain in hardware a context for each command it has queued at the devices that are attached to it. FIS-based switching also requires the SATA controller to maintain a queue per attached device ensuring that the command issue order for each device is maintained. It can be used in SATA controllers, as well as storage area network (SAN), network attached storage (NAS), and RAID (redundant array of independent/inexpensive disks) devices.

SATA controller has the following features:

- Designed to comply with *Serial ATA 2.5 Specification*
- Supports speeds: 1.5 Gbps (first-generation SATA), 3 Gbps (second-generation SATA and eSATA)
- Supports advanced technology attachment packet interface (ATAPI) devices
- Contains high-speed descriptor-based DMA controller
- Supports native command queuing (NCQ) commands
- Supports port multiplier operation
- Supports hot plug including asynchronous signal recovery

Figure 15-1 shows a block diagram of the SATA.




Figure 15-1. SATA Block Diagram

There are four layers in the SATA architecture: application, transport, link, and PHY. The application layer is responsible for overall ATA command execution, including controlling command block register accesses. The transport layer is responsible for placing control information and data to be transferred between the host and device in a packet/frame, known as a frame information structure (FIS). The link layer is responsible for taking data from the constructed frames, inserting control characters and moving data to the PHY layer. The PHY layer is responsible for 8B/10B encoding/decoding, then transmitting and receiving the encoded information as a serial data stream on the wire.

15.2 Command Operation

The SATA controller maintains a queue consisting of up to 16 commands. These commands can be distributed to a single attached device or, if the system contains a port multiplier, over each of the attached devices. It is possible to queue queued commands and non-queued commands into the SATA controller, provided the host software does not break protocol to any particular device (it is illegal to issue a non-queued command to a device that still has a queued command active as per ATAPI/ATA protocol).

15.2.1 Command Issue

When the host software is preparing to issue a command, it first builds a command descriptor as shown in Figure 15-29. The format of the command FIS is defined in the Serial ATA 2.5 standard shown in Figure 15-30. The software is also responsible for the creation of a scatter/gather list for data movement. This list should be defined to exactly match the transfer length as programmed into the command FIS. If the 16 entries are not sufficient, then an extended entry (ext) can be used to refer to an alternate table. When

the descriptor is built, the host software locates a free command slot within the SATA controller by examining the command queue register. To issue the command, the host software programs the address of the command descriptor and the attributes into the appropriate command header locations and then issues the command by writing the PMP and setting the appropriate CQ bit in the command queue register.

15.2.2 Command Service

After a command is issued, the SATA controller takes ownership of the command descriptor, transferring the command FIS to the targeted device when required, servicing the data transfer using the scatter/gather list provided and transferring the status back into the command descriptor (if programmed).

15.2.3 Command Completion Interrupt Timing

When a command completes, it is possible to enable the SATA controller to generate an interrupt. Associated with some commands there will be a command completion status FIS. The SATA controller will always transfer the status FIS to memory whether it indicates an error or good command completion.

15.2.4 DMA Context (Read Data)

When receiving FIS's from attached devices, the SATA controller has to support interleaving from various devices. Data FIS's from device 0 could be interleaved with data FIS's for device 1. In order to accomplish this, the SATA controller maintains in hardware a context for each command which is pushed onto and pulled from the DMA controller when needed to service the transfers.

15.2.5 DMA Context (Write Data)

When the SATA controller receives an FIS indicating that the next operation to a particular device should be a data write transfer, the SATA controller will lock the interface by forcing the link layer to transition to X_RDY immediately and not go through idle SYNC. This will mean that write transfers will not have to be interleaved, which simplifies the transmit data path and eliminates the need for a complex scheduler.

15.2.6 DMAT Primitive Processing

The SATA controller supports the reception of the DMAT primitive. When the SATA controller receives a DMAT primitive from the device, it will perform the following actions.

The DMA controller will complete the current read burst and transfer the data to the transport layer FIFO. The DMA controller signals an EOF on the last data of the burst, which causes the link layer to insert the CRC and EOF. The context for this transfer is returned to the context store. Once this action is completed, the device can terminate the transfer or re-initiate the transfer as per Serial ATA Revision 2.5 Section 9.4.4.

15.3 Command Layer Overview

The function of the SATA command layer is to allow host software queue commands. It then manages the command issue and service using context to complete the queued commands.

15.3.1 SATA Memory Map/Register Definition

Table 15-1 shows the memory map for the SATA registers. The offsets to the memory map table are defined for both SATA hosts. That is, SATA1 starts at 0x1_8000 address offset and SATA2 at 0x1_9000. Undefined 4-byte address spaces within offset 0x000–0xFFF are reserved.

NOTE

All registers (except SYSPR) described in this section and descriptors described in Section 15.3.6, “Command Header,” and Section 15.3.7, “Command Descriptor,” use little-endian byte ordering. Software running on the local processor in big-endian mode must byte-swap the data.

In this table, and in the register figures and field descriptions, the following access definitions apply:

- Reserved fields are always ignored for the purposes of determining access type.
- ‘R/W’, ‘R’, and ‘W’ (read/write, read only, and write only) indicate that all the non-reserved fields in a register have the same access type.
- ‘w1c’ indicates that all of the non-reserved fields in a register are cleared by writing ones to them.
- ‘Mixed’ indicates a combination of access types.

Table 15-1. SATA Register Summary

| Offset | Register | Access | Reset Value | Section/Page |
|--|---|--------|-------------|-----------------|
| SATA1—Block Base Address 0x1_8000 | | | | |
| SATA Command Registers | | | | |
| 0x000 | CQR—Command queue register | R/W | 0x0000_0000 | 15.3.2.1/15-5 |
| 0x008 | CAR—Command active register | R | 0x0000_0000 | 15.3.2.2/15-6 |
| 0x010 | CCR—Command completed register | w1c | 0x0000_0000 | 15.3.2.3/15-6 |
| 0x018 | CER—Command error register | w1c | 0x0000_0000 | 15.3.2.4/15-7 |
| 0x020 | DE—Device error register | w1c | 0x0000_0000 | 15.3.2.5/15-8 |
| 0x024 | CHBA—Command header base address | R/W | 0x0000_0000 | 15.3.2.6/15-8 |
| 0x028 | HStatus—Host status register | w1c | 0x2000_0000 | 15.3.2.7/15-9 |
| 0x02C | HControl—Host control register | Mixed | 0x0000_0100 | 15.3.2.8/15-12 |
| 0x030 | CQPMP—Port number queue register | R/W | 0x0000_0000 | 15.3.2.9/15-13 |
| 0x034 | SIG—Signature register | R | 0xFFFF_FFFF | 15.3.2.10/15-14 |
| 0x038 | ICC—Interrupt coalescing control register | R/W | 0x0100_0000 | 15.3.2.11/15-14 |
| SATA1 Superset Registers | | | | |
| 0x100 | SStatus—SATA interface status register | R | 0x0000_0000 | 15.3.3.1/15-15 |
| 0x104 | SError—SATA interface error register | w1c | 0x0000_0000 | 15.3.3.2/15-16 |

Table 15-1. SATA Register Summary (continued)

| Offset | Register | Access | Reset Value | Section/Page |
|---|--|--------|-------------|-----------------|
| 0x108 | SControl—SATA interface control register | R/W | 0x0000_0300 | 15.3.3.3/15-18 |
| 0x10C | SNotification—SATA interface notification register | w1c | 0x0000_0000 | 15.3.3.4/15-19 |
| SATA1 Control Status Registers | | | | |
| 0x140 | TransCfg—Transport layer configuration | R/W | 0x0800_0016 | 15.3.4.1/15-20 |
| 0x144 | TransStatus—Transport layer status | R | 0x0000_0000 | 15.3.4.2/15-21 |
| 0x148 | LinkCfg—Link layer configuration | R/W | 0x0000_FF34 | 15.3.4.3/15-21 |
| 0x14C | LinkCfg1—Link layer configuration1 | R/W | 0x0000_0000 | 15.3.4.4/15-22 |
| 0x150 | LinkCfg2—Link layer configuration2 | R/W | 0x0000_0000 | 15.3.4.5/15-23 |
| 0x154 | LinkStatus—Link layer status | R | 0x0000_0000 | 15.3.4.6/15-23 |
| 0x158 | LinkStatus1—Link layer status1 | R | 0x0000_0000 | 15.3.4.7/15-24 |
| 0x15C | PhyCtrlCfg1—PHY control configuration1 | R/W | 0x0000_3800 | 15.3.4.8/15-26 |
| 0x160 | CommandStatus—Link layer command status | R | 0x0000_0000 | 15.3.4.9/15-27 |
| 0x164– 0x3FC | Reserved | — | — | — |
| 0x400 | PhyCtrlCfg2—PHY control configuration2 | R/W | 0x0000_8000 | 15.3.4.10/15-29 |
| SATA1 System Control Register | | | | |
| 0x410 | SYSR—System priority register | R/W | 0x0000_0000 | 15.3.5.1/15-30 |
| 0x40C– 0xFFF | Reserved | — | — | — |
| SATA2—Block Base Address 0x1_9000 | | | | |
| SATA2 has the same memory-mapped registers that are described for SATA1; the offsets of SATA2 registers are the same except they have a different block base address. | | | | |
| SATA3 has the same memory-mapped registers that are described for SATA1; the offsets of SATA3 registers are the same except they have a different block base address. | | | | |
| SATA4 has the same memory-mapped registers that are described for SATA1; the offsets of SATA4 registers are the same except they have a different block base address. | | | | |

15.3.2 Command Registers

15.3.2.1 Command Queue Register (CQR)

Before queuing a command into the SATA controller, the CQR (shown in [Figure 15-2](#)) is first examined to detect a free command queue (CQ) slot. A free CQ slot is indicated by a 0 in a bit position. To queue a command, the bit corresponding to the CQ slot to use is set. At this point the SATA controller takes ownership of the command header space and command descriptor associated with the command slot. While the command is queued in the SATA controller or at the device, the command queue bit remains 1.

When the command completes, this bit is cleared to 0 by the hardware. For a device error, the CQR holds the command queue bits at 1 for each command queued or issued to the device in error. When the host software clears the device error, the hardware in turn clears each of the commands queued.




Figure 15-2. Command Queue Register (CQR)

Table 15-2 describes the CQR fields.

Table 15-2. CQR Field Descriptions

| Bit | Name | Description |
|-------|--------|-----------------------|
| 31–16 | — | Reserved |
| 15–0 | CQ_n | Command n queue bit |

15.3.2.2 Command Active Register (CAR)

When a command is issued from the SATA controller to the device, the command is marked as active by the hardware setting the appropriate command active bit of the CAR (shown in Figure 15-3). Once a command completes, the hardware clears the appropriate bit of the CAR.

For a device error, the CAR holds the command active bits at 1 for each command issued to the device in error. When the host software clears the device error, the hardware in turn clears each of the commands queued.




Figure 15-3. Command Active Register (CAR)

Table 15-3 describes the CAR fields.

Table 15-3. CAR Field Descriptions

| Bit | Name | Description |
|-------|-------|------------------------|
| 31–16 | — | Reserved |
| 15–0 | CAn | Command n active bit |

15.3.2.3 Command Completed Register (CCR)

When a command completes, the hardware sets the command completed bit for that command in the CCR (shown in Figure 15-4) to a 1. The hardware also clears both the command queue and the command active

bit for that command. When the software needs to acknowledge the reception of the command complete, it can do so in two ways:

- Writing a 1 to the command complete bit
- Issuing a command to the command slot

An interrupt coalescing scheme runs on the CCR. When the register contains a value other than 0x0000_0000, an interrupt coalescing timer runs. Each time a command completion is acknowledged, the timer is reset. When the timer times out, an interrupt is generated.




Figure 15-4. Command Completed Register (CCR)

Table 15-4 describes the CCR fields.

Table 15-4. CCR Field Descriptions

| Bit | Name | Description |
|-------|-----------------|--------------------------------|
| 31–16 | — | Reserved |
| 15–0 | CC _n | Command <i>n</i> completed bit |

15.3.2.4 Command Error Register (CER)

When a device errors a command by setting the error bit in the status register, this is detected by the SATA controller as a single device error. The associated command completing due to error is indicated by the hardware setting the command error bit for that command in the CER (shown in Figure 15-5). For safe operation under both command queuing and non-queuing operation, all commands queued into the SATA controller and at the device are considered aborted. The queue for that device is stopped. The values of the registers CQR, CAR, and CCR will allow the host software to know which commands have completed without error and those that were queued at the SATA controller and at the device.

When the host software clears the device error (by setting DER), the software is also responsible to clear CER by writing a 1 to the command error bit for the command that was in error. After the error condition at the device has been cleared, the host application software can reissue the commands to the SATA controller, which were aborted on the reception of the single device error.



Figure 15-5. Command Error Register (CER)

Table 15-8. HStatus Field Descriptions (continued)

| Bit | Name | Description |
|-------|---------|--|
| 29 | BE | BIST error. When the protocol is placed into BIST this bit maps the BIST error. 0 Indicates the link layer is passing BIST 1 Indicates that the link layer is not passing BIST When the protocol is not in BIST this bit will assert high and can be ignored. |
| 28–19 | — | Reserved |
| 18 | ME | SATA controller master error. Indicates if the host received an error on the CSB master interface during the access to memory. 0 No error response is received when a transfer was made into the memory 1 Error response is received during the transfer into the memory |
| 17 | PTx_Err | SATA controller parity Tx error. 0 No parity errors were detected on Tx data path 1 One or more parity errors were detected on Tx data path |
| 16 | PRx_Err | SATA controller parity Rx error. 0 No parity errors were detected on Rx data path 1 One or more parity errors were detected on Rx data path |
| 15–14 | — | Reserved |
| 13 | DUE | Data underrun. 0 No underrun encountered (data was retrieved from external memory in time to send a complete FIS) 1 The SATA controller encountered an underrun condition while sending the FIS |
| 12 | DOE | Data overrun. 0 No overrun condition encountered 1 The SATA controller encountered an overrun condition while receiving the FIS |
| 11 | CET | CRC error Tx. When set, this bit indicates that one or more CRC errors occurred in Tx data path. |
| 10 | CER | CRC error Rx. When set, this bit indicates that one or more CRC errors occurred in Rx data path. |
| 9 | FOT | FIFO overflow Tx. When set, this bit indicates that Tx FIFO is in overflow condition while sending FIS. |
| 8 | FOR | FIFO overflow Rx. When set, this bit indicates that Rx FIFO is in overflow condition while receiving FIS. |
| 7–6 | — | Reserved |
| 5 | FE | Fatal error. When set, this bit indicates that fatal error occurred in SATA controller. In this state, the interrupt will be generated if FATAL_INT is set in the host control register. Write '1' to clear the interrupt source. |
| 4 | PR | PHY ready. When set, this bit indicates that PHY READY signal was changed. In this state, the interrupt will be generated if PHYRDY_INT is set in the host control register. Write '1' to clear the interrupt source. |
| 3 | SIGU | Signature update. When set, this bit indicates that the signature is updated in the host signature register. In this state, the interrupt will be generated if SIG_INT is set in the host control register. Write '1' to clear the interrupt source. |
| 2 | SNTFU | SNotification update. When set, this bit indicates that the SNotification register has at least one bit set. In this state, the interrupt will be generated if SNTFY_INT is set in the host control register. Write '1' to clear the interrupt source. |

Table 15-8. HStatus Field Descriptions (continued)

| Bit | Name | Description |
|-----|------|---|
| 1 | DE | Device error. When set, this bit indicates that the DE register has at least one bit set. In this state, the interrupt will be generated if DE_INT is set in the host control register. Write '1' to clear the interrupt source. |
| 0 | CC | Command complete. When set, this bit indicates that the register CCR has at least one bit set. In this state, the interrupt will be generated if CC_INT is set in the host control register. Write '1' to clear the interrupt source. |

15.3.2.7.1 Error Processing

On single device error:

1. Examine the register DER to determine which device is in error state. There might be multiple devices in error.
2. Examine the register CER to determine which command was in error. The software knows which command belongs to which device.
3. Examine the status location of the descriptor of the command in error and determine the reason for the error.
4. If needed, the software should send commands to the device to clear down the error condition on device or for further examination of the device's status.
5. Clear the DER n bit by writing 1 to bit n , where n indicates the device in error. This will also clear out the outstanding commands for that device.
6. Clear the CER n bit by writing 1 to bit n , where n indicates the associated command in error. After that, the software can reissue command to the device if needed.

On fatal error:

1. Read the error register and other registers to determine how many commands are outstanding and how many have completed without error.
2. Bring the SATA controller offline. When this happens all queues within the SATA controller will be cleared.
3. Perform what corrective action the software determines is necessary.
4. Bring the SATA controller online. This will cause an out-of-bounds (OOB) to be run at the PHY level which will clear down any attached device.

15.3.2.8 Host Control Register (HControl)

HControl, shown in Figure 15-9, is written to control the operation of the SATA controller. To enable an interrupt, the associated bit must be set; to disable the interrupt, the associated bit must be cleared.

Offset 0x1_802C

Access: Mixed




Figure 15-9. Host Control (HControl) Register

Table 15-9 describes the HControl fields.

Table 15-9. HControl Field Descriptions

| Bit | Name | Description |
|-------|----------|---|
| 31 | HC_ON | Online/offline control. 0 Offline. Bring the SATA controller offline and place the PHY in reset 1 Online. Bring the SATA controller online |
| 30 | HC_OFF | Offline request status. 1 The SATA controller is currently completing an operation and will go offline when the operation completes. When this bit is set, the SATA controller can be forced to go offline by aborting its current operation by writing 0 to the HC_ON bit. |
| 29–15 | — | Reserved |
| 14–13 | — | Reserved |
| 12 | PT | Parity type select. 0 Select odd parity on data paths 1 Select even parity on data paths |
| 11 | — | Reserved |
| 10 | SNOOP_EN | Snoop enable during header fetch. 0 Snoop not enabled during command header fetch. 1 Snoop enabled during command header fetch. |
| 9 | PM_EN | Port multiplier attached. This bit is used to indicate if the HBA is attached to a port multiplier. This bit is set or cleared by software. 0 This SATA controller is directly attached to a SATA device. The SATA controller hardware does not auto-detect the presence of a port multiplier; this is to allow for future changes in signature type for the port multiplier. 1 A port multiplier is attached to the SATA controller. |
| 8 | SFIS_EN | Copy out status FIS. 0 If the command completes with a good status ERR = 0, do not copy out the FIS. 1 Always copy out the status FIS of the completing command. |

Table 15-9. HControl Field Descriptions (continued)

| Bit | Name | Description |
|-----|------------|---|
| 7-6 | — | Reserved. Reset value must be preserved when writing to the register. |
| 5 | FATAL_INT | Enable interrupt on fatal error. |
| 4 | PHYRDY_INT | Enable interrupt on PHY ready change. |
| 3 | SIG_INT | Enable interrupt signature update. |
| 2 | SNTFY_INT | Enable interrupt on SNotify register update. |
| 1 | DE_INT | Enable interrupt on single device error. |
| 0 | CC_INT | Enable interrupt on command complete. |

15.3.2.8.1 Bringing the SATA Controller Online/Offline

The HC_ON bit in HControl allows the host software to bring the SATA controller online or offline. The SATA controller online status should only be changed when there are no commands queued in the SATA controller or at any attached device.

When the host application wishes to bring the SATA controller offline it clears the HC_ON control bit. This acts as a request to the SATA controller to go offline. The SATA controller will signal it has completed this operation by clearing the HS_ON bit in the HStatus register. If any commands are outstanding at SATA controller or device then the SATA controller will wait for the operation to complete before going offline.

If the host application wishes to bring the SATA controller offline regardless of the queue status, it clears the HC_ON bit while the HS_OFF bit of the HStatus register is set.

When the host application wishes to bring the SATA controller online, it sets the HC_ON control bit. This acts as a request to the SATA controller to go online. The SATA controller will signal it has completed this operation by setting the HS_ON status bit.

15.3.2.9 Port Number Queue Register (CQPMP)

When queuing a command into the SATA controller, the CQPMP, shown in [Figure 15-10](#), is written with the value of the PMP field that addresses the device to which the command will be issued. If the device is directly attached (that is, there is no port multiplier in the system), then this register is not required and should be cleared.




Figure 15-10. Port Number Queue Register (CQPMP)

Table 15-10 describes the CQPMP fields.

Table 15-10. CQPMP Field Descriptions

| Bit | Name | Description |
|------|-------|-------------------------------------|
| 31–4 | — | Reserved |
| 3–0 | CQPMP | Command queue port multiplier field |

15.3.2.10 Signature Register (SIG)

The 32-bit SIG register, shown in Figure 15-11, contains the initial signature of an attached device when the first D2H register FIS is received from that device.




Figure 15-11. Signature Register (SIG)

Table 15-11 describes the SIG register fields.

Table 15-11. SIG Register Field Descriptions

| Bit | Name | Description |
|-------|----------|-----------------------|
| 31–24 | LBA_HIGH | LBA high register |
| 23–16 | LBA_MID | LBA mid register |
| 15–8 | LBA_LOW | LBA low register |
| 7–0 | SEC_CNT | Sector count register |

15.3.2.11 Interrupt Coalescing Control Register (ICC)

When a command completes, the SATA controller sets the corresponding bit in the command completed register. The interrupt coalescing scheme runs on the SIG register, shown in Figure 15-12. The scheme runs in two ways:

- If the number of completed commands exceeds the threshold, then the interrupt will be signaled.
- If any command complete bit has been set in the register for a number of HCLK ticks equal to the programmed count, then the interrupt will be set. This timer will be reset each time a command completion is acknowledged by the application layer software.




Figure 15-12. Interrupt Coalescing Control Register (ICC)

Table 15-12 describes the ICC fields.

Table 15-12. ICC Field Descriptions

| Bit | Name | Description |
|-------|-------|--|
| 31–29 | — | Reserved |
| 28–24 | ITC | Interrupt threshold count. The number of command completions that will raise the interrupt. 00000 Implied no threshold and the interrupt will be signaled based on the threshold timer. 00001, 01111 The number of command complete bits which, if set, will cause the interrupt to be signaled. |
| 23–19 | — | Reserved |
| 18–0 | ITTCV | Interrupt threshold timer compare value. A value of 0 indicates that whenever a command complete bit is set the interrupt should be signaled. |

15.3.3 SATA Superset Registers

Serial ATA provides an additional block of registers to control the interface and to retrieve interface state information.

15.3.3.1 SATA Interface Status Register (SStatus)

SStatus, shown in Figure 15-13, is a 32-bit read-only register that conveys the current state of the interface and host adapter. The register conveys the interface state at the time it is read and is updated continuously and asynchronously by the host adapter. Writes to this register have no effect.



Figure 15-13. SATA Interface Status Register (SStatus)

Table 15-13 describes the SStatus fields.

Table 15-13. SStatus Field Descriptions

| Bit | Name | Description |
|-------|------|--|
| 31–12 | — | Reserved |
| 11–8 | IPM | Interface power management state. Indicates the current interface power management state. 0000 Device not present or communication not established 0001 Interface in active state 0010 Interface in partial power management state 0110 Interface in slumber power management state All other values reserved |

Table 15-13. SStatus Field Descriptions (continued)

| Bit | Name | Description |
|-----|------|---|
| 7–4 | SPD | Speed. Indicates the negotiated interface communication speed established. 0000 No negotiated speed (device not present or communication not established) 0001 First-generation communication rate negotiated 0010 Second-generation communication rate negotiated All other values reserved |
| 3–0 | DET | Detection. Indicates the interface device detection and PHY state. 0000 No device detected and PHY communication not established 0001 Device presence detected but PHY communication not established 0011 Device presence detected and PHY communication established 0100 PHY in offline mode as a result of the interface being disabled or running in a BIST loopback mode All other values reserved |

15.3.3.2 SATA Interface Error Register (SError)

SError, shown in Figure 15-14, is a 32-bit register that conveys supplemental interface error information to complement the error information available in the shadow register block error register. The register represents all the detected errors accumulated since the last time the SError register was cleared (whether recovered by the interface or not). Set bits in the error register are explicitly cleared by a write operation to the SError register or by a reset operation. The error bits that have been set in this register are cleared by writing a 1 to the corresponding field. Host software should clear the interface SError register at appropriate checkpoints in order to best isolate error conditions and the commands they impact.

Bits 31–16 of this register represent the DIAG decode bits, which contain diagnostic error information, for use by diagnostic software in validating correct operation or isolating failure modes. Bits 15–0 represent the ERR decode bits, which contain information for use by the host software in determining the appropriate response to the error condition.

Offset 0x1_8104

Access: w1c




Figure 15-14. SATA Interface Error Register (SError)

Table 15-14 describes the SError field descriptions.

Table 15-14. SError Field Descriptions

| Bit | Name | Description |
|--------------------|------|---|
| DIAG Decode | | |
| 31–28 | — | Reserved bit for future use. Should be cleared. |

Table 15-14. SError Field Descriptions (continued)

| Bit | Name | Description |
|-------------------|------|---|
| 27 | A | Port selector presence detected. This bit is set when COMWAKE is received while the host is in state HP2: HR_AwaitCOMINIT. On power-up reset this bit is cleared. The bit is cleared when the host writes a 1 to this bit location. |
| 26 | X | Exchanged. When set to 1 this bit indicates that device presence has changed since the last time this bit was cleared. The means by which the implementation determines that the device presence has changed is vendor specific. This bit may be set anytime a PHY reset initialization sequence occurs as determined by reception of the COMINIT signal, whether in response to a new device being inserted, to a COMRESET having been issued, or to power-up. |
| 25 | F | Unrecognized FIS type. When set to 1, this bit indicates that since the bit was last cleared, one or more FIS's were received by the transport layer with good CRC, but they had a type field that was not recognized. |
| 24 | — | Reserved. |
| 23 | S | Link sequence error. When set to 1, this bit indicates that one or more link state machine error conditions were encountered since the last time this bit was cleared. The link layer state machine defines the conditions under which the link layer detects an erroneous transition. |
| 22 | H | Handshake error. When set to 1, this bit indicates that one or more R_ERRP handshake responses were received in response to frame transmission. Such errors may be the result of a CRC error detected by the recipient, of a disparity or 10b/8b decoding error, or of other error conditions leading to a negative handshake on a transmitted frame. |
| 21 | C | CRC error. When set to 1, this bit indicates that one or more CRC errors occurred with the link layer since the bit was last cleared. |
| 20 | D | Disparity error. When set to 1, this bit indicates that incorrect disparity was detected one or more times since the last time the bit was cleared. |
| 19 | B | 10b to 8b decode error. When set to 1, this bit indicates that one or more 10-bit to 8-bit decoding errors occurred since the bit was last cleared. |
| 18 | W | COMWAKE detected. When set to 1, this bit indicates that a COMWAKE signal was detected by the PHY since the last time this bit was cleared. |
| 17 | IN | PHY internal error. When set to 1, this bit indicates that the PHY detected some internal error since the last time this bit was cleared. |
| 16 | N | PHYRDY change. When set to 1, this bit indicates that the PHYRDY signal changed state since the last time this bit was cleared. |
| ERR Decode | | |
| 15–12 | — | Reserved bit for future use; should be cleared. |
| 11 | E | E Internal error. The host bus adapter experienced an internal error that caused the operation to fail and may have put the host bus adapter into an error state. Host software should reset the interface before retrying the operation. If the condition persists, the host bus adapter may suffer from a design issue rendering it incompatible with the attached device. |
| 10 | — | Reserved. |

Table 15-14. SError Field Descriptions (continued)

| Bit | Name | Description |
|-----|------|---|
| 9 | C | Non-recovered persistent communication or data integrity error. A communication error that was not recovered occurred that is expected to be persistent. Because the error condition is expected to be persistent, the operation need not be retried by the host software. Persistent communications errors may arise from faulty interconnect with the device, from a device that has been removed or has failed, or a number of other causes. |
| 8 | T | Non-recovered transient data integrity error: A data integrity error occurred that was not recovered by the interface. Because the error condition is not expected to be persistent, the operation should be retried by the host software. |
| 7–2 | — | Reserved |
| 1 | M | Recovered communications error. Communications between the device and host were temporarily lost but were re-established. This can arise from a device temporarily being removed, from a temporary loss of PHY synchronization, or from other causes, and may be derived from the PHYRDY _n signal between the PHY and link layers. No action is required by the host software, because the operation ultimately succeeded. However, the host software may elect to track such recovered errors to gauge overall communications integrity and potentially step down the negotiated communication speed. |
| 0 | ITG | Recovered data integrity error. A data integrity error occurred that was recovered by the interface through a retry operation or other recovery action. This can arise from a noise burst in the transmission, a voltage supply variation, or other causes. No action is required by host software, because the operation ultimately succeeded. However, the host software may elect to track such recovered errors to gauge overall communications integrity and potentially step down the negotiated communication speed. |

15.3.3.3 SATA Interface Control Register (SControl)

SControl, shown in [Figure 15-15](#), is a 32-bit read-write register that provides the interface by which software controls SATA interface capabilities. Writes to the SControl register result in an action being taken by the host adapter or interface. Reads from the register return the last value written to it.




Figure 15-15. SATA Interface Control Register (SControl)

Table 15-15 describes the SControl fields.

Table 15-15. SControl Field Descriptions

| Bit | Name | Description |
|-------|------|---|
| 31–16 | — | Reserved, should be cleared. |
| 15–12 | SPM | Select power management. Used to select a power management state. A non-zero value written to this field will cause the power management state specified to be initiated. A value written to this field is treated as a one-shot. 0000 No power management state transition requested 0001 Transition to the partial power management state initiated 0010 Transition to the slumber power management state initiated 0100 Transition to the active power management state initiated All other values reserved |
| 11–8 | IPM | Interface power management. The enabled interface power management states can be invoked via the SATA interface power management capabilities. 0000 No interface power management state restrictions 0001 Transitions to the partial power management state disabled 0010 Transitions to the slumber power management state disabled 0011 Transitions to both the partial and slumber power management states disabled All other values reserved |
| 7–4 | SPD | Speed. Highest allowed communication speed the interface is allowed to negotiate when interface communication speed is established. 0000 No speed negotiation restrictions 0001 Limit speed negotiation to a rate not greater than first-generation communication rate 0010 Limit speed negotiation to a rate not greater than second-generation communication rate All other values reserved |
| 3–0 | DET | Detection. Controls the host adapter device detection and interface initialization. 0000 No device detection or initialization action requested 0001 Perform interface communication initialization sequence to establish communication. This is functionally equivalent to a hard reset and results in the interface being reset and communications re-initialized. Upon a write to the SControl register that sets the DET field to 0001, the host interface should transition to the HP1: HR_Reset [Delete space after state and should remain in that state until the DET field is set to a value other than 0001 by a subsequent write to the SControl register. 0100 Disable the SATA interface and put PHY in offline mode All other values reserved |

15.3.3.4 SATA Interface Notification Register (SNotification)

SNotification, shown in Figure 15-16, is a 32-bit, write-one-to-clear register that conveys the devices that have sent the host a set device bits FIS with the notification bit. When the host receives a set device bits FIS with the notification bit set to 1, the host should set the bit in SNotification corresponding to the value of the PM port field in the received FIS. For example, if the PM port field is set to 7 then the host should clear bit 7 by writing a 1 to it. Next, the host should generate an interrupt if the I bit of the set device bits FIS is set to 1 and interrupts are enabled.

In this register, bits previously set are explicitly cleared by a write operation or by a power-on-reset operation. If the register is not cleared due to a COMRESET, the software is responsible for clearing the register as appropriate.

Table 15-19. LinkCfg Field Descriptions (continued)

| Bit | Name | Description |
|-------|--------------|--|
| 25–16 | PRT | PHY ready timer. These ten bits specify the timeout value of the PHY_READY timer. If EN_PHY_TO is set, the link layer will count down on every rising edge of scanTxClk, as long as PHY_READY is de-asserted. When the counter reaches 0, a PHY_RESET will be issued to the PHY to try and re-establish communications with the far end. The timer is initially loaded with a value equal to the concatenation of {PHY_READY_TIMER, 9b0_0000_0000}. |
| 15–8 | AR | Align insertion rate. The SATA specification requires that the link layer send a pair of ALIGN primitives at least every 254 words of data. This is achieved by setting ALIGN_RATE to '11111111'. However, for test purposes it is possible to send ALIGNs at a higher rate. This can be achieved by setting ALIGN_RATE to a lower value (that is, ALIGN_RATE-1); words will be sent by the link layer between each set of ALIGN primitive pairs. Note: If SEND_4_ALIGNs is set, one should not set the ALIGN_RATE to be four or less. If SEND_4_ALIGNs is not set, one should not set the ALIGN_RATE to be two or less. |
| 7 | EPNRT | Enable PHY not ready timer. If PHY_READY is de-asserted for a length of time, as specified by CFG_PHY_READY_TIMER, then this bit, when asserted, enables the link layer to re-issue a PHY_RESET, thereby re-initiating OOB. |
| 6 | S4A | Send four ALIGNs. When asserted, four ALIGN primitives are transmitted at the specified rate, instead of the normal two ALIGNs. |
| 5 | RX_SCR_EN | Rx scramble enable. If this bit is asserted, then descrambling of the receive data is enabled as per the SATA specification. |
| 4 | TX_SCR_EN | Tx scramble enable. If this bit is asserted, then scrambling of the transmit data is enabled as per the SATA specification. |
| 3 | TX_PRIM_JUNK | TX prim junk. If this bit is de-asserted, then scrambled junk data is sent after a CONT primitive, as per the SATA specification. If this bit is asserted, then the single character 0xDEADBEEF is sent continuously instead. This is to aid debug. |
| 2 | TX_CONT_EN | TX CONT. If this bit is asserted, then the transmission of CONT primitives is enabled. If de-asserted, then long sequences of repeated primitives can be sent by the link layer. |
| 1 | RX_BAD_CRC | RX bad CRC. When a rising edge is detected on this bit, it causes a bad CRC to be detected for the current frame. This bit has to be toggled from a 0 to a 1 to enable this feature. |
| 0 | TX_BAD_CRC | Tx bad CRC. A bad CRC (inverted value of the correct CRC) value will be transmitted for one FIS only by the link layer when a rising edge is detected on this signal. This bit has to be toggled from a 0 to a 1 to enable this feature. |

15.3.4.4 Link Layer Configuration Register1 (LinkCfg1)

LinkCfg1, shown in Figure 15-20, controls the configuration of the link layer.




Figure 15-20. Link Layer Configuration Register1 (LinkCfg1)

Table 15-20 describes the LinkCfg1 fields.

Table 15-20. LinkCfg1 Field Descriptions

| Bit | Name | Description |
|------|----------------|--|
| 31–7 | — | Reserved |
| 6 | CD | This bit specifies whether the data used during the primitive override should be a data character or a primitive. For example, if CD = 1, PRIM_OVR_STATE = L_SendEOF and PRIM = WTRM, then a WTRM primitive will be inserted into the datastream instead of an EOF (whenever a rising edge is seen on PRIM_OVR_EN). If CD = 0, then a normal data character (as specified by PRIM) is inserted into the datastream instead of the EOF. |
| 5–0 | PRIM_OVR_STATE | Prim override state. These 6 bits are used in the primitive override debug functionality. When the link layer detects a positive edge on PRIM_OVR_EN, it overrides the next primitive that would be inserted during the PRIM_OVR_STATE, with the data specified by the PRIM and CD configuration bits. |

15.3.4.5 Link Layer Configuration Register2 (LinkCfg2)

LinkCfg2, shown in Figure 15-21, controls the configuration of the link layer.




Figure 15-21. Link Layer Configuration Register1 (LinkCfg1)

Table 15-21 describes the LinkCfg2 fields.

Table 15-21. LinkCfg2 Field Descriptions

| Bit | Name | Description |
|------|------|--|
| 31–0 | PRIM | This 32-bit bus specifies the data to be used in the overriding primitive debug logic, described in the definition of LinkCfg1 register. |

15.3.4.6 Link Layer Status Register (LinkStatus)

LinkStatus, shown in Figure 15-22, indicates the status of the link layer.




Figure 15-22. Link Layer Status Register (LinkStatus)

Table 15-22 describes the LinkStatus fields.

Table 15-22. LinkStatus Field Descriptions

| Bit | Name | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|--------------------|------------------------|--|-------------|--------------------|------------|----------------|-------------------|----------------|-------------------|------------------|------------------|----------------|------------------|---------------|-------------------|---------------|-------------|------------------|--------------|-------------|-----------------|----------------|---------------|----------------|------------------|----------------|----------------|----------------|-----------------|----------------|--------------------|------------------------|----------------|------------------------|-----------------|------------------------|----------------|------------------------|----------------|------------------------|-------------|------------------------|------------------|------------|------------|--|
| 31–6 | — | Reserved | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 5–0 | LINK_STATE | Current value of the link layer state machine at the time the LinkStatus register is read. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | <table border="0"> <tr> <td>L_Reset = 0</td> <td>L_NoCommPower = 21</td> </tr> <tr> <td>L_Idle = 1</td> <td>L_WakeUp1 = 22</td> </tr> <tr> <td>HL_SendChkRdy = 2</td> <td>L_WakeUp2 = 23</td> </tr> <tr> <td>DL_SendChkRdy = 3</td> <td>L_RcvChkRdy = 24</td> </tr> <tr> <td>L_TPMPartial = 4</td> <td>L_RcvData = 25</td> </tr> <tr> <td>L_TPMSlumber = 5</td> <td>L_BadEnd = 26</td> </tr> <tr> <td>L_RcvWaitFifo = 6</td> <td>L_RcvEOF = 27</td> </tr> <tr> <td>L_PMOff = 7</td> <td>L_SendHoldA = 28</td> </tr> <tr> <td>L_PMDeny = 8</td> <td>L_Hold = 29</td> </tr> <tr> <td>L_NoCommErr = 9</td> <td>L_GoodCRC = 30</td> </tr> <tr> <td>L_NoComm = 10</td> <td>L_GoodEnd = 31</td> </tr> <tr> <td>L_SendAlign = 11</td> <td>L_PMOff_2 = 32</td> </tr> <tr> <td>L_SendSOF = 12</td> <td>L_PMOff_3 = 33</td> </tr> <tr> <td>L_SendData = 13</td> <td>L_PMOff_4 = 34</td> </tr> <tr> <td>WAIT_FOR_SYNC = 14</td> <td>WAIT_PMACK_SENT_1 = 35</td> </tr> <tr> <td>L_SendCRC = 15</td> <td>WAIT_PMACK_SENT_2 = 36</td> </tr> <tr> <td>L_SendHold = 16</td> <td>WAIT_PMACK_SENT_3 = 37</td> </tr> <tr> <td>L_RcvHold = 17</td> <td>WAIT_PMACK_SENT_4 = 38</td> </tr> <tr> <td>L_SendEOF = 18</td> <td>WAIT_PMACK_SENT_5 = 39</td> </tr> <tr> <td>L_Wait = 19</td> <td>WAIT_PMACK_SENT_6 = 40</td> </tr> <tr> <td>L_ChkPhyRdy = 20</td> <td>BIST0 = 41</td> </tr> <tr> <td>BIST1 = 42</td> <td></td> </tr> </table> | L_Reset = 0 | L_NoCommPower = 21 | L_Idle = 1 | L_WakeUp1 = 22 | HL_SendChkRdy = 2 | L_WakeUp2 = 23 | DL_SendChkRdy = 3 | L_RcvChkRdy = 24 | L_TPMPartial = 4 | L_RcvData = 25 | L_TPMSlumber = 5 | L_BadEnd = 26 | L_RcvWaitFifo = 6 | L_RcvEOF = 27 | L_PMOff = 7 | L_SendHoldA = 28 | L_PMDeny = 8 | L_Hold = 29 | L_NoCommErr = 9 | L_GoodCRC = 30 | L_NoComm = 10 | L_GoodEnd = 31 | L_SendAlign = 11 | L_PMOff_2 = 32 | L_SendSOF = 12 | L_PMOff_3 = 33 | L_SendData = 13 | L_PMOff_4 = 34 | WAIT_FOR_SYNC = 14 | WAIT_PMACK_SENT_1 = 35 | L_SendCRC = 15 | WAIT_PMACK_SENT_2 = 36 | L_SendHold = 16 | WAIT_PMACK_SENT_3 = 37 | L_RcvHold = 17 | WAIT_PMACK_SENT_4 = 38 | L_SendEOF = 18 | WAIT_PMACK_SENT_5 = 39 | L_Wait = 19 | WAIT_PMACK_SENT_6 = 40 | L_ChkPhyRdy = 20 | BIST0 = 41 | BIST1 = 42 | |
| L_Reset = 0 | L_NoCommPower = 21 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_Idle = 1 | L_WakeUp1 = 22 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| HL_SendChkRdy = 2 | L_WakeUp2 = 23 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| DL_SendChkRdy = 3 | L_RcvChkRdy = 24 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_TPMPartial = 4 | L_RcvData = 25 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_TPMSlumber = 5 | L_BadEnd = 26 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_RcvWaitFifo = 6 | L_RcvEOF = 27 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_PMOff = 7 | L_SendHoldA = 28 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_PMDeny = 8 | L_Hold = 29 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_NoCommErr = 9 | L_GoodCRC = 30 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_NoComm = 10 | L_GoodEnd = 31 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_SendAlign = 11 | L_PMOff_2 = 32 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_SendSOF = 12 | L_PMOff_3 = 33 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_SendData = 13 | L_PMOff_4 = 34 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| WAIT_FOR_SYNC = 14 | WAIT_PMACK_SENT_1 = 35 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_SendCRC = 15 | WAIT_PMACK_SENT_2 = 36 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_SendHold = 16 | WAIT_PMACK_SENT_3 = 37 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_RcvHold = 17 | WAIT_PMACK_SENT_4 = 38 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_SendEOF = 18 | WAIT_PMACK_SENT_5 = 39 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_Wait = 19 | WAIT_PMACK_SENT_6 = 40 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| L_ChkPhyRdy = 20 | BIST0 = 41 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| BIST1 = 42 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

15.3.4.7 Link Layer Status Register1 (LinkStatus1)

LinkStatus1, shown in Figure 15-23, indicates the status of the link layer.




Figure 15-23. Link Layer Status Register1 (LinkStatus1)

Table 15-23 describes the LinkStatus1 fields.

Table 15-23. LinkStatus1 Field Descriptions

| Bit | Name | Description |
|-------|-------|--|
| 31–24 | KEGC | Kchar error gray count. The number of words received from the PHY, where one or more control character errors have been detected. A value of 255 indicates an error count of 255 or more as this counter does not wrap around to 0. The count value is updated with its current value each time the Status1 register is read. The count is represented in gray code. |
| 23–16 | PIEGC | PHY internal error gray count. The number of words received from the PHY, where one or more internal errors have been detected. A value of 255 indicates an error count of 255 or more as this counter does not wrap around to 0. The count value is updated with its current value each time the Status1 register is read. The count is represented in gray code. |
| 15–8 | CEGC | Code error gray count: The number of words received from the PHY, where one or more code errors have been detected. A value of 255 indicates an error count of 255 or more as this counter does not wrap around to 0. The count value is updated with its current value each time the Status1 register is read. The count is represented in gray code. |
| 7–0 | DEGC | Disparity error gray count. The number of words received from the PHY, where one or more disparity errors have been detected. A value of 255 indicates an error count of 255 or more as this counter does not wrap around to 0. The count value is updated with its current value each time the Status1 register is read. The count is represented in gray code. |

Sample C code to convert gray counts to binary:

```
int Gray2Binary(int gray)
{
int ish;
unsigned long ans, idiv;
ish=1; This is the more complicated direction: In hierarchical stages, starting with a one-bit
right shift, cause each bit to be XORed with all more significant bits.
Ans=gray;
for (;;)
{
ans ^= (idiv=ans >> ish);
if (idiv <= 1 || ish == 16) return ans;
ish <<= 1; Double the amount of shift on the next cycle.
}
}
```


Table 15-24. PhyCtrlCf1 Field Descriptions (continued)

| Bit | Name | Description |
|-----|-------------|---|
| 8–7 | FPRFTI | <p>Force PHY ready, force Tx idle. This pair of signals determines how phyRdy is driven, how the output buffer IDLE condition is controlled and how disparity errors in ALIGN primitives should be tolerated during OOB. The IDLE condition is defined in SATA as both traces of the transmit differential pair being driven to common mode.</p> <ul style="list-style-type: none"> • frcPhyRdy = 0 • frcTxIdle = 0 <p>In this mode phyRdy and Tx buffer IDLE control driven by OOB state machine. Disparity errors in ALIGN primitives are not tolerated during OOB.</p> <ul style="list-style-type: none"> • frcPhyRdy = 0 • frcTxIdle = 1 <p>In this mode phyRdy and Tx buffer IDLE control driven by OOB state machine. Disparity errors in ALIGN primitives are tolerated during OOB.</p> <ul style="list-style-type: none"> • frcPhyRdy = 1 • frcTxIdle = 0 <p>In this mode phyRdy is asserted high and Tx buffer IDLE control is forced off, causing the output buffer to be enabled. Tolerance of disparity errors in ALIGN primitives is of no consequence, because OOB is bypassed.</p> <ul style="list-style-type: none"> • frcPhyRdy = 1 • frcTxIdle = 1 <p>In this mode phyRdy is asserted high and Tx buffer IDLE control is forced on, causing the output buffer to the IDLE condition. Tolerance of disparity errors in ALIGN primitives is of no consequence as OOB is bypassed.</p> |
| 6–4 | REF_CLK_SEL | <p>Reference clock frequency select. This field defines the PHY's input reference clock frequency.</p> <p>000 75 MHz input frequency 010 150 MHz input frequency 101 50 MHz input frequency 110 100 MHz input frequency 111 125 MHz input frequency All other values are reserved.</p> <p>Note: This field is available only for SATA1.</p> |
| 3–1 | LPB_EN | <p>Loopback enable. These bits control both loopback modes and power management modes.</p> <p>000 No loopback and in normal power mode 001 Far end re-timed (parallel) loopback enabled 010 Near end analog (serial) loopback enabled 011 Invalid 100 Invalid 101 goPartial. This encoding results in the OOB state machine entering the partial state. Note that in the PCS, partial and slumber have the same effect. 110 goSlumber. This encoding results in the OOB state machine entering the slumber state. Note that in the PCS, partial and slumber have the same effect. 111 Invalid</p> <p>Note: This field is available only for SATA1.</p> |
| 0 | BIST_EN | <p>PHY BIST enable. Built-in self test mode for the PCS and SerDes. When asserted high, the PCS transmitter continually sends a 256-byte incrementing data pattern, and the PCS receiver signals correct reception of the test pattern by asserting bistErr pin low.</p> |

15.3.4.9 Link Layer Command Status Register (CommandStatus)

CommandStatus, shown in [Figure 15-25](#), indicates the status of the command layer.

Offset 0x1_8160

Access: Read only




Figure 15-25. Link Layer Command Status Register (CommandStatus)

Table 15-25 describes the CommandStatus fields.

Table 15-25. CommandStatus Field Descriptions

| Bit | Name | Description | | |
|-------|-------------|-----------------------|-------------------|---------------------------|
| 31–24 | — | Reserved | | |
| 23–20 | DMATCmdNo | | | |
| 19–16 | DMATPMPNo | | | |
| 15–12 | RecPMPLocal | | | |
| 11–8 | CSSState | CSSIdle = 0x0 | CSSNPS = 0x4 | CSSPF = 0x7 |
| | | CSSGP = 0x1 | CSSSC = 0x5 | CSSTO = 0x8 |
| | | CSSPCA = 0x2 | CSSNCS = 0x6 | CSSWCI = 0x9 |
| 7–0 | CMState | CMIdle = 0x00 | CMSDBCCT = 0x13 | CMSATAPI = 0x26 |
| | | CMFatalError = 0x01 | CMSDBPRT = 0x14 | CMWFC = 0x27 |
| | | CMWE = 0x02 | CMSDBNT = 0x15 | CMDF = 0x28 |
| | | CMRF = 0x03 | CMSDBLT = 0x16 | CMDFWD = 0x29 |
| | | CMSNDF = 0x04 | CMSDBFCC = 0x17 | CMDFWDW = 0x2A |
| | | CMWSNDF = 0x05 | CMSDBNCC = 0x18 | CMDFWGRC = 0x2B |
| | | CMWSNDFWUCA = 0x06 | CMSDBLCC = 0x19 | CMDFBSY = 0x2C |
| | | CMCIWNE = 0x07 | CMSDBWFT = 0x1A | CMDMAA = 0x2D |
| | | CMCIWIF = 0x08 | CMRDMA = 0x1B | CMDMAAWFTF = 0x2E |
| | | CMCISNDS = 0x09 | CMRDMASTC = 0x1C | CMDMAADW = 0x2F |
| | | CMCIWDMAC = 0x0A | CMRDMASTNC = 0x1D | CMSD = 0x30 |
| | | CMRUF = 0x0B | CMRDMASTLC = 0x1E | CMDC = 0x31 |
| | | CMRUFUS = 0x0C | CMRDMASTT = 0x1F | CMWDC = 0x32 |
| | | CMRUFWUS = 0x0D | CMRDMASTL = 0x20 | CMWU = 0x33 |
| | | CMRWSU = 0x0E | CMRDMAWFTF = 0x21 | CMWRFD = 0x34 |
| | | CMRUFWMW = 0x0F | CMRDMAWDW = 0x22 | CMWCC = 0x35 |
| | | CMSDB = 0x10 | CMPIOS = 0x23 | CMRUNF = 0x36 |
| | | CMSDBWSN = 0x11 | CMPIOSWFTF = 0x24 | CMRUNFC = 0x37 |
| | | CMSDBCcleanACK = 0x12 | CMPIOSDW = 0x25 | CMFatalErrorUpdate = 0x38 |

15.3.4.10 PHY Control Configuration Register2 (PhyCtrlCfg2)

PhyCtrlCfg2, shown in Figure 15-26, provides user control for the SATA PHY. Note that the bit ordering is 0–31, rather than 31–0 of the other registers in this chapter.




Figure 15-26. PHY Control Configuration Register2 (PhyCtrlCfg2)

Table 15-24 describes the PhyCtrlCfg2 fields.

Table 15-26. PhyCtrlCfg2 Field Descriptions

| Bit | Name | Description |
|-------|--------------|---|
| 0 | PHY_RST | SATA PHY reset. This active high reset behaves like a power on reset for SATA PHY. This is used to reset all data paths, OOB processor, and PLLs. This should be used if the reference clock value needs to be changed from its default value (75 MHz) in PhyCtrlCfg1 register. Note: This bit is available only for SATA1. |
| 1–15 | — | Reserved |
| 16 | PHY_MODE | SATA PHY mode select. 0 “m” mode 1 “i” mode Note: This bit is available only for SATA1. |
| 17–23 | — | Reserved |
| 24–25 | TX_AMP_CNTRL | Tx amplitude control. 00 No transmit amplitude control 01 Transmit amplitude control of +6% 10 Transmit amplitude control of +10% 11 Transmit amplitude control of +16% Note: This bit is available only for SATA1. |

Table 15-26. PhyCtrlCfg2 Field Descriptions (continued)

| Bit | Name | Description |
|-------|---------------|---|
| 26–28 | TX_PRE_EMP_G2 | Transmit pre-emphasis for GEN2 speeds. 000 0.0% pre-emphasis 001 3.7% pre-emphasis 010 7.4% pre-emphasis 011 11.1% pre-emphasis 100 14.8% pre-emphasis 101 18.5% pre-emphasis 110 22.2% pre-emphasis 111 Reserved |
| 29–31 | TX_PRE_EMP_G1 | Transmit pre-emphasis for GEN1 speeds. 000 0.0% pre-emphasis 001 3.7% pre-emphasis 010 7.4% pre-emphasis 011 11.1% pre-emphasis 100 14.8% pre-emphasis 101 18.5% pre-emphasis 110 22.2% pre-emphasis 111 Reserved |

15.3.5 System Control Registers

15.3.5.1 System Priority Register (SYSPR)

SYSPR, shown in [Figure 15-27](#), can be used to control various settings that affect the system response to DMA operations. Note that the bit ordering is 0–31, rather than 31–0 of the other registers in this chapter.




Figure 15-27. System Priority Register (SYSPR)

Table 15-27 describes the SYSPR fields.

Table 15-27. SYSPR Field Descriptions

| Bits | Name | Description |
|-------|----------|--|
| 0–15 | — | Reserved |
| 16–17 | PRI_DATA | Priority. This field will be used to present priority level for CSB arbitration for SATA controller’s dma requests. Bits 16-17 will be used when the request belongs to data transfer. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) |
| 18–19 | PRI_DES | Priority. This field will be used to present priority level for CSB arbitration for the SATA controller’s DMA requests. Bits 18–19 will be used when the request belongs to descriptor fetch. 00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority) |
| 10–28 | — | Reserved |
| 29 | RD_SAFE | Read safe. This bit should be set only if the target of the read DMA operation is a well behaved memory that is not affected by the read operation and which will return the same data if read again from the same location. This means that unaligned reading operation can be rounded up to enable more efficient read operations. 0 It is not safe to read more bytes that were intended. 1 It is safe to read more bytes that were intended. |

15.3.6 Command Header

Each entry in the command header table consists of the structure shown in [Figure 15-28](#).

NOTE

In this chapter, “word” refers to 4 bytes or 32 bits.

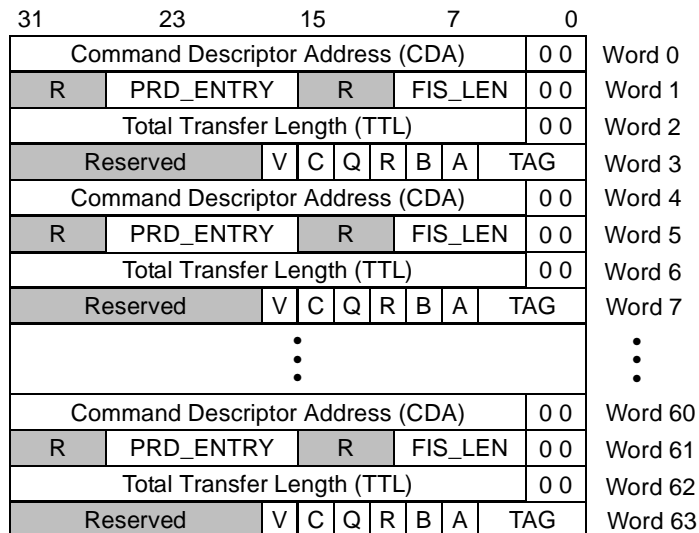


Figure 15-28. Command Header

Table 15-28 shows the command header word 0—data base address.

Table 15-28. Command Header Word 0—Data Base Address

| Bit | Name | Description |
|------|------|---|
| 31–2 | CDA | Command descriptor base address. Indicates the 32-bit physical address of the command descriptor block. The block must be word-aligned, indicated by bits 1–0 being reserved. |
| 1–0 | — | Reserved |

Table 15-29 shows the command header word 1—FIS_LEN.

Table 15-29. Command Header Word 1—FIS_LEN

| Bit | Name | Description |
|-------|-----------|---|
| 31–22 | — | Reserved |
| 21–16 | PRD_ENTRY | Number of PRD entries including indexed entries. |
| 7 | — | Reserved |
| 6–2 | FIS_LEN | FIS length. This is a 5-bit word count of the total length of the control or vendor-specific FIS to transfer. |
| 1–0 | — | Reserved |

Table 15-30 shows the command header word 2—data base address.

Table 15-30. Command Header Word 2—Data Base Address

| Bit | Name | Description |
|------|------|---|
| 31–2 | TTL | Total transfer length. This is a 30-bit word count of the total length of the data transfer. It is used to detect overruns/underruns between the transfer lengths programmed in the command and the PRDT. |
| 1–0 | — | Reserved |

Table 15-31 shows the command header word 3—description information.

Table 15-31. Command Header Word 3—Description Information

| Bit | Name | Description |
|-------|------|--|
| 31–12 | — | Reserved |
| 11 | — | Reserved, should be 1. |
| 10 | V | Vendor BIST. When this bit is set, it indicates that the command is a Vendor BIST, thus FIS will loop back at the PHY local test. |
| 9 | C | Snoop enable during all descriptor read/write operations associated with this command. |
| 8 | Q | Queued. Command is an FPDMA queued command. |
| 7 | R | Reset. The command is a SRST or device reset. |
| 6 | B | BIST. The command will require the host to enter BIST mode. |
| 5 | A | ATAPI command. The command is an ATAPI command and thus will require that the host uses the ATAPI portion of the command descriptor to issue the command. The CFIS also has to be written with the packet command. |
| 4–0 | TAG | The 5-bit TAG assigned by software for command tracking. It is the same as the value written to the command register host-to-device. |

15.3.7 Command Descriptor

As shown in [Figure 15-29](#), each entry in the command list points to a structure called the command descriptor.

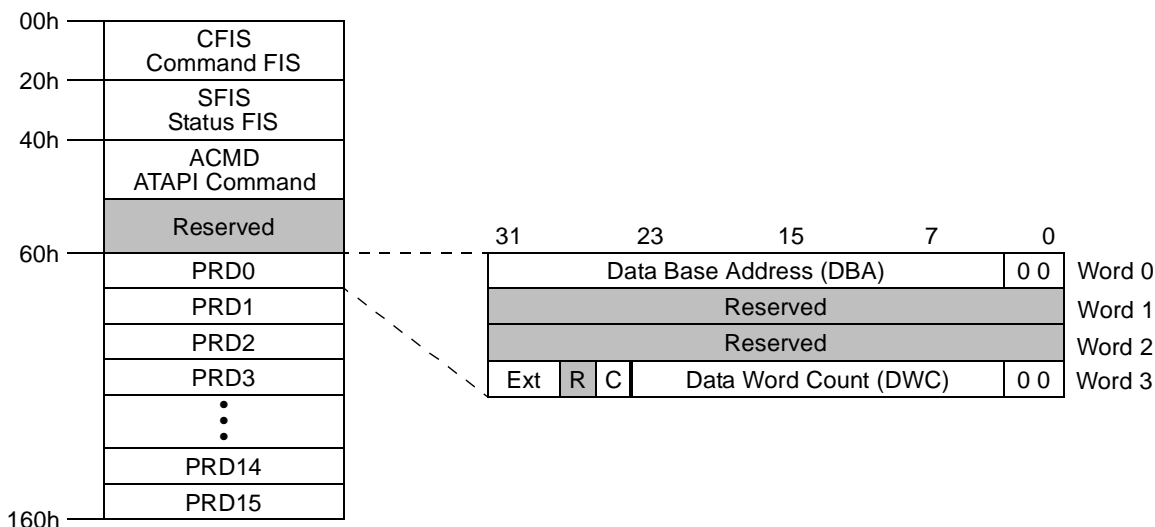


Figure 15-29. Command Descriptor

15.3.7.1 Command FIS Non-Queued Commands (CFIS)

Command FIS is a software constructed FIS. For data transfer operations, this is the H2D Register FIS format as specified in the Serial ATA 2.5 standard. The SATA controller fetches this from memory and sends the appropriate amount of data to the attached port. If a port multiplier is attached, this field must have the port multiplier port number in the FIS itself. CFIS lengths are two to eight words and must be in word granularity. A typical command FIS is shown in [Figure 15-30](#).




Figure 15-30. Register Host-to-Device

15.3.7.2 Command FIS First Party DMA Commands NCQ

Figure 15-31 shows register host-to-device first party DMA commands NCQ. The shaded components show where this FIS differs for the non-NCQ register host-to-device FIS.

Register FIS host to device—Read/write FPDMA queued

| | | | | | | | | | | | | |
|---|-------------------------------------|----|----------------|----|------------------------------------|----|----|---------------------|---------|----------------|---|---|
| | 31 | 24 | 23 | 16 | 15 | 12 | 11 | 8 | 7 | 3 | 2 | 0 |
| 0 | Features Sector Count 7:0 | | Command | | C | R | R | R | PM Port | FIS Type (27h) | | |
| 1 | Device | | LBA High | | LBA Mid | | | LBA Low | | | | |
| 2 | Features (exp) Sector Count 15:8 | | LBA High (exp) | | LBA Mid (exp) | | | LBA Low (exp) (0) | | | | |
| 3 | Control | | Reserved (0) | | Sector Count (exp) Reserved (0) | | | Sector Count TAG | | | | |
| 4 | Reserved (0) | | Reserved (0) | | Reserved (0) | | | Reserved (0) | | | | |

Figure 15-31. Register Host-to-Device First Party DMA Commands NCQ

15.3.7.3 Status FIS (SFIS)

This FIS is created in hardware. For normal operations, this is the D2H register FIS format as specified in the Serial ATA 2.5 standard. SFIS lengths are two to eight words and must be of word granularity. A typical status FIS is shown in Figure 15-32.

Register FIS device to host

| | | | | | | | | | | | | |
|---|--------------|----|----------------|----|--------------------|----|----|-------------------|---------|----------------|---|---|
| | 31 | 24 | 23 | 16 | 15 | 12 | 11 | 8 | 7 | 3 | 2 | 0 |
| 0 | Error | | Status | | R | I | R | R | PM Port | FIS Type (34h) | | |
| 1 | Device | | LBA High | | LBA Mid | | | LBA Low | | | | |
| 2 | Reserved (0) | | LBA High (exp) | | LBA Mid (exp) | | | LBA Low (exp) (0) | | | | |
| 3 | Reserved (0) | | Reserved (0) | | Sector Count (exp) | | | Sector Count | | | | |
| 4 | Reserved (0) | | Reserved (0) | | Reserved (0) | | | Reserved (0) | | | | |

Figure 15-32. Register Device-to-Host

15.3.7.4 ATAPI Command (ACMD)

This is a software constructed region of three or four words in length that contains the ATAPI command to transmit if the ‘A’ bit is set in the command header. The ATAPI command must be either 12 or 16 bytes in length. The length transmitted by the SATA IP is determined by the PIO setup FIS that is sent by the device requesting the ATAPI command.

15.3.7.5 Physical Region Descriptor Table (PRDT)

This is a software constructed table of addresses to use to complete the data transfer. Up to 16 structures can be supported in the current command descriptor. The format of the address entry is defined by the “Block vector structures for passing segmented data type of the IEEE Std. 1212.1-1993”. The total definable length supported in the 16 entries is 64 Mbytes.

Table 15-32 shows the PRDT word 0—data base address.

Table 15-32. PRDT Word 0—Data Base Address

| Bit | Name | Description |
|------|------|---|
| 31–2 | DBA | Data base address. Indicates the 32-bit physical address of the data block. The block must be word aligned, indicated by bits 1–0 being “reserved, must be 00.” |
| 1–0 | — | Reserved |

Table 15-33 shows the PRDT word 3—description information.

Table 15-33. PRDT Word 3—Description Information

| Bit | Name | Description |
|-------|------|---|
| 31 | EXT | If the extension flag is set to 1, then the DBA field contains the address of the extension segment, and the DWC field contains the size of this extension segment (this is called an “indirect descriptor”). |
| 30–23 | — | Reserved |
| 22 | C | Data snoop enable bit. When this bit is set, all data read/write operations associated with the PRD entry for this command will be snoopable. |
| 21–2 | DDC | Data word count. A 0-based value that indicates the length, in words, of the data block. A maximum length of 4 Mbytes may exist for any entry. Bits 1–0 of this field must always be 0 to indicate that size is in words. A value of 0x0_0000 indicates a full 4 Mbytes transfer. |
| 1–0 | — | Reserved |

15.3.8 Vendor-Specific BIST Operation

As part of the host self-diagnostic operation, a vendor-specific BIST mode is supported. This mode, in conjunction with a PHY that supports serial loopback, allows for the test of the SATA controller operation.

The mode exercises the following paths:

- DMA controller FIS transmission
- Command layer FIS transmission
- Transport layer Tx FIFO FIS transmission
- Link layer FIS transmission
- PHY modes
- Link layer FIS reception
- Transport layer Rx FIFO and FIS reception
- Command layer FIS reception
- Command layer FIS reception
- Host DMA controller FIS reception

To run this self-test on SATA1, lane A, the software performs the following operations:

1. Builds the vendor-specific header and descriptor as detailed in the vendor-specific BIST header and descriptor.

2. Modifies the SerDes control register 1 to place the PHY into analog loopback mode.
3. Issues this command to the SATA controller.
4. When the SATA controller indicates the command has completed, the software examines the contexts of the command descriptor's stats field. If the pre-programmed values are present, the test has passed.
5. The contents of the vendor-specific FIS to be created in memory are as follows:
 - Build command header in the memory. See [Table 15-34](#):

Table 15-34. Vendor BIST Test—Command Header

| Word Number | Hexadecimal Value |
|-------------|-------------------|
| Word 0 | CDA |
| Word 1 | 0x0000_000C |
| Word 2 | 0x0000_0000 |
| Word 3 | 0x0000_0400 |

- Build command descriptor in the memory. See [Table 15-35](#).

Table 15-35. Vendor BIST Test—Command Descriptor

| Word Number | Hexadecimal Value | Comments |
|-------------|-------------------|--|
| Word 0 | 0x0001_0058 | — |
| Word 1 | 0xAAAA_A034 | AAAA_A is the first test pattern (can be any value) |
| Word 2 | 0xB BBBB_B034 | BBBB_B is the second test pattern (can be any value) |
| Word 3 | Reserved | Reserved, must be all zeros |
| Word 4 | Reserved | Reserved, must be all zeros |
| Word 5 | Reserved | Reserved, must be all zeros |
| Word 6 | Reserved | Reserved, must be all zeros |
| Word 7 | Reserved | Reserved, must be all zeros |

The values of A and B can be filled in by the user with the pattern to test for. See [Figure 15-33](#).

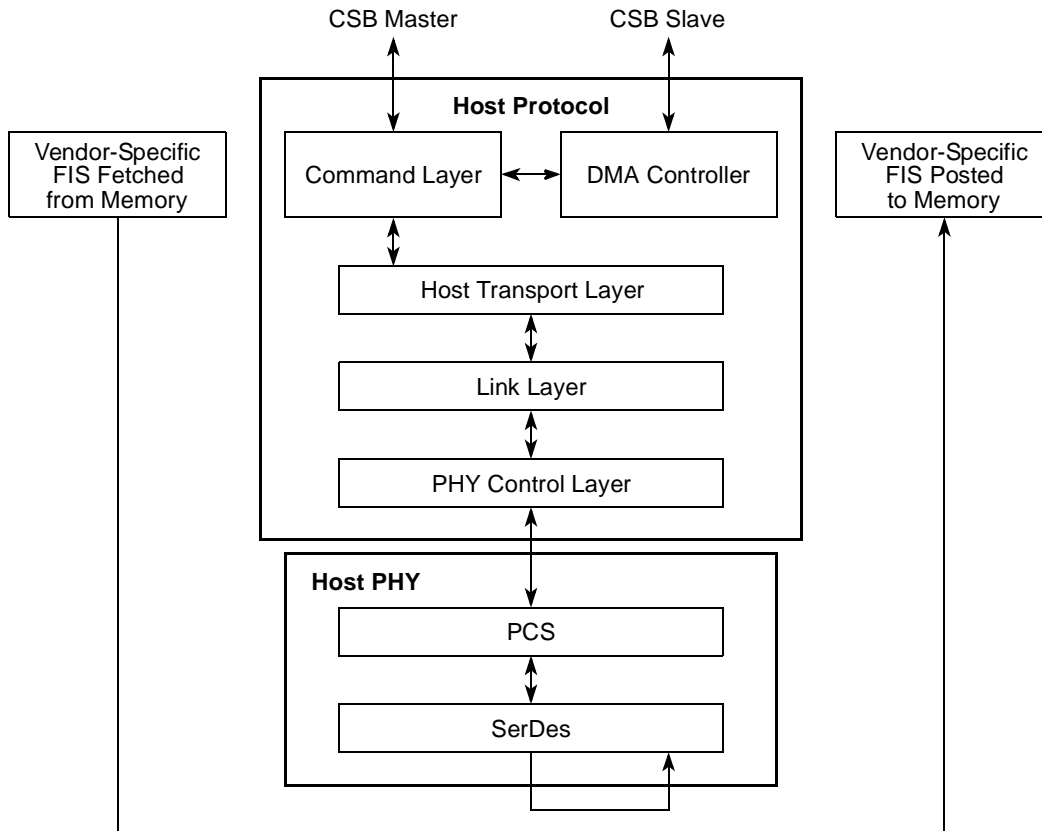


Figure 15-33. Vendor-Specific BIST Operation

15.4 Transport Layer Architectural Overview

The function of the SATA transport layer is to interface between the command and link layers in the transmission and reception of FIS.

On the transmit path, the transport layer frames the FIS's placed into the Tx FIFO. The FIS's are framed based on a programmed length for non-data FIS and are a configurable length for data FIS. When the transport layer is instructed to send a non-data FIS, it employs a retry policy until the far end signals acceptance of the transmitted FIS.

On the reception path, the transport layer deframes the FIS's and places them into the Rx FIFO. When an FIS is received, the transport layer informs the command layer. For a non-data FIS, the FIS is considered received when the end-of-frame (EOF) is signaled by the link layer and the FIS has been received with a good CRC. For a short vendor-specific FIS, the FIS is considered as a non-data FIS. For a longer vendor-specific FIS, the FIS reception is signaled when the RX FIFO reaches its water mark. For a data FIS, the FIS is considered received when the first word (header) is written into the FIFO.

The receive FIFO is written with data contained in the FIS sent by the link layer. When the data is stable at the output of the receive FIFO, the command layer can take the data. If the command layer is not ready to accept the data, the data builds up in the receive FIFO. When the receive FIFO exceeds its threshold, the transport layer stalls the link layer, which will in turn send HOLD primitives to the far end. This

threshold takes into consideration the latency involved in getting the far end to stop transmitting the data. This threshold is programmable to allow for the use of high-latency repeaters or retainers in between the host and device.

The transmit FIFO is written with data to be sent in the FIS transferred by the DMA controller. When the data is stable at the output of the transmit FIFO, the link layer can take the data. If the transmit FIFO cannot supply data to the link layer, the transport layer stalls the link layer, which will in turn send HOLD primitives to the far end.

15.5 Link Layer Overview

The function of the SATA link layer is to interface between the transport and physical layers in the transmission and reception of frames and primitives. The link layer utilizes the two unidirectional links provided by the SATA interface to maintain coordinated communication between the host and the device. Payload data can only be transmitted in one direction at a time. The link layer can work at either SATA first-generation (1.5 Gbps) or second-generation 2 (3 Gbps) speeds.

On transmit, the link layer first communicates with the peer far end link layer to determine if it is ready to receive. Assuming the far end link layer can receive data, the local link layer can then begin to take data in the form of words from its transport layer. It inserts start-of-frame (SOF) before the start of the data portion of a frame, calculates and inserts the CRC after the data portion of a frame, and inserts the EOF primitive at the end. The link layer scrambles the contents of the frame, including the calculated CRC, but excluding the SOF and EOF diameters and any other embedded primitives. The 8B/10B encoding of the data is done in the PHY layer. At the end of the transmission, the link layer reports transmission status to the transport layer.

On receive, the link layer first acknowledges its readiness to receive with its peer link layer. Then it awaits reception of the SOF primitive that marks the start of the received data. Following detection of the SOF primitive, the link layer proceeds to accept the incoming data. The 8B/10B decoding of the data is done in the PHY layer. Next, the link layer removes all primitives including the SOF and EOF diameters. It then descrambles the contents of the frame. The link layer also calculates the CRC on the incoming frame between the SOF and EOF delimiters, and compares this calculated value to the received value. Any mismatch is reported to the transport layer. During frame reception, disparity or code errors are reported to the command layer, and appropriate action is taken in the link layer. The descrambled and decoded receive data stream is passed to the transport layer as the frame is being received. Finally, at the end of the frame, the link layer reports reception status to the transport layer.

The link layer also partakes in flow control between the local and remote ends. The layer supports flow control actions based on the local FIFO status (located in the transport layer), or in response to receiving flow control messages from the remote end.

The transmit side of the link layer is also responsible for inserting a pair of ALIGN primitives every 254 words, or more frequently if programmed by the user.

15.5.1 Link Layer functionality

The link layer is composed of a number of functions:

- Link layer state machines
- Frame content scrambler and descrambler
- CRC generation and checking
- Bus interfaces to PHY and transport layer
- CONT primitive processing
- ALIGN insertion on transmit
- Debug functionality
- BIST support
- Link layer state machines

The four link layer state machines are described in the following sections.

15.5.1.1 Link Idle State Machine

The link idle state machine is responsible for detecting a transmit request from the transport layer or a frame reception request from the far end. The state machine arbitrates whether these two events coincide. The SATA specification defines that the host end always backs down in this case. Furthermore, this machine interprets power mode change requests from both the transport and PhyCtrl layers and initiates actions to enable the power mode change. Power mode change can only occur if the feature is enabled via the PhyCtrlCfg register LPB_EN bits. Finally, this state machine also detects the negation and assertion of PHY_READY from the PHY and notifies the transport layer of the change.

15.5.1.2 Transmit State Machine

This state machine is responsible for frame transmission to the PHY. The state machine places the SOF and EOF headers on each frame, calculates the CRC, and inserts it before the EOF delimiter. Between the SOF and CRC markers, the link layer accepts the current word from the transport layer and uses this as the next word of the frame. The link layer also inserts a pair of ALIGN primitives every 254 words of frame data. Finally, at the end of the frame transmission, the state machine waits for status from the far end link layer via received R_OK or R_ERR primitives. If the far end received the frame correctly, the local link layer signals TX_OK to the transport layer; otherwise, it signals TX_NOT_OK to the transport layer.

The transmit state machine also partakes in flow control actions, if necessary, during packet transmission. If the transport layer cannot supply a new word and the frame is not finished, the transmit state machine responds by sending HOLD primitives until the transport layer is ready with valid frame data. Also, during frame transmission, if the state machine detects a received HOLD primitive from the PHY layer, it interrupts the current frame transmission and sends HOLDA primitives to the PHY to be transmitted to the far end.

The current frame transmission can only be aborted by two events. The first is on reception of a DMAT primitive from the far end. In this case, the link layer state machine stops the current transfer and calculates

and inserts the current CRC. This is a controlled termination. The second is when the transport layer wishes to send a control register frame signaled via TRANSMIT_CRF.

If at any point in the frame transmission process, the link layer detects error conditions, it signals these to the command layer. The errors can occur if the link layer detects the following conditions:

- PHY_READY negates
- SYNC primitive is received during frame transmission

15.5.1.3 Receive State Machine

This state machine is responsible for frame reception from the PHY layer. The state machine removes the SOF and EOF headers and other primitives from each frame, calculates the CRC, and compares it to the received CRC. Between the SOF and CRC markers, the link layer accepts the current word from the Phy layer and uses this as the next word of the frame, transferring it to the transport layer. At the end of the frame reception, if the calculated CRC is not the same as the received CRC, the link layer signals an error to the transport layer. This is done via RX_CRC_OK and RX_CRC_NOT_OK. During frame reception, if no errors are detected, the link layer transmits R_IP primitives to the far end peer link layer. Finally, at the end of the frame reception, the link layer sends the R_OK primitive if no error was detected during reception. If an error was detected, it sends a R_ERR primitive instead.

The receive state machine also partakes in flow control actions if necessary, during FIS reception. If the transport layer cannot accept a new word, (because its receive FIFO has reached its watermark level), and the FIS is not finished, the receive state machine responds by sending HOLD primitives on the back channel until such time as the transport layer is ready to accept FIS data again. Also, during FIS reception, if the state machine detects a received HOLD primitive from the far end, it responds by sending HOLDA primitives to the far end.

The current frame reception can be interrupted if the transport layer wishes to send a control register frame, signaled via TRANSMIT_CRF.

If at any point in the frame reception process, the link layer detects error conditions, it signals these to the command layer. The errors can occur if the link layer detects the following conditions:

- PHY_READY negates
- SYNC primitives is received during frame transmission
- WTRM primitive is received before EOF

15.5.1.4 Power Mode Change State Machine

This state machine is responsible for handling change of power mode requests. These requests can come from the command layer superset registers or the far end. This state machine responds by transmitting PMREQ_P/PMREQ_S primitives to the far end and waiting for PMACK primitives from it in response. Once PMACK is received, the state machine instructs the PHY layer to enter either a partial or slumber state.

A write to the SControl register SPM field or reception of a COMWAKE from the far end will initiate a resume to active power mode.

If the link layer receives a PMREQ_P/PMREQ_S primitive from a peer link layer and is enabled to perform power management modes (SControl IPM bits are cleared), it responds by sending at least four PMACK primitives. A write to the SControl register SPM field or reception of a COMWAKE from the far end will initiate a resume to active power mode.

If the link layer receives an XRDY primitive from the far end while it is in the partial or slumber state, it returns to idle and signals a link sequence error to the command layer, that is, SError[S] = 1.

15.5.1.5 Frame Content Scrambler and Descrambler

There are two separate scramblers used in the SATA controller, one for the data payload and the other for repeated primitive suppression. The contents of each word of data (excluding all primitives) between SOF and EOF must be scrambled before 8B/10B encoding. Scrambling is performed on word quantities according to the following polynomial:

$$G(X) = X^{16} + X^{15} + X^{13} + X^4 + 1$$

The scrambler is initialized with a seed value of 0xFFFF at each SOF transmission and rolls over every 2048 words. Payload data is scrambled prior to transmission, by XORing the data to be transmitted with the output of this scrambler.

If a CONT primitive is transmitted, then the intervening data between the last CONT primitive and a subsequent primitive must be scrambled also. This scrambler uses the same polynomial as defined above for data payload scrambling and is reset to the initial value upon detection of a COMINIT or COMRESET event. If a CONT primitive is transmitted or received during a frame transfer, then the current data payload scrambler value at the last word is held.

When payload data is received by the link layer it is descrambled by XORing it with the output of its descrambler. The descrambler is re-seeded at the beginning of the received data payload, that is, at each SOF reception. The descrambler uses the same polynomial as the scrambler.

15.5.1.6 CRC Generator and Checker

A 32-bit CRC is calculated on the data contents of each frame and is inserted in the word before the EOF. The CRC covers all data bytes in the frame excluding any primitives such as SOF, EOF, HOLD, HOLDA, DMAT, SYNC, X_RDY, R_RDY, or ALIGNs.

The CRC generator works on word quantities. Any padding to the boundary is done in the transport layer. The polynomial used for the CRC is as follows:

$$G(X) = X^{32} + X^{26} + X^{23} + X^{22} + X^{16} + X^{12} + X^{11} + X^{10} + X^8 + X^7 + X^5 + X^4 + X^2 + X + 1$$

The CRC is initialized with a seed value of 0x52325032 at each SOF.

The CRC generation or checking does not apply to primitives (as stated above) or to CONT'ed primitives. If a CONT primitive is transmitted or received, then the intervening data between the last CONT primitive and a subsequent primitive is not included in the CRC calculation for a frame. If this happens during a frame transfer, then the current CRC scrambler value at the last word is held.

15.5.1.7 8B/10B Encode and Decode

All data and primitives must be encoded prior to transmission on the line. The 8B/10B encode/decode occur in the PHY layer.

15.5.1.8 CONT Primitive Processing

Using the CONT primitive, the link layer is capable of replacing repetitive primitive streams with scrambled data. This reduces EMI emissions because primitives are not scrambled.

The link layer can transmit a CONT primitive at a point where it knows it must transmit a number of repeating primitives. After a CONT primitive has been transmitted, the link layer then transmits scrambled junk data to the PHY layer. The content of this junk data is disregarded. At the far end link layer, the reception of a CONT primitive will cause the last received valid primitive to be implied to be repeated until it receives the next valid non-ALIGN primitive. Transmission of a new valid primitive halts the current CONT processing; reception of a new valid non-ALIGN primitive halts the current CONT processing.

This action can occur on transmit and receive. The link layer supports both the transmission and reception of CONT primitives.

15.5.1.9 ALIGN Insertion

The link layer is responsible for ALIGN insertion and removal at a fixed frequency. A pair of ALIGN primitives are inserted into the transmit data stream every 254 words. At the receive end, the ALIGN primitives are stripped from the incoming data stream in the link layer.

For diagnostic purposes, the rate of ALIGNs can be increased as much as two ALIGNs per one word; for example, ALIGN, ALIGN, data, ALIGN, ALIGN. In addition, the SEND_4_ALIGNs bit can be set to instruct the link layer to send four ALIGNs at a time instead of two.

15.5.1.10 Debug Functionality

There are a number of useful features designed into the link layer to aid debug, as follows:

- The align insertion rate can be increased using the ALIGN_RATE register field in the command layer.
- Four error counters can be monitored by issuing register reads to the command layer: the disparity error counter, the code error counter, the PHY internal error counter, and the control character error counter.
- A number of configuration bits in the command layer can be used to override normal primitive insertion. For example,
 - Set PRIM_OVR_STATE = (L_SendHold state (16))
 - Set PRIM = 0xb5b5957c, that is, a SYNC primitive
 - During the transfer, set PRIM_OVRD_EN = 1
- When the link layer detects a rising edge on PRIM_OVRD_EN, it will insert one SYNC primitive into the datastream in place of the HOLD, when the LINK_STATE reaches the L_SendHold state.

Only one HOLD primitive will be overridden; the PRIM_OVRD_EN must be cleared and written to again to force another override to occur.

15.5.1.11 BIST Support

The transmit and receive subblocks of the link layer contain logic to support BIST activate FIS functionality.

When a BIST activate FIS is either received or transmitted successfully by the transport layer, it issues a request to the link layer to enter BIST mode. This forces the link layer to enter a BIST state in its state machine as soon as it receives a SYNC primitive from the far end. In the BIST state, the link layer transmits a data sequence as specified by the two BIST data patterns in the BIST activate FIS. The link layer also monitors the incoming data from the PHY to detect the BIST data pattern is as specified in the BIST activate FIS. When it detects the correct data sequence, the HStatus[BIST_Err] is deasserted. The BIST_Err bit will stay deasserted unless an error occurs in the datastream from the far end.

15.6 PHY Control Layer Overview

The PHY control layer operates between the PHY and link layers. On receive, the PHY control layer converts the 16-bit parallel data from the PHY to a 32-bit word, which it presents to the link layer. The PHY control layer aligns the control word of the SATA primitive to the lowest word position of the word. The PHY control layer takes in the per-byte error signals and the per-byte control/data bits output by the PHY and converts them into 4-bit buses, with each bit of the bus corresponding to a byte in the word.

On transmit, the PHY control layer takes in the 32-bit transmit data from the link layer and converts it to 16 bits of data which it presents to the PHY. The control/data bit from the link layer (which is always assumed to be associated with the lowest byte position of the transmit word) is also passed onto the PHY with the appropriate word.

15.7 Initialization/Application Information

The typical initialization sequence is described in the following sections.

15.7.1 SATA Controller Initialization Steps

These steps bring the SATA controller online, synchronize the SATA controller with the attached device, and issue typical command for execution.

1. Write HControl[HC_ON] = 1 to bring the SATA controller online.
2. Poll the HStatus[HS_ON] till HS_ON = 1, indicating that the controller is online.
3. Poll the SStatus[DET] till DET = 4'b0011 meaning that the device presence is detected and PHY communication is established. In this state, SStatus[SPD] indicates the negotiated communication speed.
4. To read the device's signature, poll HStatus[SIG_UPD] till it goes up. Read the signature from the SIG register.
5. Initialize the CHBA register to point to the command header block.

6. Build a command header block in memory. Refer to [Section 15.3.6, “Command Header.”](#)
7. Build a command descriptor block in memory. Refer to [Section 15.3.7, “Command Descriptor.”](#)
8. Build a number of PRD tables in memory as defined by the PRD_NUM field in the command header. Refer to [Section 15.3.7.5, “Physical Region Descriptor Table \(PRDT\).”](#)
9. Initialize the CQPMP register with the device’s PM number. If the port multiplier is not used, clear this field.
10. Poll the CQR[CQn] to determine which command can be issued.
11. After CQn is determined, write 1 to CQR[CQn] to issue this command and start execution.
12. Poll the CCR[CCn] till CCn goes up, indicating that the command is completed.

The following example presents the structure of descriptors to issue the ReadDMA command:

- Build the command header in memory. See [Table 15-36](#), below.

Table 15-36. Read DMA Command—Command Header

| Word Number | Hexadecimal Value | Description |
|-------------|-------------------|--|
| Word 0 | CDA | Pointer to memory where command descriptor begins |
| Word 1 | 0x0002_0014 | Two PRD tables contain the data, FIS length = 20 bytes |
| Word 2 | 0x0000_0200 | Length of data associated with this command = 0x200 |
| Word 3 | 0x0000_0000 | Tag = 0 |

- Build command descriptor in memory. See [Table 15-37](#), below.

Table 15-37. Read DMA Command—Command Descriptor

| Word Number | Hexadecimal Value | Description |
|-------------|-------------------|--------------------|
| Word 0 | 0x00C8_8027 | Command = Read DMA |
| Word 1 | 0x0000_0010 | LBA = 24’h10 |
| Word 2 | 0x0000_0000 | LBA(exp) = 24’h0 |
| Word 3 | 0x0000_0001 | Sector Count = 1 |
| Word 4 | 0x0000_0000 | Reserved |

- Build two PRD entries in memory. See [Table 15-38](#).

Table 15-38. Read DMA Command—PRD Entries

| Word Number | Hexadecimal Value | Description |
|-------------|-------------------|------------------------------------|
| Word 0 | PRD1 | Address of first portion of data. |
| Word 1 | 0x0000_0000 | Reserved |
| Word 2 | 0x0000_0000 | Reserved |
| Word 3 | 0x0000_0100 | PRD1 contains 0x100 bytes of data |
| Word 4 | PRD2 | Address of second portion of data. |
| Word 5 | 0x0000_0000 | Reserved |

Table 15-38. Read DMA Command—PRD Entries (continued)

| Word Number | Hexadecimal Value | Description |
|-------------|-------------------|-----------------------------------|
| Word 6 | 0x0000_0000 | Reserved |
| Word 7 | 0x0000_0100 | PRD2 contains 0x100 bytes of data |

- Fill the PRD1 and PRD2 with user-defined data; see [Figure 15-34](#), below.




Figure 15-34. PRD Entry

Chapter 16

SerDes PHY

NOTE

On the MPC8315E there are two SerDes lanes, A and B. This chapter describes a SerDes with lanes A and E. Whenever lane E is mentioned in this chapter, the MPC8315E lane B is intended.

16.1 Introduction

The SerDes PHY block includes the following components:

- SerDes PHY
- Protocol multiplexer and converter per protocol
- Control registers and control logic
- Power-down/reset state machine for cold (power-on) or warm (software-initiated) reset of the SerDes, PCVTR, and controllers
- Interface with the clock controls

16.1.1 Overview

Figure 16-1 is a block diagram showing the functional blocks inside the SerDes PHY block and its connections to other modules in the processor.




Figure 16-1. SerDes PHY Block Diagram

16.1.2 Features

The SerDes PHY block has the following features:

- Support for either two x1 PCI Express or two x1 SGMII
 - Gen1i, Gen1m, Gen2i, and Gen2m electrical specifications are supported in SATA mode, compliant to *Serial ATA 2.5 Specification*
- Link-layer interfaces to IP controller
- Memory-mapped registers with 256-byte address region
- SerDes power-down/reset state machine for cold (power-on) or warm (software-initiated) reset of SerDes, PHY, and controllers

16.1.3 Modes of Operation

The SerDes PHY block supports the following modes of operation:

- SerDes
 - Two lanes running x1 SGMII at 1.25 Gbps, or
 - Two lanes running x1 PCI Express at 2.5 Gbps




Figure 16-2. Modes of Operation

16.1.4 Clock

The SerDes control has one clock which is running at platform speed. This is an internally generated clock based off of the system clock.

16.2 External Signals

Table 16-1 describes the external signals to the SerDes PHY block.

Table 16-1. SerDes External Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|----------------------------------|-----|--|
| SD_IMP_CAL_RX | I | Receiver impedance calibration control signal. This pin requires an external resistor to ground to set the differential input impedance of the receivers. |
| | | State Meaning Assertion/Negation—The SerDes acts as an integrated active impedance calibration circuit to ensure the best possible impedance control of the receiver's link termination resistors. The calibration circuit uses an externally established impedance against which internal impedance is calibrated. The user connects a 200 Ω , 1% tolerance resistor between the sd_imp_cal_rx input and ground. If the user wishes to set the impedance control to its nominal value, the sd_imp_cal_rx input maybe tied directly to the xc0revdd supply. |
| SD_IMP_CAL_TX | I | Transmitter impedance calibration control signal. This pin requires an external resistor to ground to set the differential output impedance of the transmitters. |
| | | State Meaning Asserted/Negated—The SerDes acts as an integrated active impedance calibration circuit to ensure the best possible impedance control of the transmitter's output impedance resistors. The calibration circuit uses and externally established impedance against which internal impedance is calibrated. The user connects a 100 Ω , 1% tolerance resistor between the sd_imp_cal_tx input and ground. If the user wishes to set the impedance control to its nominal value, the sd_imp_cal_tx input maybe tied directly to the xpadvdd supply. |
| SD_REF_CLK | I | SerDes PLL reference clock, along with $\overline{\text{SD_REF_CLK}}$, is used by the SerDes PLL to generate all of the necessary clocks for the SerDes. |
| | | Timing Assertion/Negation—Choices of input reference clock values are limited by specification, output bit rate and PLL functionality. |
| $\overline{\text{SD_REF_CLK}}$ | I | SerDes PLL reference clock complement, along with SD_REF_CLK, is used by the SerDes PLL to generate all of the necessary clocks for the SerDes. |
| | | Timing Assertion/Negation—Choices of input reference clock values are limited by specification, output bit rate and PLL functionality. |
| SD_RXA | I | Serial receiver data, lane A, positive. |
| | | Timing Assertion/Negation—Serial differential receiver input which can be configured to meet either the PCI Express, or SGMII specifications. |
| $\overline{\text{SD_RXA}}$ | I | Serial receiver data, lane A, complement. |
| | | Timing Assertion/Negation—Serial differential receiver input which can be configured to meet either the PCI Express, or SGMII specifications. |
| SD_RXE | I | Serial receiver data, lane E, positive. |
| | | Timing Assertion/Negation—Serial differential receiver input which can be configured to meet either the PCI Express, or SGMII specifications. |
| $\overline{\text{SD_RXE}}$ | I | Serial receiver data, lane E, complement. |
| | | Timing Assertion/Negation—Serial differential receiver input which can be configured to meet either the PCI Express, or SGMII specifications. |

Table 16-1. SerDes External Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|--------|-----|--|
| SD_TXA | O | Serial transmitter data, lane A, positive. |
| | | Timing Assertion/Negation—Serial differential transmitter output which can be configured to meet either the PCI Express, or SGMII specifications. |
| SD_TXA | O | Serial transmitter data, lane A, complement. |
| | | Timing Assertion/Negation—Serial differential transmitter output which can be configured to meet either the PCI Express, or SGMII specifications. |
| SD_TXE | O | Serial transmitter data, lane E, positive. |
| | | Timing Assertion/Negation—Serial differential transmitter output which can be configured to meet either the PCI Express, or SGMII specifications. |
| SD_TXE | O | Serial transmitter data, lane E, complement. |
| | | Timing Assertion/Negation—Serial differential transmitter output which can be configured to meet either the PCI Express, or SGMII specifications. |

16.3 Memory Map/Registers

Table 16-2 lists the SerDes PHY block registers and their addresses. Reading undefined portions of the memory map returns all zeros; writing has no effect. All the registers are 32 bits wide.

Table 16-2. SerDes PHY Block Memory Map

| Offset | Register | Access | Reset | Section/Page |
|---|--|--------|-------------|------------------------------|
| SerDes PHY—Block Base Address 0xE_3000 | | | | |
| 0xE_3000 | SRDSCR0—SerDes Control Register 0 | R/W | 0x1100_CC30 | 16.3.1/16-5 |
| 0xE_3004 | SRDSCR1—SerDes Control Register 1 | R/W | 0x0000_0040 | 16.3.2/16-8 |
| 0xE_3008 | SRDSCR2—SerDes Control Register 2 | R/W | 0x0080_0000 | 16.3.3/16-9 |
| 0xE_300C | SRDSCR3—SerDes Control Register 3 | R/W | 0x0101_0000 | 16.3.4/16-10 |
| 0xE_3010 | SRDSCR4—SerDes Control Register 4 | R/W | 0xnn00_0n0n | 16.3.5/16-11 |
| 0xE_3014– 0xE_301C | Reserved | — | — | — |
| 0xE_3020 | SRDSRSTCTL—SerDes Reset Control Register | R/W | 0x0044_4500 | 16.3.6/16-13 |
| 0xE_3024– 0xE_31FC | Reserved | — | — | — |

NOTE

Reserved bits should always be written with the value they return when read. That is, the register should be programmed by reading the value, modifying appropriate fields, and writing back the value.

Table 16-3. SRDSCR0 Field Descriptions (continued)

| Bits | Name | Description |
|-------|-------|---|
| 6–7 | RXEQE | Receive equalization selection bus for lane E—when asserted in PCI Express mode: 00 No equalization 01 2 dB of equalization 10 4 dB of equalization 11 Reserved Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 01 • SGMII: 01 |
| 8–15 | — | Reserved |
| 16 | DPPA | Diff pk-pk swing for lane A. Sets the peak value for output swing of transmitters and the amount of transmit equalization for lane A. 0 $V_{DD-diff-pk-pk}$ 1 Reserved Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 0 • SGMII: 0 |
| 17–19 | TXEQA | Sets the peak value for output swing of transmitters and the amount of transmit equalization for lane A. Transmit equalization selection bus for lane A. 000 No equalization 001 1.09x relative amplitude 010 1.2x relative amplitude 011 1.33x relative amplitude 100 1.5x relative amplitude 101 1.71x relative amplitude 110 2.0x relative amplitude 111 Reserved Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 100 • SGMII: 100 |
| 20 | DPPE | Diff pk-pk swing for lane E. Sets the peak value for output swing of transmitters and the amount of transmit equalization for lane E. 0 $V_{DD-diff-pk-pk}$ 1 Reserved Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 0 • SGMII: 0 |

Table 16-3. SRDSCR0 Field Descriptions (continued)

| Bits | Name | Description |
|-------|-------|--|
| 21–23 | TXEQE | Sets the peak value for output swing of transmitters and the amount of transmit equalization for lane E. Transmit equalization selection bus for lane E. 000 No equalization 001 1.09x relative amplitude 010 1.2x relative amplitude 011 1.33x relative amplitude 100 1.5x relative amplitude 101 1.71x relative amplitude 110 2.0x relative amplitude 111 Reserved Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 100 • SGMII: 100 |
| 24 | SDPD | SerDes power down. This power down signal shuts down the PLL, all of the receiver amplifiers, all of the samplers and places the transmitters in 3-state. 0 Application mode 1 Block power down |
| 25 | — | Reserved |
| 26 | IACCA | Used to set on-chip AC coupling in the receiver in lane A. 0 Disable on-chip AC coupling 1 Enable on-chip AC coupling Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 1 • SGMII: 1 |
| 27 | IACCE | Used to set on-chip AC coupling in the receiver in lane E. 0 Disable on-chip AC coupling 1 Enable on-chip AC coupling Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 1 • SGMII: 1 |
| 28–29 | — | Reserved |
| 30 | RXEIA | When asserted, places lane A into receiver electrical idle state. 0 Lane A is not 'forced' into receive electrical idle state 1 Place lane A into electrical idle state Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 0 • SGMII: 0 |
| 31 | RXEIE | When asserted, places lane E into receiver electrical idle state. 0 Lane E is not 'forced' into receive electrical idle state 1 Place lane E into electrical idle state Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 0 • SGMII: 0 |

16.3.2 SerDes Control Register 1 (SRDSCR1)

SRDSCR1, shown in Figure 16-4, contains the functional control bits for the SerDes logic.

Individual lanes can be powered down using SRDSCR1[0] and SRDSCR1[4]. The entire SerDes must be reset to activate a lane from power-down. Refer to Section 16.3.6, “SerDesn Reset Control Register (SRDSRSTCTL).”

Figure 16-4. SerDes Control Register 1 (SRDSCR1)




Table 16-4 describes the SRDSCR1.

Table 16-4. SRDSCR1 Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0 | PDA | Lane A power down. 0 Normal 1 Power down lane A |
| 1–3 | — | Reserved |
| 4 | PDE | Lane E power down. 0 Normal 1 Power down lane E |
| 5–9 | — | Reserved |
| 10 | X3SA | Lane A transmitter three-state. 0 Normal 1 The transmitter output is disabled and place in a three-state condition |
| 11–13 | — | Reserved |
| 14 | X3SE | Lane E transmitter three-state. 0 Normal 1 The transmitter output is disabled and place in a three-state condition |
| 15–20 | — | Reserved |
| 21–24 | LBSEL | Select type loop-back. 0000 Application mode 0001 Digital loopback both lanes 0010 Analog loopback both lanes All other modes reserved for test |

Table 16-5 describes the SRDSCR2.

Table 16-5. SRDSCR2 Field Descriptions

| Bits | Name | Description |
|-------|----------|---|
| 0–13 | — | Reserved |
| 14 | TIMPCALS | Transmitter impedance calibration stop command. Allows user to stop calibration of transmitter impedances. 0 Run transmit impedance calibration 1 Stop transmit impedance calibration Recommended setting per protocol: • PCI Express: 0 • SGMII: 0 |
| 15 | RIMPCALS | Receiver impedance calibration stop command. Allows user to stop calibration of receiver impedances. 0 Run receive impedance calibration 1 Stop receive impedance calibration Recommended setting per protocol: • PCI Express: 0 • SGMII: 0 |
| 16–21 | — | Reserved |
| 22–23 | PEICA | PCI-EXP Receiver electrical idle detection control for lanes A–D 00 Exit from idle ~85 UI and unexpected idle detect ~1 μs (application mode) 01 Exit from idle ~85 UI and unexpected idle detect ~10 μs 10 Exit from idle ~45 UI and unexpected idle detect ~1 μs 11 Bypass |
| 24–29 | — | Reserved |
| 30–31 | PEICE | PCI Express receiver electrical idle detection control for lanes E–H. 00 Exit from idle ~85 UI and unexpected idle detect ~1 μs (application mode) 01 Exit from idle ~85 UI and unexpected idle detect ~10 μs 10 Exit from idle ~45 UI and unexpected idle detect ~1 μs 11 Bypass |

16.3.4 SerDes Control Register 3 (SRDSCR3)

SRDSCR3, shown in Figure 16-6, contains the functional control bits for the SerDes logic.




Figure 16-6. SerDes Control Register 3 (SRDSCR3)

Table 16-6 describes the SRDSCR3.

Table 16-6. SRDSCR3 Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0–5 | — | Reserved |
| 6–7 | SDFMA | Sets the bandwidth of the digital filter to optimize for given frequency offset specification for lane A. 00 200 ppm (SGMII) 01 600 ppm (PCI Express or SATA) 10 Reserved 11 Reserved Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 01 • SGMII: 00 |
| 8–13 | — | Reserved |
| 14–15 | SDFME | Sets the bandwidth of the digital filter to optimize for given frequency offset specification for lane E. 00 200 ppm (SGMII) 01 600 ppm (PCI Express or SATA) 10 Reserved 11 Reserved Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 01 • SGMII: 00 |
| 16–31 | — | Reserved[|

16.3.5 SerDes Control Register 4 (SRDSCR4)

SRDSCR4, shown in [Figure 16-7](#), contains the functional control bits for the SerDes logic.

NOTE

The protocol select configuration should be identical for both lane A and lane E. The user can only select identical protocol on both lanes. To power down a lane, use SRDSCR1[0] or SRDSCR1[4] to power down individual lanes (refer to [Section 16.3.2, “SerDes Control Register 1 \(SRDSCR1\)”](#)).

Valid combinations for protocol select:

- PCI Express mode (PROTA/PROTE = 001) is only a x1 lane. The reference clock can be either 100 or 125 MHz.
- SGMII mode (PROTA/PROTE = 101) is only a x1 lane. The reference clock can be either 100 or 125 MHz.

Valid combinations for protocol select:

- PCI Express mode (AD_PROTO_SEL/EH_PROTO_SEL[2:0] = 001) is only a x1 lane. The reference clock can be either 100 or 125 MHz.
- SGMII mode (AD_PROTO_SEL/EH_PROTO_SEL[2:0] = 101) is only a x1 lane. The reference clock can be either 100 or 125 MHz.

NOTE

The entire SerDes need to be reset in order to activate a lane from power-down. The SRDSCR4 register is initialized base on RCWH[TSEC1M] and RCWH[TSEC2M]. If SGMII mode is selected, the values will be according to SGMII protocol; otherwise, values will be initialized to PCI Express protocol. The SRDSCR3 register is not initialized based on RCWH[TSEC1M] and RCWH[TSEC2M]; it needs to be done explicitly based on the usage scenario. (Refer to [Section 4.3.2.2.5, “eTSEC1 Mode,”](#) and [Section 4.3.2.2.6, “eTSEC2 Mode.”](#))

16.5 Power Management: Power Down

The SerDes is capable of several different power management states depending on the settings of the protocol selection and power down signals.

By setting the register field SRDSCR0[24] powers down the entire SerDes and is comparable to the PCI Express L2 low power link state. The steps for powering down the SerDes are as follows:

1. Apply power to all XCOREVDD, XPADVDD, SDAVDD supplies.
2. Be sure all XCOREVSS, XPADVSS, and SDAVSS supplies are grounded.
3. Assert the SRDSCR0[24] control from the chip logic to whichever SerDes block is not in use. This safely parks all the analog circuitry and stops all SerDes-generated clocks.
4. Tie all unused SD_RX n and $\overline{\text{SD_RX}}_n$ serial differential inputs to XCOREVSS.
5. Float all unused SD_TX n and $\overline{\text{SD_TX}}_n$ serial differential outputs.
6. If a SerDes block is not being used and has its own receiver and transmitter calibration external resistors, then these can be tied to XCOREVDD for the receiver (SD_IMP_CAL_RX) and XPADVDD for the transmitter (SD_IMP_CAL_TX).
7. Tie SD_REF_CLK and $\overline{\text{SD_REF_CLK}}$ both to XCOREVSS.

NOTE

If the entire SerDes is powered down, or even if parts of the SerDes is powered down, all power pads in the SerDes must be powered.

Powering down the SerDes includes the following steps:

1. Disable the PLL and place its output clocks into a known, static state.
2. Power down the receiver termination and amplifier cells. In PCI Express mode, there is still a differential termination, but its value is no longer calibrated. Also, there is no longer a termination to ground. In SGMII mode, there is still a termination to ground and its value is calibrated.

3. Power down the transmitter and receiver impedance control amplifiers so that the termination impedances are uncalibrated.
4. Power down the transmitter cell. The $V_{TX-DIFFp} < 20$ mV. The DC common mode voltage is not held.



Chapter 17

Universal Serial Bus Interface

This chapter describes the universal serial bus (USB) interface of the device. The USB interface implements many industry standards. However, it is beyond the scope of this document to document the intricacies of these standards. Instead, it is left to the reader to refer to the governing specifications.

The following documents are available from the USB Implementers Forum web page at <http://www.usb.org/developers/docs/>.

- *Universal Serial Bus Revision 2.0 Specification*
- *On-The-Go Supplement to the USB 2.0 Specification, Revision 1.0a*

The following documents are available from the Intel USB Specifications web page at <http://www.intel.com/technology/usb/spec.htm>.

- *Enhanced Host Controller Interface (EHCI) Specification for Universal Serial Bus, Revision 1.0*
- *USB 2.0 Transceiver Macrocell Interface (UTMI) Specification, Version 1.05*

The following documents are available from the ULPI web page at <http://www.ulpi.org>.

- *UTMI + Specification, Revision 1.0*
- *UTMI Low Pin-Count Interface (ULPI) Specification, Revision 1.0*

17.1 Introduction

The device implements a dual-role (DR) USB module. This module may be connected to an external port. Collectively the module and external port are called the USB interface. The USB interface is shown in [Figure 17-1](#).




Figure 17-1. USB Interface Block Diagram

17.1.1 Overview

The USB DR module is a USB 2.0-compliant serial interface engine for implementing a USB interface. The registers and data structures for the module are based on the *Enhanced Host Controller Interface Specification for Universal Serial Bus* (EHCI) from Intel Corporation. The DR module can act as a device or host controller. Interfaces to negotiate the host or device role on the bus in compliance with the On-The-Go (OTG) supplement to the USB specification are also provided.

The DR module supports the required signaling for USB transceiver macrocell interface (UTMI) and UTMI low pin count interface (ULPI) transceivers (PHYs). The PHY interfacing to the UTMI is an internal PHY, and the PHY interfacing to the ULPI is an external PHY.

The module contains a chaining DMA (direct memory access) engine that reduces the interrupt load on the application processor and reduces the total system bus bandwidth that must be dedicated to servicing the USB interface requirements.

17.1.2 Features

The USB DR module includes the following features:

- Complies with USB specification rev 2.0
- Supports operation as a standalone USB host controller
 - Supports enhanced host controller interface (EHCI)
- Supports high-speed (480 Mbps), full-speed (12 Mbps), and low-speed (1.5 Mbps) operation. Low speed is only supported in host mode.

- Supports external PHY with ULPI (UTMI + low-pin interface) OR
- On-chip USB-2.0 full-speed/high-speed PHY with UTMI
- Supports operation as a standalone USB device
 - Supports one upstream facing port
 - Supports three programmable, bidirectional USB endpoints
- Host and device support
- OTG (on-the-go) support, which includes both device and host capability, with external PHY (ULPI)

17.1.3 Modes of Operation

The USB DR module has three basic operating modes: host, device, and OTG. The module can be configured to use one of two different PHY interfaces: ULPI or UTMI. However, due to pin limitations of the integrated UTMI PHY, OTG is not supported through the UTMI interface.

NOTE

Only high-speed and full-speed operations are supported in device mode.

17.2 External Signals

This section contains detailed descriptions of all the USB dual-role controller signals.

Table 17-1. USB External Signals

| Signal | I/O | Description |
|--------------------|-----|---------------------------------------|
| USBDR_D0_ENABLEN | I/O | ULPI—Use as USBDR_D0 USB_DP |
| USBDR_D1_SER_TXD | I/O | ULPI—Use as USBDR_D1 USB_DM |
| USBDR_D2_VMO_SE0 | I | ULPI—Use as USBDR_D2 USB_VBUS |
| USBDR_D3_SPEED | I/O | ULPI—Use as USBDR_D3 |
| USBDR_D4_DP | O | ULPI—Use as USBDR_D4 USB_TPA |
| USBDR_D5_DM | I | ULPI—Use as USBDR_D5 USB_RBIAS |
| USBDR_D6_SER_RCV | I | ULPI—Use as USBDR_D6 USB_PLL_PWR3 |
| USBDR_D7_DRVVBUS | I | ULPI—Use as USBDR_D7 USB_PLL_GND0 |
| USBDR_SESS_VLD_NXT | I | ULPI—Use as USBDR_NXT USB_PLL_GND1 |

Table 17-1. USB External Signals (continued)

| Signal | I/O | Description |
|--------------------|-----|---|
| USBDR_DIR_DPPULLUP | I | ULPI—Use as USBDR_DIR USB_PLL_PWR1 |
| USBDR_STP_SUSPEND | I | ULPI—Use as USBDR_STP USB_VSSA_BIAS |
| USBDR_PWRFAULT | I | ULPI—Use as USBDR_PWRFAULT USB_VDDA_BIAS |
| USBDR_PCTL0 | I | ULPI—Use as USBDR_PCTL0 USB_VSSA |
| USBDR_PCTL1 | I | ULPI—Use as USBDR_PCTL1 USB_VDDA |
| USBDR_CLK | I | ULPI—Use as USBDR_CLK |

17.2.1 UTMI Interface

UTMI was developed to specify a standard interface between USB 2.0 controllers and USB 2.0 PHY's. This interface is made available to support applications that may use a UTMI-compliant PHY. UTMI+ extensions are not supported by the USB DR module. Functionality added by UTMI+ is available in the ULPI interface. Only the integrated PHY uses the UTMI interface; therefore, there are no external UTMI signals to be documented in this manual. The integrated USB PHY has four dedicated external signals, which are only used when the device is a host: USBDR_DRIVE_VBUS, USBDR_PWRFAULT, USBDR_PCTL0, and USBDR_PCTL1. These are sideband signals that are also used as port of ULPI interface. [Table 17-2](#) describes the signals for the UTMI interface.

Table 17-2. UTMI Signal Descriptions

| Signal | I/O | Description | | | | |
|------------------|----------------------|--|---------------|-------------|----------|---------------------|
| USB_DP | I/O | USB 2.0 data positive signal | | | | |
| USB_DM | I/O | USB 2.0 data negative signal | | | | |
| USB_VBUS | I/O | USB 2.0 VBUS Signal | | | | |
| USB_RBIAS | I | The USB_RBIAS input is used by the analog bias block for generating an accurate bias current. An external precision resistor of $10K\Omega \pm 1\%$ is connected between this pin and ground. | | | | |
| USB_TPA | I/O | This is a test pin. It should be left unconnected. | | | | |
| USBDR_DRIVE_VBUS | I/O | This pin is used to enable/disable power (VBus) on applications supporting port power switching. This signal is not applicable to device only applications. | | | | |
| | | <table border="1"> <thead> <tr> <th>State Meaning</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Asserted</td> <td>Vbus power enabled.</td> </tr> <tr> <td>Negated</td> <td>Vbus power disabled.</td> </tr> </tbody> </table> | State Meaning | Description | Asserted | Vbus power enabled. |
| State Meaning | Description | | | | | |
| Asserted | Vbus power enabled. | | | | | |
| Negated | Vbus power disabled. | | | | | |

Table 17-2. UTMI Signal Descriptions (continued)

| Signal | I/O | Description |
|----------------|-----|--|
| USBDR_PWRFAULT | I | Power fault. USBDR_PWRFAULT indicates whether a power fault occurred on the USB port Vbus. |
| | | State Meaning Asserted—Indicates that a Vbus fault occurred. Applications that support power switching must shut down Vbus power. Negated—Indicates normal operation. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_PCTL0 | O | Port control 0. USBDR_PCTL0 controls the port status indicator LED 0 when in host mode. |
| | | State Meaning Asserted—LED on. Negated—LED off. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_PCTL1 | O | Port control 1. USBDR_PCTL1 controls the port status indicator LED 1 when in host mode. |
| | | State Meaning Asserted—LED on. Negated—LED off. |
| | | Timing Synchronous to PHY_CLK. |

17.2.2 ULPI Interface

The ULPI (UTMI low pin count interface) is a reduced pin-count (12 signals) extension of the UTMI+ specification. Pin count is reduced by converting relatively static signals to register bits, and providing a bidirectional, generic data bus that carries USB and register data. This interface minimizes pin count requirements for external PHYs. [Table 17-3](#) describes the signals for the ULPI interface.

Table 17-3. ULPI Signal Descriptions

| Signal | I/O | Description |
|-----------|-----|---|
| USBDR_DIR | I | Direction. USBDR_DIR controls the direction of the data bus. When the PHY has data to transfer to USB port, it drives USBDR_DIR high to take ownership of the bus. When the PHY has no data to transfer it drives USBDR_DIR low and monitors the bus for link activity. The PHY pulls USBDR_DIR high whenever the interface cannot accept data from the link, for example, when the internal PHY PLL is not stable. |
| | | State Meaning Asserted—PHY has data to transfer to the link. Negated—PHY has no data to transfer. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_NXT | I | Next data. The PHY asserts USBDR_NXT to throttle the data. When USB port is sending data to the PHY, USBDR_NXT indicates when the current byte has been accepted by the PHY. The USB port places the next byte on the data bus in the following clock cycle. When the PHY is sending data to USB port, USBDR_NXT indicates when a new byte is available for USB port to consume. |
| | | State Meaning Asserted—PHY is ready to transfer byte. Negated—PHY is not ready. |
| | | Timing Synchronous to PHY_CLK. |

Table 17-3. ULPI Signal Descriptions (continued)

| Signal | I/O | Description |
|------------------|-----|---|
| USBDR_STP | O | Stop. USBDR_STP indicates the end of a transfer on the bus. |
| | | State Meaning Asserted—USB asserts this signal for 1 clock cycle to stop the data stream currently on the bus. If USB port is sending data to the PHY, USBDR_STP indicates the last byte of data was previously on the bus. If the PHY is sending data to USB port, USBDR_STP forces the PHY to end its transfer, negate USBDR_DIR and relinquish control of the data bus to the USB port. Negated—Indicates normal operation. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_PWRFAULT | I | Power fault. USBDR_PWRFAULT indicates whether a power fault occurred on the USB port Vbus. |
| | | State Meaning Asserted—Indicates that a Vbus fault occurred. Applications that support power switching must shut down Vbus power. Negated—Indicates normal operation. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_PCTL0 | O | Port control 0. USBDR_PCTL0 controls the port status indicator LED 0 when in host mode. |
| | | State Meaning Asserted—LED on. Negated—LED off. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_PCTL1 | O | Port control 1. USBDR_PCTL1 controls the port status indicator LED 1 when in host mode. |
| | | State Meaning Asserted—LED on. Negated—LED off. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_TXRXD[7:0] | I/O | Data bit <i>n</i> . USBDR_TXRXD _{<i>n</i>} is bit <i>n</i> of the 8-bit (USBDR_TXRXD7–USBDR_TXRXD0), uni-directional data bus used to carry USB register, and interrupt data between the PHY and the USB controller. |
| | | State Meaning Asserted—Data bit <i>n</i> is 1. Negated—Data bit <i>n</i> is 0. |
| | | Timing Synchronous to PHY_CLK. |

17.2.3 PHY Clocks

The USBDR_CLK input provides the clocking signal for the ULPI PHY interface. The clock is 60 MHz. Detailed clock specifications are given in the appropriate hardware specifications document.

17.3 Memory Map/Register Definitions

This section provides the memory map and detailed descriptions of all USB interface registers. The memory map of the USB interface is shown in [Table 17-4](#).

Table 17-4. USB Interface Memory Map

| Offset | Register | Access | Reset | Section/Page |
|--|--|--------|-------------|---------------------------------|
| USB DR Controller—Block Base Address 0x2_3000 | | | | |
| 0x000–0x0FF | Reserved, should be cleared | — | — | — |
| 0x100 | CAPLENGTH—Capability register length | R | 0x40 | 17.3.1.1/17-9 |
| 0x102 | HCVERSION—Host interface version number ¹ | R | 0x0100 | 17.3.1.2/17-9 |
| 0x104 | HCSPARAMS—Host controller structural parameters ¹ | R | 0x0001_0011 | 17.3.1.3/17-10 |
| 0x108 | HCCPARAMS—Host controller capability parameters ¹ | R | 0x0000_0006 | 17.3.1.4/17-10 |
| 0x120 | DCVERSION—Device interface version number | R | 0x0001 | 17.3.1.5/17-11 |
| 0x124 | DCCPARAMS—Device controller parameters | R | 0x0000_0183 | 17.3.1.6/17-12 |
| 0x140 | USBCMD—USB command | Mixed | 0x0008_0000 | 17.3.2.1/17-13 |
| 0x144 | USBSTS—USB status | Mixed | 0x0000_0000 | 17.3.2.2/17-15 |
| 0x148 | USBINTR—USB interrupt enable | R/W | 0x0000_0000 | 17.3.2.3/17-17 |
| 0x14C | FRINDEX—USB frame index | R/W | 0x0000_0000 | 17.3.2.4/17-19 |
| 0x154 | PERIODICLISTBASE—Frame list base address ¹ | R/W | 0x0000_0000 | 17.3.2.6/17-20 |
| | DEVICEADDR—USB device address | R/W | 0x0000_0000 | 17.3.2.7/17-21 |
| 0x158 | ASYNCLISTADDR—Next asynchronous list addr (host mode) ¹ | R/W | 0x0000_0000 | 17.3.2.8/17-22 |
| | ENDPOINTLISTADDR—Address at endpoint list (device mode) | R/W | 0x0000_0000 | 17.3.2.9/17-22 |
| 0x160 | BURSTSIZE—Programmable burst size | R/W | 0x0000_1010 | 17.3.2.10/17-23 |
| 0x164 | TXFILLTUNING—Host TT transmit pre-buffer packet tuning | R/W | 0x0000_0000 | 17.3.2.11/17-24 |
| 0x170 | ULPI VIEWPORT—ULPI Register Access | Mixed | 0x0000_0000 | 17.3.2.12/17-25 |
| 0x180 | CONFIGFLAG—Configured flag register | R | 0x0000_0001 | 17.3.2.13/17-27 |
| 0x184 | PORTSC—Port status/control | Mixed | 0x1000_0000 | 17.3.2.14/17-27 |
| 0x1A4 | OTGSC—On-The-Go status and control ¹ | Mixed | 0x200C_0000 | 17.3.2.15/17-32 |
| 0x1A8 | USBMODE—USB device mode | R/W | 0x0000_0000 | 17.3.2.16/17-35 |
| 0x1AC | ENDPTSETUPSTAT—Endpoint setup status | R/W | 0x0000_0000 | 17.3.2.17/17-36 |
| 0x1B0 | ENDPOINTPRIME—Endpoint initialization | R/W | 0x0000_0000 | 17.3.2.18/17-36 |
| 0x1B4 | ENDPTFLUSH—Endpoint de-initialize | R/W | 0x0000_0000 | 17.3.2.19/17-37 |
| 0x1B8 | ENDPTSTATUS—Endpoint status | R | 0x0000_0000 | 17.3.2.20/17-38 |
| 0x1BC | ENDPTCOMPLETE—Endpoint complete | w1c | 0x0000_0000 | 17.3.2.21/17-38 |
| 0x1C0 | ENDPTCTRL0—Endpoint control 0 | Mixed | 0x0080_0080 | 17.3.2.22/17-39 |
| 0x1C4 | ENDPTCTRL1—Endpoint control 1 | R/W | 0x0000_0000 | 17.3.2.23/17-40 |
| 0x1C8 | ENDPTCTRL2—Endpoint control 2 | R/W | 0x0000_0000 | 17.3.2.23/17-40 |

Table 17-4. USB Interface Memory Map (continued)

| Offset | Register | Access | Reset | Section/Page |
|-----------------|------------------------------------|--------|-------------|---------------------------------|
| 0x1CA– 0x1D4 | Reserved | — | — | — |
| 0x400 | SNOOP1—Snoop 1 | R/W | 0x0000_0000 | 17.3.2.24/17-42 |
| 0x404 | SNOOP2—Snoop 2 | R/W | 0x0000_0000 | 17.3.2.24/17-42 |
| 0x408 | AGE_CNT_THRESH—Age count threshold | R/W | 0x0000_0000 | 17.3.2.25/17-43 |
| 0x40C | PRI_CTRL—Priority control | R/W | 0x0000_0000 | 17.3.2.26/17-44 |
| 0x410 | SI_CTRL—System interface control | R/W | 0x0000_0000 | 17.3.2.27/17-45 |
| 0x500 | CONTROL—Control | Mixed | 0x0000_0000 | 17.3.2.28/17-46 |
| 0x504– 0xFFF | Reserved, should be cleared | — | — | — |

¹ This register has separate functions for the host and device operation; the host function is listed first in the table.

The following sections provide details about the registers in the USB memory map.

NOTE

Memory may be viewed from either a big-endian or little-endian byte ordering perspective depending on the processor configuration. In big-endian mode, the most-significant byte of word 0 is located at address 0 and the least-significant byte of word 0 is located at address 3. In little-endian mode, the least-significant byte of word 0 is located at address 0 and the most-significant byte of word 0 is located at address 3. Within registers, bits are numbered within a word starting with bit 31 as the most-significant bit. By convention USB registers use little-endian byte ordering. In the USB DR module, these are the registers from offsets 0x00 to 0x1FF. The registers associated with the internal system interface (0x400 and above) use big-endian byte ordering.

17.3.1 Capability Registers

The capability registers specify the software limits, restrictions, and capabilities of the host/device controller implementation. Most of these registers are defined by the EHCI specification. Registers that are not defined by the EHCI specification are noted in their descriptions.

17.3.1.1 Capability Registers Length (CAPLENGTH)

CAPLENGTH is used as an offset to add to the register base address to find the beginning of the operational register space, that is, the location of the USBCMD register. Figure 17-2 shows CAPLENGTH.




Figure 17-2. Capability Registers Length (CAPLENGTH)

Table 17-5 provides bit descriptions for the CAPLENGTH register.

Table 17-5. CAPLENGTH Register Field Descriptions

| Bits | Name | Description |
|------|-----------|---|
| 7-0 | CAPLENGTH | Capability registers length. Value is 0x40. |

17.3.1.2 Host Controller Interface Version (HCIVERSION)

HCIVERSION contains a BCD encoding of the EHCI revision number supported by this host controller. The most-significant byte of the register represents a major revision and the least-significant byte is the minor revision. Figure 17-3 shows the HCIVERSION register.




Figure 17-3. Host Controller Interface Version (HCIVERSION)

Table 17-6 provides bit descriptions for the HCIVERSION register.

Table 17-6. HCIVERSION Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 15-0 | — | EHCI revision number. Value is 0x0100 indicating version 1.0. |

17.3.1.3 Host Controller Structural Parameters (HCSPARAMS)

HCSPARAMS contains structural parameters such as the number of downstream ports. Figure 17-4 shows the HCSPARAMS register.




Figure 17-4. Host Controller Structural Parameters (HCSPARAMS)

Table 17-7 provides bit descriptions for the HCSPARAMS register.

Table 17-7. HCSPARAMS Register Field Descriptions

| Bits | Name | Description |
|-------|---------|---|
| 31–28 | — | Reserved, should be cleared. |
| 27–24 | N_TT | Number of transaction translators. This is a non-EHCI field. This field indicates the number of embedded transaction translators associated the module. Always 1. See Section 17.9.1, “Embedded Transaction Translator Function.” |
| 23–20 | N_PTT | Ports per transaction translator. This is a non-EHCI field. The number of ports assigned to each transaction translator. This is equal to N_PORTS. |
| 19–17 | — | Reserved, should be cleared. |
| 16 | PI | Port indicators. Indicates whether the ports support port indicator control. Always 1. 1 The port status and control registers include a R/W field for controlling the state of the port indicator. |
| 15–12 | N_CC | Number of companion controllers associated with the DR controller. Always 0. |
| 11–8 | N_PCC | Number ports per CC. This field indicates the number of ports supported per internal companion controller. Always 0. |
| 7–5 | — | Reserved, should be cleared. |
| 4 | PPC | Power port control. Indicates whether the host controller supports port power control. Always 1. 1 Ports have power port switches. |
| 3–0 | N_PORTS | Number of ports. Number of physical downstream ports implemented for host applications. The value of this field determines how many port registers are addressable in the operational register. Always 1. |

17.3.1.4 Host Controller Capability Parameters (HCCPARAMS)

HCCPARAMS identifies multiple mode control (time-base bit functionality) addressing capability. Figure 17-5 shows the HCCPARAMS register.




Figure 17-5. Host Control Capability Parameters (HCCPARAMS)

Table 17-8 provides bit descriptions for the HCCPARAMS register.

Table 17-8. HCCPARAMS Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–16 | — | Reserved, should be cleared. |
| 15–8 | EECP | EHCI extended capabilities pointer. Indicates the existence of a capabilities list. A value of 0x00 indicates no extended capabilities are implemented. A non-zero value in this register indicates the offset in PCI configuration space of the first EHCI extended capability. The pointer value must be 0x40 or greater if implemented to maintain the consistency of the PCI header defined for this class of device. This field is always 0. |
| 7–4 | IST | Isochronous scheduling threshold. Indicates, relative to the current position of the executing host controller, where software can reliably update the isochronous schedule. When bit 7 is zero, the value of the least significant 3 bits indicates the number of microframes a host controller can hold a set of isochronous data structures (one or more) before flushing the state. When bit 7 is a one, then host software assumes the host controller may cache an isochronous data structure for an entire frame. This field is always 0. |
| 3 | — | Reserved, should be cleared. |
| 2 | ASP | Asynchronous schedule park capability. Indicates whether the USB DR module supports the park feature for high-speed queue heads in the asynchronous schedule. The feature can be disabled or enabled and set to a specific level by using the asynchronous schedule park mode enable and asynchronous schedule park mode count fields in the USBCMD register. This field is always 1 (park feature supported). |
| 1 | PFL | Programmable frame list flag. Indicates whether system software can specify and use a frame list length less than 1024 elements. Frame list size is configured via the USBCMD register frame list size field. The frame list must always be aligned on a 4-K page boundary. This requirement ensures that the frame list is always physically contiguous. This field is always 1. |
| 0 | ADC | 64-bit addressing capability. Always 0; 64-bit addressing is not supported. 0 Data structures use 32-bit address memory pointers |

17.3.1.5 Device Controller Interface Version (DCIVERSION)—Non-EHCI

This register is not defined in the EHCI specification. DCIVERSION is a two-byte register containing a BCD encoding of the device controller interface. The most-significant byte of the register represents a major revision and the least-significant byte is the minor revision. Figure 17-6 shows the DCIVERSION register.

Offset 0x120

Access: Read-only




Figure 17-6. Device Interface Version (DCIVERSION)

Table 17-9 provides bit descriptions for the DCIVERSION register.

Table 17-9. DCIVERSION Register Field Descriptions

| Bits | Name | Description |
|------|------------|-----------------------------------|
| 15–0 | DCIVERSION | Device interface revision number. |

17.3.1.6 Device Controller Capability Parameters (DCCPARAMS)—Non-EHCI

This register is not defined in the EHCI specification. This register describes the overall host/device capability of the DR module. Figure 17-7 shows the DCCPARAMS register.

Offset 0x124

Access: Read-only




Figure 17-7. Device Control Capability Parameters (DCCPARAMS)

Table 17-10 provides bit descriptions for the DCCPARAMS register.

Table 17-10. DCCPARAMS Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 31–9 | — | Reserved, should be cleared. |
| 8 | HC | Host capable. Always 1, indicating the USB DR controller can operate as an EHCI compatible USB 2.0 host |
| 7 | DC | Device capable. Always 1, indicating the USB DR controller can operate as an USB 2.0 device. 1 Device capability. 0 No device capability (host only). |
| 6–5 | — | Reserved, should be cleared. |
| 4–0 | DEN | Device endpoint number. Indicates the number of endpoints built into the device controller. Always 0x3. |

17.3.2 Operational Registers

The operational registers are comprised of dynamic control or status registers that may be read-only, read/write, or read/write-1-to-clear. The following sections define the operational registers.

Table 17-11. USBCMD Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 10 | — | Reserved, should be cleared. |
| 9–8 | ASP | Asynchronous schedule park mode count. This field defaults to 0x3 and is R/W. It contains a count of the number of successive transactions the host controller is allowed to execute from a high-speed queue head on the Asynchronous schedule before continuing traversal of the Asynchronous schedule. Valid values are 0x1H to 0x3H. Software must not write a zero to this field when ASPE is set as this results in undefined behavior. |
| 7 | LR | Light host/device controller reset (OPTIONAL). Not implemented. Always 0. |
| 6 | IAA | Interrupt on async advance doorbell. Used as a doorbell by software to tell the USB DR controller to issue an interrupt the next time it advances asynchronous schedule. Software must write a 1 to this bit to ring the doorbell. When the controller has evicted all appropriate cached schedule states, it sets USBSTS[AAI]. If USBINTR[AAE] is set, the host controller asserts an interrupt at the next interrupt threshold. The controller clears this bit after it has set USBSTS[AAI]. Software should not set this bit when the asynchronous schedule is inactive. Doing so yields undefined results. This bit is only used in host mode. Setting this bit when the USB DR module is in device mode is selected results in undefined results. |
| 5 | ASE | Asynchronous schedule enable. Controls whether the controller skips processing the asynchronous schedule. Only used in host mode. 0 Do not process the asynchronous schedule 1 Use the ASYNCLISTADDR register to access the asynchronous schedule. |
| 4 | PSE | Periodic schedule enable. Controls whether the controller skips processing the periodic schedule. Only used in host mode. 0 Do not process the periodic schedule. 1 Use the PERIODICLISTBASE register to access the periodic schedule. |
| 3–2 | FS | Frame list size. Together with bit 15 these bits make the FS[2:0] field. This field is read/write only if programmable frame list flag in the HCCPARAMS registers is set to 1. This field specifies the size of the frame list that controls which bits in FRINDEX should be used for the frame list current index. Only used in host mode. Note that values below 256 elements are not defined in the EHCI specification. 000 1024 elements (4096 bytes) 001 512 elements (2048 bytes) 010 256 elements (1024 bytes) 011 128 elements (512 bytes) 100 64 elements (256 bytes) 101 32 elements (128 bytes) 110 16 elements (64 bytes) 111 8 elements (32 bytes) |

Table 17-12 shows the USBSTS register field descriptions.

Table 17-12. USBSTS Register Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 31–16 | — | Reserved, should be cleared. |
| 15 | AS | Asynchronous schedule status. Reports the current real status of the asynchronous schedule. The USB DR controller is not required to immediately disable or enable the asynchronous schedule when software transitions USBCMD[ASE]. When this bit and USBCMD[ASE] have the same value, the asynchronous schedule is either enabled (1) or disabled (0). Only used in host mode. 0 Disabled 1 Enabled |
| 14 | PS | Periodic schedule status. Reports the current real status of the periodic schedule. The USB DR controller is not required to immediately disable or enable the periodic schedule when software transitions USBCMD[PSE]. When this bit and USBCMD[PSE] have the same value, the periodic schedule is either enabled (1) or disabled (0). Only used in host mode. 0 Disabled 1 Enabled |
| 13 | RCL | Reclamation. Used to detect an empty asynchronous schedule. Only used by the host mode. 0 Non-empty asynchronous schedule 1 Empty asynchronous schedule |
| 12 | HCH | HC halted. This bit is a zero whenever USBCMD[RS] is a one. The USB DR controller sets this bit to one after it has stopped executing because of USBCMD[RS] being cleared, either by software or by the host controller hardware (for example, internal error). Only used in host mode. 0 Running 1 Halted |
| 11 | — | Reserved, should be cleared. |
| 10 | ULPII | ULPI interrupt. An event completion to the viewport register sets this bit. If the ULPI enables the USBINTR[ULPIE] to be set, the USB interrupt (UI) occurs. |
| 9 | — | Reserved, should be cleared. |
| 8 | SLI | DCSuspend. This is a non-EHCI bit. When a device controller enters a suspend state from an active state, this bit is set. The device controller clears the bit upon exiting from a suspend state. Only used by the device controller. 0 Active 1 Suspended |
| 7 | SRI | Host mode: <ul style="list-style-type: none"> This is a non-EHCI status bit. In host mode, this bit is set every 125 us, provided the PHY clock is present and running (for example, the port is NOT suspended), and can be used by the host controller driver as a time base. Device mode: <ul style="list-style-type: none"> SOF received. When the USB DR controller detects a Start Of (Micro)Frame, this bit is set. When a SOF is extremely late, the DR controller automatically sets this bit to indicate that an SOF was expected. Therefore, this bit is set roughly every 1 msec in device FS mode and every 125 msec in HS mode and is synchronized to the actual SOF that is received. Because the controller is initialized to FS before connect, this bit is set at an interval of 1 msec during the prelude to the connect and chirp. Software writes a 1 to this bit to clear it. |

Table 17-12. USBSTS Register Field Descriptions (continued)

| Bits | Name | Description |
|------|--------------------|--|
| 6 | URI | USB reset received. This is a non-EHCI bit. When the USB DR controller detects a USB reset and enters the default state, this bit is set. Software can write a 1 to this bit to clear the USB reset received status bit. Only used by the device mode. 0 No reset received 1 Reset received |
| 5 | AAI | Interrupt on async advance. System software can force the controller to issue an interrupt the next time the USB DR controller advances the asynchronous schedule by writing a one to USBCMD[IAA]. This status bit indicates the assertion of that interrupt source. Only used by the host mode. 0 No async advance interrupt 1 Async advance interrupt |
| 4 | SEI | System error. This bit is set whenever an error is detected on the system bus. If USBINTR[SEE] is set, an interrupt is generated. The interrupt and status bits remain asserted until cleared by writing a 1 to this bit. Additionally, when in host mode, USBCMD[RS] is cleared, effectively disabling the USB DR controller. For the USB DR controller in device mode, an interrupt is generated, but no other action is taken. 0 Normal operation 1 Error |
| 3 | FRI | Frame list rollover. The controller sets this bit to a one when the frame list index rolls over from its maximum value to zero. The exact value at which the rollover occurs depends on the frame list size. For example, if the frame list size (as programmed in USBCMD[FS]) is 1024, FRINDEX rolls over every time FRINDEX [1 3] toggles. Similarly, if the size is 512, the USB DR controller sets this bit to a one every time FHINDEX [12] toggles. Only used by the host mode. |
| 2 | PCI | Host mode: <ul style="list-style-type: none"> Port change detect. The controller sets this bit when a connect status occurs on any port, a port enable/disable change occurs, an over current change occurs, or PORTSC[FPR] is set as the result of a J-K transition on the suspended port. Device mode: <ul style="list-style-type: none"> The USB DR controller sets this bit when it enters the full or high-speed operational state. When the it exits the full or high-speed operation states due to reset or suspend events, the notification mechanisms are USBSTS[URI] and USBSTS[SLI], respectively. This bit is not EHCI compatible. |
| 1 | UEI (USBERRINT) | USB error interrupt (USBERRINT). When completion of a USB transaction results in an error condition, this bit is set by the controller. This bit is set along with the UI, if the TD on which the error interrupt occurred also had its interrupt on complete (IOC) bit set. See Section 4.15.1 in EHCI for a complete list of host error interrupt conditions. Also see Table 17-92 in this chapter for more information on device error matrix. For the USB DR controller in device mode, only resume signaling is detected, all others are ignored. 0 No error 1 Error detected |
| 0 | UI (USBINT) | USB interrupt (USBINT). This bit is set by the controller when the cause of an interrupt is a completion of a USB transaction where the transfer descriptor (TD) has an interrupt on complete (IOC) bit set. This bit is also set by the controller when a short packet is detected. A short packet is when the actual number of bytes received was less than the expected number of bytes. |

17.3.2.3 USB Interrupt Enable Register (USBINTR)

The interrupts to software are enabled with the USB interrupt enable register, shown in [Figure 17-10](#). An interrupt is generated when a bit is set and the corresponding interrupt is active. The USB status register

(USBSTS) still shows interrupt sources even if they are disabled by the USBINTR register, allowing polling of interrupt events by the software.

Offset 0x148

Access: Read/Write

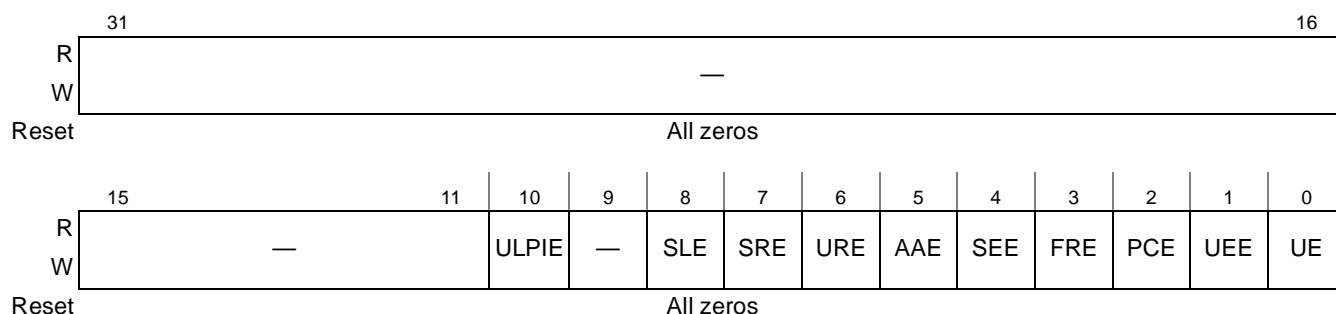


Figure 17-10. USB Interrupt Enable (USBINTR)

Table 17-13 shows the USBINTR register field descriptions.

Table 17-13. USBINTR Register Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 31–11 | — | Reserved, should be cleared. |
| 10 | ULPIE | ULPI interrupt enable. An event completion to the viewport register sets the USBSTS[ULPII]. If the ULPI enables ULPIE bit to be set, then the USBINT (USBSTS[UI]) occurs. 0 Disable 1 Enable |
| 9 | — | Reserved, should be cleared. |
| 8 | SLE | Sleep enable. This is a non-EHCI bit. When this bit is a one, and USBSTS[SLI] transitions, the USB DR controller issues an interrupt. The interrupt is acknowledged by software writing a one to USBSTS[SLI]. Only used in device mode. 0 Disable 1 Enable |
| 7 | SRE | SOF received enable. This is a non-EHCI bit. When this bit is a one, and USBSTS[SRI] is a one, the controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[SRI]. 0 Disable 1 Enable |
| 6 | URE | USB reset enable. This is a non-EHCI bit. When this bit is a one, USBSTS[URI] is a one, the device controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[URI] bit. Only used in device mode. 0 Disable 1 Enable |
| 5 | AAE | Interrupt on async advance enable. When this bit is a one, and USBSTS[AAI] is a one, the controller issues an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[AAI]. Only used in host mode. 0 Disable 1 Enable |
| 4 | SEE | System error enable. When this bit is a one, and USBSTS[SEI] is a one, the controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[SEI]. 0 Disable 1 Enable |

Table 17-13. USBINTR Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 3 | FRE | Frame list rollover enable. When this bit is a one, and USBSTS[FRI] is a one, the controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[FRI]. Only used by the host mode. 0 Disable 1 Enable |
| 2 | PCE | Port change detect enable. When this bit is a one, and USBSTS[PCI] is a one, the controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[PCI]. 0 Disable 1 Enable |
| 1 | UEE | USB error interrupt enable. When this bit is a one, and USBSTS[UEI] is a one, the controller issues an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[UEI]. 0 Disable 1 Enable |
| 0 | UE | USB interrupt enable. When this bit is a one, and USBSTS[UI] is a one, the DR controller issues an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[UI]. 0 Disable 1 Enable |

17.3.2.4 Frame Index Register (FRINDEX)

In host mode, the frame index register is used by the controller to index the periodic frame list. The register updates every 125 microseconds (once each microframe). Bits N–3 are used to select a particular entry in the periodic frame list during periodic schedule execution. The number of bits used for the index depends on the size of the frame list as set by system software in USBCMD[FS].

This register must be written as a DWord. Byte writes produce undefined results. This register cannot be written unless the USB DR controller is in the Halted state as indicated by the USBSTS[HCH]. A write to this register while USBCMD[RS] is set produces undefined results. Writes to this register also affect the SOF value.

In device mode, this register is read-only and the USB DR controller updates the FRINDEX[13–3] register from the frame number indicated by the SOF marker. Whenever a SOF is received by the USB bus, FRINDEX[13–3] is checked against the SOF marker. If FRINDEX[13–3] is different from the SOF marker, FRINDEX[13–3] is set to the SOF value and FRINDEX[2–0] is cleared (that is, SOF for 1 msec frame). If FRINDEX[13–3] is equal to the SOF value, FRINDEX[2–0] is incremented (that is, SOF for 125- μ sec microframe).

Figure 17-11 shows the USB frame index register.

Offset 0x14C

Access: Read/Write



Figure 17-11. USB Frame Index (FRINDEX)

Table 17-14 shows the FRINDEX register field descriptions.

Table 17-14. FRINDEX Register Field Descriptions

| Bits | Name | Description |
|-------|---------|---|
| 31–14 | — | Reserved, should be cleared. |
| 13–0 | FRINDEX | Frame index. The value in this register increments at the end of each time frame (for example, microframe). Bits N–3 are used for the Frame List current index. This means that each location of the frame list is accessed 8 times (frames or microframes) before moving to the next index. In device mode, the value is the current frame number of the last frame transmitted. It is not used as an index. In either mode, bits 2–0 indicate the current microframe. |

Table 17-15 illustrates values of N based on the value of the Frame List Size in the USBCMD register, when used in host mode.

Table 17-15. FRINDEX N Values

| USBCMD[FS] | Frame List Size | FRINDEX N value |
|------------|----------------------------|-----------------|
| 000 | 1024 elements (4096 bytes) | 12 |
| 001 | 512 elements (2048 bytes) | 11 |
| 010 | 256 elements (1024 bytes) | 10 |
| 011 | 128 elements (512 bytes) | 9 |
| 100 | 64 elements (256 bytes) | 8 |
| 101 | 32 elements (128 bytes) | 7 |
| 110 | 16 elements (64 bytes) | 6 |
| 111 | 8 elements (32 bytes) | 5 |

17.3.2.5 Control Data Structure Segment Register (CTRLDSSEGMENT)

The CTRLDSSEGMENT register is not implemented on this device.

17.3.2.6 Periodic Frame List Base Address Register (PERIODICLISTBASE)

This register contains the beginning address of the Periodic Frame List in the system memory. The host controller driver loads this register prior to starting the schedule execution by the controller. The memory structure referenced by this physical memory pointer is assumed to be 4-Kbyte aligned. The contents of this register are combined with the frame index register (FRINDEX) to enable the controller to step through the Periodic Frame List in sequence.

Note that this register is shared between the host and device mode functions. In host mode, it is the PERIODICLISTBASE register; in device mode, it is the DEVICEADDR register. See [Section 17.3.2.7, “Device Address Register \(DEVICEADDR\)—Non-EHCI,”](#) for more information.

17.3.2.8 Current Asynchronous List Address Register (ASYNCLISTADDR)

This 32-bit register contains the address of the next asynchronous queue head to be executed by the host. Bits 4–0 of this register cannot be modified by the system software and always return zeros when read.

Note that this register is shared between the host and device mode functions. In host mode, it is the ASYNCLISTADDR register; in device mode, it is the ENDPOINTLISTADDR register. See [Section 17.3.2.9, “Endpoint List Address Register \(ENDPOINTLISTADDR\)—Non-EHCI,”](#) for more information.

Figure 17-14 shows the current asynchronous list address register.




Figure 17-14. Current Asynchronous List Address (ASYNCLISTADDR)

Table 17-18 describes the current asynchronous list address register.

Table 17-18. ASYNCLISTADDR Register Field Descriptions

| Bits | Name | Description |
|------|---------|--|
| 31–5 | ASYBASE | Link pointer low (LPL). These bits correspond to memory address signals [31:5]. This field may only reference a queue head (QH). Only used by the host controller. |
| 4–0 | — | Reserved, should be cleared. |

17.3.2.9 Endpoint List Address Register (ENDPOINTLISTADDR)—Non-EHCI

The endpoint list address register is not defined in the EHCI specification. In device mode, this register contains the address of the top of the endpoint list in system memory. Bits 10–0 of this register cannot be modified by the system software and always return zeros when read. The memory structure referenced by this physical memory pointer is assumed to be 64-bytes. The queue head is actually a 48-byte structure, but must be aligned on 64-byte boundary. However, the ENDPOINTLISTADDR[EPBASE] has a granularity of 2 Kbytes, so in practice the queue head should be 2-Kbyte aligned.

Note that this register is shared between the host and device mode functions. In device mode, it is the ENDPOINTLISTADDR register; in host mode, it is the ASYNCLISTADDR register. See [Section 17.3.2.8, “Current Asynchronous List Address Register \(ASYNCLISTADDR\),”](#) for more information.

Figure 17-15 shows the endpoint list address register.

Offset 0x158

Access: Read/Write




Figure 17-15. Endpoint List Address (ENDPOINTLISTADDR)

Table 17-19 describes the endpoint list address register fields.

Table 17-19. ENDPOINTLISTADDR Register Field Descriptions

| Bits | Name | Description |
|-------|--------|---|
| 31–11 | EPBASE | Endpoint list address. Address of the top of the endpoint list. |
| 10–0 | — | Reserved, should be cleared. |

17.3.2.10 Master Interface Data Burst Size Register (BURSTSIZE)—Non-EHCI

The master interface data burst size register, shown in Figure 17-16, is not defined in the EHCI specification. This register is used to control and dynamically change the burst size used during data movement on the initiator (master) interface.

Offset 0x160

Access: Read/Write




Figure 17-16. Master Interface Data Burst Size (BURSTSIZE)

Table 17-20 describes the master interface data burst size register fields.

Table 17-20. BURSTSIZE Register Field Descriptions

| Bits | Name | Description |
|-------|----------|--|
| 31–16 | — | Reserved, should be cleared. |
| 15–8 | TXPBURST | Programable TX burst length. This register represents the maximum length of a burst in 32-bit words while moving data from system memory to the USB bus. Must not be set to greater than 16. |
| 7–0 | RXPBURST | Programable RX burst length. This register represents the maximum length of a burst in 32-bit words while moving data from the USB bus to system memory. Must not be set to greater than 16. |

Table 17-21. TXFILLTUNING Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|-------------|---|
| 21–16 | TXFIFOTHRES | FIFO burst threshold. Control the number of data bursts that are posted to the TX latency FIFO in host mode before the packet begins on to the bus. The minimum value is 2 and this value should be a low as possible to maximize USB performance. A higher value can be used in systems with unpredictable latency and/or insufficient bandwidth where the FIFO may underrun because the data transferred from the latency FIFO to USB occurs before it can be replenished from system memory. This value is ignored if USBMODE[SDIS] (stream disable bit) is set. When USBMODE[SDIS] is set, the host controller behaves as if TXFIFOTHRES is set to the maximum value. |
| 15–13 | — | Reserved, should be cleared. |
| 12–8 | TXSCHHEALTH | Scheduler health counter. Increment when the host controller fails to fill the TX latency FIFO to the level programmed by TXFIFOTHRES before running out of time to send the packet before the next Start-Of-Frame. This health counter measures the number of times this occurs to provide feedback to selecting a proper TXSCHOH. Writing to this register clears the counter and this counter stops counting after reaching the maximum of 31. |
| 7–0 | TXSCHOH | Scheduler overhead. These bits add an additional fixed offset to the schedule time estimator described above as T_{ff} . As an approximation, the value chosen for this register should limit the number of back-off events captured in the TXSCHHEALTH to less than 10 per second in a highly utilized bus. Choosing a value that is too high for this register is not desired as it can needlessly reduce USB utilization. The time unit represented in this register is 1.267 μ s when a device is connected in high-speed mode. The time unit represented in this register is 6.333 μ s when a device is connected in low-/full-speed mode. For most applications, TXSCHOH can be set to 4 or less. A good value to begin with is: $TXFIFOTHRES \times (BURSTSIZE \times 4 \text{ bytes-per-word}) \div (40 \times TimeUnit)$, always rounded to the next higher integer. <i>TimeUnit</i> is either 1.267 or 6.333 as noted earlier in this description. For example, if TXFIFOTHRES is 5 and BURSTSIZE is 8, then set TXSCHOH to $5 \times (8 \times 4) \div (40 \times 1.267) = 4$ for a high-speed link. If this value of TXSCHOH results in a TXSCHHEALTH count of 0 per second, try lowering the value by 1 if optimizing performance is desired. If TXSCHHEALTH exceeds 10 per second, try raising the value by 1. If streaming mode is disabled via the USBMODE register, treat TXFIFOTHRES as the maximum value for purposes of the TXSCHOH calculation. |

17.3.2.12 ULPI Register Access (ULPI VIEWPORT)

The ULPI register access provides indirect access to the ULPI PHY register set. Although the controller modules perform access to the ULPI PHY register set, there may be extraordinary circumstances where software may need direct access. Be advised that writes to the ULPI through the ULPI viewport can substantially harm standard USB operations. Currently no usage model has been defined where software should need to execute writes directly to the ULPI. Note that executing read operations through the ULPI viewport should have no harmful side effects to standard USB operations. Also note that if the ULPI interface is not enabled, this register will always read zeros.

force signaling and allows software to put the PHY into low power suspend mode and disable the PHY clock.

Offset 0x184

Access: Mixed

| | | | | | | | | | | | | | | | | |
|-------|-----|----|----|----|------|----|----|------|------|------|------|------|-----|----|-----|-----|
| | 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 16 | | |
| R | PTS | | — | — | PSPD | | — | PFSC | PHCD | WKOC | WKDS | WKCN | PTC | | | |
| W | | | | | | | | | | | | | | | | |
| Reset | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | |
| | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| R | PIC | | PO | PP | LS | | — | PR | SUSP | FPR | OCC | OCA | PEC | PE | CSC | CCS |
| W | | | | | | | | | | | w1c | | w1c | | w1c | |
| Reset | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

Figure 17-20. Port Status and Control (PORTSC)

Table 17-24 describes the PORTSC register fields.

Table 17-24. PORTSC Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–30 | PTS | Port transceiver select. This register bit is used to control which parallel transceiver interface is selected. 00 UTMI parallel interface 01 Reserved, should be cleared 10 ULPI parallel interface 11 Reserved This bit is not defined in the EHCI specification. |
| 29 | — | Reserved, should be cleared |
| 28 | — | Reserved |
| 27–26 | PSPD | Port speed. This read-only register field indicates the speed at which the port is operating. This bit is not defined in the EHCI specification. 00 Full-speed 01 Low-speed 10 High-speed 11 Undefined |
| 25 | — | Reserved, should be cleared |
| 24 | PFSC | Port force full-speed connect. Used to disable the chirp sequence that allows the port to identify itself as a HS port. This is useful for testing FS configurations with a HS host, hub or device. 0 Allow the port to identify itself as high speed. 1 Force the port to only connect at full speed. This bit is not defined in the EHCI specification. This bit is for debugging purposes. |

Table 17-24. PORTSC Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 23 | PHCD | PHY low power suspend. This bit is not defined in the EHCI specification. Host mode: <ul style="list-style-type: none"> The PHY can be put into low power suspend – when the downstream device has been put into suspend mode or when no downstream device is connected. Low power suspend is completely under the control of software. Device mode: <ul style="list-style-type: none"> The PHY can be put into low power suspend – when the device is not running (USBCMD[RS] = 0b) or suspend signaling is detected on the USB. Low power suspend is cleared automatically when the resume signaling has been detected or when forcing port resume. 0 Normal PHY operation. 1 Signal the PHY to enter low power suspend mode Reading this bit indicates the status of the PHY. Note: If there is no clock connected to the USBDR_CLK signals, PHCD must be set and the following registers should not be written: DEVICE_ADDR/PERIODICLISTBASE, PORTSC, ENDPTCTRL0, ENDPTCTRL1, ENDPTCTRL2. |
| 22 | WKOC | Wake on over-current enable. Writing this bit to a one enables the port to be sensitive to over-current conditions as wake-up events. This field is zero if Port Power (PP) is zero. This bit is used only in OTG/Host mode. |
| 21 | WKDS | Wake on disconnect enable. Writing this bit to a one enables the port to be sensitive to device disconnects as wake-up events. This field is zero if Port Power(PP) is zero or in device mode. This bit is used only in OTG/Host mode. |
| 20 | WKCN | Wake on connect enable. Writing this bit to a one enables the port to be sensitive to device connects as wake-up events. This field is zero if Port Power(PP) is zero or in device mode. This bit is used only in OTG/Host mode. |
| 19–16 | PTC | Port test control. Any other value than zero indicates that the port is operating in test mode. 0000 Not Enabled 0001 J_STATE 0010 K_STATE 0011 SEQ_NAK 0100 Packet 0101 FORCE_ENABLE 0110–1111 Reserved, should be cleared Refer to Chapter 7 of the USB Specification Revision 2.0 [3] for details on each test mode. |
| 15–14 | PIC | Port indicator control. Control the link indicator signals. These signals are valid for host mode only. 00 Off 01 Amber 10 Green 11 Undefined Refer to the USB Specification Revision 2.0 [3] for a description on how these bits are to be used. This field is output from the module on the USB port control signals for use by an external LED driving circuit. |
| 13 | PO | Port owner. Unconditionally goes to a 0 when the configured bit in the CONFIGFLAG register makes a 0 to 1 transition. This bit unconditionally goes to 1 whenever the Configured bit is zero. System software uses this field to release ownership of the port to a selected the module (in the event that the attached device is not a high-speed device). Software writes a one to this bit when the attached device is not a high-speed device. A one in this bit means that an internal companion controller owns and controls the port. Port owner hand-off is not implemented in this design, therefore this bit is always 0. |

Table 17-24. PORTSC Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 12 | PP | <p>Port power. Represents the current setting of the switch (0=off, 1=on). When power is not available on a port (that is, PP equals a 0), the port is non-functional and will not report attaches, detaches, and so on. When an over-current condition is detected on a powered port, the PP bit in each affected port is transitioned by the host controller driver from a one to a zero (removing power from the port). This feature is implemented in the host/OTG controller (PPC = 1). In a device-only implementation port power control is not necessary, thus PPC and PP = 0.</p> |
| 11–10 | LS | <p>Line status. Reflect the current logical levels of the USB D+ (bit 11) and D– (bit 10) signal lines. The use of line status by the host controller driver is not necessary (unlike EHCI), because the connection of FS and LS is managed by hardware.</p> <p>00 SE0 10 J-state 01 K-state 11 Undefined</p> |
| 9 | — | Reserved, should be cleared |
| 8 | PR | <p>Port reset.</p> <p>Host mode:</p> <ul style="list-style-type: none"> When software writes a one to this bit the bus-reset sequence as defined in the USB Specification Revision 2.0 is started. This bit will automatically change to zero after the reset sequence is complete. This behavior is different from EHCI where the host controller driver is required to set this bit to a zero after the reset duration is timed in the driver. <p>Device mode:</p> <ul style="list-style-type: none"> This bit is a read only status bit. Device reset from the USB bus is also indicated in the USBSTS register. <p>1 Port is in reset. 0 Port is not in reset. This field is zero if Port Power(PP) is zero.</p> |
| 7 | SUSP | <p>Suspend.</p> <p>Host mode:</p> <ul style="list-style-type: none"> The port enabled bit (PE) and suspend (SUSP) bit define the port states as follows: <p>0x Disable 10 Enable 11 Suspend</p> <ul style="list-style-type: none"> When in suspend state, downstream propagation of data is blocked on this port, except for port reset. The blocking occurs at the end of the current transaction if a transaction was in progress when this bit was written to 1. In the suspend state, the port is sensitive to resume detection. Note that the bit status does not change until the port is suspended and that there may be a delay in suspending a port if there is a transaction currently in progress on the USB. The module unconditionally sets this bit to zero when software clears the FPR bit. A write of zero to this bit is ignored by the host controller. If host software sets this bit to a one when the port is not enabled (that is, port enabled bit is a zero) the results are undefined. This field is zero if Port Power (PP) is zero in host mode. <p>Device mode:</p> <p>1 Port in suspend state. 0 Port not in suspend state. Default. In device mode this bit is a read-only status bit.</p> |

Table 17-24. PORTSC Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 6 | FPR | <p>Force port resume. This bit is not-EHCI compatible.</p> <p>1 Resume detected/driven on port. 0 No resume (K-state) detected/driven on port.</p> <p>Host mode:</p> <ul style="list-style-type: none"> • Software sets this bit to one to drive resume signaling. The controller sets this bit to one if a J-to-K transition is detected while the port is in the Suspend state. When this bit transitions to a one a J-to-K transition is detected, USBSTS[PCI] (port change detect) is also set. This bit will automatically change to zero after the resume sequence is complete. This behavior is different from EHCI where the host controller driver is required to set this bit to a zero after the resume duration is timed in the driver. • Note that when the controller owns the port, the resume sequence follows the defined sequence documented in the USB Specification Revision 2.0. The resume signaling (Full-speed 'K') is driven on the port as long as this bit remains a one. This bit will remain a one until the port has switched to the high-speed idle. Writing a zero has no affect because the port controller will time the resume operation clear the bit the port control state switches to HS or FS idle. • This field is zero if Port Power (PP) is zero in host mode. <p>Device mode:</p> <ul style="list-style-type: none"> • After the device has been in Suspend State for 5 msec or more, software must set this bit to one to drive resume signaling before clearing. The USB DR controller will set this bit to one if a J-to-K transition is detected while the port is in the Suspend state. The bit is cleared when the device returns to normal operation. Also, when this bit transitions to a one because a J-to-K transition detected, USBSTS[PCI] is also set. |
| 5 | OCC | <p>Over-current change. This bit gets set when there is a change to over-current active. Software clears this bit by writing a one to this bit position.</p> <p>Host/OTG mode:</p> <ul style="list-style-type: none"> • The user can provide over-current detection to the USB_n_PWRFAULT signal for this condition. <p>Device mode:</p> <ul style="list-style-type: none"> • This bit must always be 0. <p>1 Over current detect. 0 No over current.</p> |
| 4 | OCA | <p>Over-current active. This bit will automatically transition from one to zero when the over current condition is removed.</p> <p>Host/OTG mode:</p> <ul style="list-style-type: none"> • The user can provide over-current detection to the USB_n_PWRFAULT signal for this condition. <p>Device mode:</p> <ul style="list-style-type: none"> • This bit must always be 0. <p>1 Port currently in over-current condition. 0 Port not in over-current condition.</p> |
| 3 | PEC | <p>Port enable/disable change.</p> <p>For the root hub, this bit gets set only when a port is disabled due to disconnect on the port or due to the appropriate conditions existing at the EOF2 point (See Chapter 11 of the USB Specification). Software clears this by writing a one to it.[</p> <p>In device mode:</p> <ul style="list-style-type: none"> • The device port is always enabled. (This bit is zero.) <p>1 Port disabled. 0 No change.</p> <p>This field is zero if Port Power(PP) is zero.</p> |

Table 17-24. PORTSC Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 2 | PE | Port enabled/disabled. Host mode: <ul style="list-style-type: none"> Ports can only be enabled by the controller as a part of the reset and enable. Software cannot enable a port by writing a one to this field. Ports can be disabled by either a fault condition (disconnect event or other fault condition) or by the host software. Note that the bit status does not change until the port state actually changes. There may be a delay in disabling or enabling a port due to other host and bus events. When the port is disabled, (0) downstream propagation of data is blocked except for reset. This field is zero if Port Power(PP) is zero in host mode. Device mode: <ul style="list-style-type: none"> The device port is always enabled. (This bit is one.) |
| 1 | CSC | Connect change status. Host mode: <ul style="list-style-type: none"> This bit indicates a change has occurred in the port's Current Connect Status. the controller sets this bit for all changes to the port device connect status, even if system software has not cleared an existing connect status change. For example, the insertion status changes twice before system software has cleared the changed condition, hub hardware is 'setting' an already-set bit (that is, the bit will remain set). Software clears this bit by writing a one to it. 1 Connect Status has changed. 0 No change. <ul style="list-style-type: none"> This field is zero if Port Power(PP) is zero. Device mode: <ul style="list-style-type: none"> This bit is undefined. |
| 0 | CCS | Current connect status. Host mode: 1 Device is present 0 No device present. This field is zero if Port Power(PP) is zero in host mode. In device mode: 1 Attached 0 Not attached. A one indicates that the device successfully attached and is operating in either high-speed or full-speed as indicated by the High Speed Port bit in this register. A zero indicates that the device did not attach successfully or was forcibly disconnected by the software writing a zero to USB_CMD[RS] (run bit). It does not state the device being disconnected or suspended. |

17.3.2.15 On-The-Go Status and Control (OTGSC)—Non-EHCI

This register is not defined in the EHCI specification. The USB DR module implements one On-The-Go (OTG) status and control register corresponding to Port 0.

The OTGSC register has four sections:

- OTG interrupt enables (Read/Write)
- OTG interrupt status (Read/Write to Clear)
- OTG status inputs (Read Only)
- OTG controls (Read/Write)

The status inputs are de-bounced using a 1-msec time constant. Values on the status inputs that do not persist for more than 1 msec will not cause an update of the status inputs, or cause an OTG interrupt.

Offset 0x1A4

Access: Mixed

| | | | | | | | | | | | | | | | | |
|-------|----|------|------|-------|-------|-------|-------|------|----|------|------|-------|-------|-------|-------|------|
| | 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 |
| R | — | DPIE | 1msE | BSEIE | BSVIE | ASVIE | AVVIE | IDIE | — | DPIS | 1msS | BSEIS | BSVIS | ASVIS | AVVIS | IDIS |
| W | | | | | | | | | | w1c | w1c | w1c | w1c | w1c | w1c | w1c |
| Reset | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

| | | | | | | | | | | | | | | | | |
|-------|----|-----|------|-----|-----|-----|-----|----|---|---|---|----|----|---|----|----|
| | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| R | — | DPS | 1msT | BSE | BSV | ASV | AVV | ID | — | | | DP | OT | — | VC | VD |
| W | | | | | | | | | | | | | | | | |
| Reset | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 |

Figure 17-21. OTG Status Control (OTGSC)

Table 17-25. OTGSC Register Field Descriptions

| Bits | Name | Description |
|------|-------|---|
| 31 | — | Reserved, should be cleared. |
| 30 | DPIE | Data pulse interrupt enable 1 Enable 0 Disable |
| 29 | 1msE | 1-millisecond timer Interrupt enable 1 Enable 0 Disable |
| 28 | BSEIE | B session end interrupt enable 1 Enable 0 Disable |
| 27 | BSVIE | B session valid interrupt enable 1 Enable 0 Disable |
| 26 | ASVIE | A session valid interrupt enable 1 Enable 0 Disable |
| 25 | AVVIE | A VBus valid interrupt enable 1 Enable 0 Disable |
| 24 | IDIE | USB ID interrupt enable. 1 Enable 0 Disable |
| 23 | — | Reserved, should be cleared. |
| 22 | DPIS | Data pulse interrupt status. Set when data bus pulsing occurs on DP or DM. Data bus pulsing is only detected when USBMODE[CM] = Host (11) and PORTSC[PP] (port power) = Off (0). Software must write a one to clear this bit. |
| 21 | 1msS | 1-millisecond timer interrupt status. Set once every millisecond. Software must write a one to clear this bit. |
| 20 | BSEIS | B session end interrupt status. Set when VBus has fallen below the B session end threshold. Software must write a one to clear this bit. |

Table 17-25. OTGSC Register Field Descriptions (continued)

| Bits | Name | Description |
|------|-------|--|
| 19 | BSVIS | B session valid interrupt status. Set when VBus has either risen above or fallen below the B session valid threshold (0.8 VDC). Software must write a one to clear this bit. |
| 18 | ASVIS | A session valid interrupt status. Set when VBus has either risen above or fallen below the A session valid threshold (0.8 VDC). Software must write a one to clear this bit. |
| 17 | AVVIS | A VBus valid interrupt status. Set when VBus has either risen above or fallen below the VBus valid threshold (4.4 VDC) on an A device. Software must write a one to clear this bit. |
| 16 | IDIS | USB ID interrupt status. Set when a change on the ID input has been detected. Software must write a one to clear this bit. |
| 15 | — | Reserved, should be cleared. |
| 14 | DPS | Data bus pulsing status 1 Pulsing detected on port 0 No pulsing on port |
| 13 | 1msT | 1 millisecond timer toggle. This bit toggles once per millisecond. |
| 12 | BSE | B session end 1 VBus is below the B session end threshold. 0 VBus is above the B session end threshold. |
| 11 | BSV | B session valid 1 VBus is above the B session valid threshold. 0 VBus is below the B session valid threshold. |
| 10 | ASV | A session valid 1 VBus is above the A session valid threshold. 0 VBus is below the A session valid threshold. |
| 9 | AVV | A VBus valid 1 VBus is above the A VBus valid threshold. 0 VBus is below the A VBus valid threshold. |
| 8 | ID | USB ID 1 B device 0 A device |
| 7–5 | — | Reserved, should be cleared. |
| 4 | DP | Data pulsing 1 The pullup on DP is asserted for data pulsing during SRP. 0 The pullup on DP is not asserted. |
| 3 | OT | OTG termination. This bit must be set when the OTG device is in device mode. 1 Enable pulldown on DM 0 Disable pulldown on DM |
| 2 | — | Reserved, should be cleared. |
| 1 | VC | VBUS charge. Setting this bit causes the VBus line to be charged. This is used for VBus pulsing during SRP. |
| 0 | VD | VBUS discharge. Setting this bit causes VBus to discharge through a resistor. |

17.3.2.16 USB Mode Register (USBMODE)—Non-EHCI

The USB mode register, shown in [Figure 17-22](#), is not defined in the EHCI specification. This register controls the operating mode of the module.




Figure 17-22. USB Mode (USBMODE)

[Table 17-26](#) describes the USB mode register fields.

Table 17-26. USBMODE Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 31–5 | — | Reserved, should be cleared. |
| 4 | SDIS | Stream disable In host mode, setting this bit ensures that overruns/underruns of the latency FIFO are eliminated for low bandwidth systems where the RX and TX buffers are sufficient to contain the entire packet. Enabling stream disable also has the effect of ensuring the TX latency is filled to capacity before the packet is launched onto the USB. Note that time duration to pre-fill the FIFO becomes significant when stream disable is active. See TXFILLTUNING to characterize the adjustments needed for the scheduler when using this feature. Also note that in systems with high system bus utilization, setting this bit will ensure no overruns or underruns during operation, at the expense of link utilization. For those who desire optimal link performance, SDIS can be left clear, and the rules used under the description of the TXFILLTUNING register to limit underruns/overruns. 1 Active. 0 Inactive. In device mode, setting this bit disables double priming on both RX and TX for low bandwidth systems. This mode ensures that when the RX and TX buffers are sufficient to contain an entire packet that the standard double buffering scheme is disabled to prevent overruns/underruns in bandwidth limited systems. Note that in high-speed mode, all packets received are responded to with a NYET handshake when stream disable is active. |
| 3 | SLOM | Setup lockout mode. In device mode, this bit controls behavior of the setup lock mechanism. See Section 17.8.3.5, “Control Endpoint Operation Model.” 1 Setup lockouts off. DCD requires use of setup data buffer tripwire in USBCMD (SUTW). 0 Setup lockouts on |
| 2 | — | Reserved, should be cleared. |
| 1–0 | CM | Controller mode This register can only be written once after reset. If it is necessary to switch modes, software must reset the controller by writing to USBCMD[RST] before reprogramming this register. 00 Idle (default for combination host/device). 01 Reserved, should be cleared. 10 Device controller (default for device only controller). 11 Host controller (default for host only controller). Defaults to the idle state and needs to be initialized to the desired operating mode after reset. |

17.3.2.17 Endpoint Setup Status Register (ENDPTSETUPSTAT)—Non-EHCI

The endpoint setup status register, shown in [Figure 17-23](#), is not defined in the EHCI specification. This register contains the endpoint setup status. It is only used in device mode.




Figure 17-23. Endpoint Setup Status (ENDPTSETUPSTAT)

[Table 17-27](#) describes the endpoint setup status register fields.

Table 17-27. ENDPTSETUPSTAT Register Field Descriptions

| Bits | Name | Description |
|------|--------------------|---|
| 31–3 | — | Reserved, should be cleared. |
| 2–0 | ENDPTSETUP STAT | Setup endpoint status. For every setup transaction that is received, a corresponding bit in this register is set. Software must clear or acknowledge the setup transfer by writing a one to a respective bit after it has read the setup data from queue head. The response to a setup packet as in the order of operations and total response time is crucial to limit bus time outs while the setup lockout mechanism is engaged. This register is only used in device mode. |

17.3.2.18 Endpoint Initialization Register (ENDPTPRIME)—Non-EHCI

The endpoint initialization register, shown in [Figure 17-24](#), is not defined in the EHCI specification. This register is used to initialize endpoints. It is only used in device mode.




Figure 17-24. Endpoint Initialization (ENDPTPRIME)

Table 17-28 describes the endpoint initialization register fields.

Table 17-28. ENDPTPRIME Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–19 | — | Reserved, should be cleared. |
| 18–16 | PETB | Prime endpoint transmit buffer. For each endpoint a corresponding bit is used to request that a buffer prepared for a transmit operation in order to respond to a USB IN/INTERRUPT transaction. Software should write a one to the corresponding bit when posting a new transfer descriptor to an endpoint. Hardware will automatically use this bit to begin parsing for a new transfer descriptor from the queue head and prepare a transmit buffer. Hardware will clear this bit when the associated endpoint(s) is (are) successfully primed. PETB[2] (bit 18 of the register) corresponds to endpoint 2. Note that these bits are momentarily set by hardware during hardware re-priming operations when a dTD is retired, and the dQH is updated. |
| 15–3 | — | Reserved, should be cleared. |
| 2–0 | PERB | Prime endpoint receive buffer. For each endpoint, a corresponding bit is used to request a buffer prepare for a receive operation in order to respond to a USB OUT transaction. Software should write a one to the corresponding bit whenever posting a new transfer descriptor to an endpoint. Hardware will automatically use this bit to begin parsing for a new transfer descriptor from the queue head and prepare a receive buffer. Hardware will clear this bit when the associated endpoint(s) is (are) successfully primed. PERB[2] corresponds to endpoint 2. Note that these bits are momentarily set by hardware during hardware re-priming operations when a dTD is retired, and the dQH is updated. |

17.3.2.19 Endpoint Flush Register (ENDPTFLUSH)—Non-EHCI

The endpoint flush register, shown in Figure 17-25, is not defined in the EHCI specification. This register is only used in device mode.




Figure 17-25. Endpoint Flush (ENDPTFLUSH)

Table 17-29 describes the endpoint flush register fields.

Table 17-29. ENDPTFLUSH Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 31–19 | — | Reserved, should be cleared. |
| 18–16 | FETB | Flush endpoint transmit buffer. Writing a one to a bit(s) in this register will cause the associated endpoint(s) to clear any primed buffers. If a packet is in progress for one of the associated endpoints, then that transfer will continue until completion. Hardware will clear this register after the endpoint flush operation is successful. FETB[2] (bit 18 of the register) corresponds to endpoint 2. |

Table 17-29. ENDPTFLUSH Register Field Descriptions (continued)


| Bits | Name | Description |
|------|------|---|
| 15–3 | — | Reserved, should be cleared. |
| 2–0 | FERB | Flush endpoint receive buffer. Writing a one to a bit(s) will cause the associated endpoint(s) to clear any primed buffers. If a packet is in progress for one of the associated endpoints, then that transfer will continue until completion. Hardware will clear this register after the endpoint flush operation is successful. FERB[2] corresponds to endpoint 2. |

17.3.2.20 Endpoint Status Register (ENDPTSTATUS)—Non-EHCI

The endpoint status register, shown in [Figure 17-26](#), is not defined in the EHCI specification. This register is only used in device mode.

Offset 0x1B8

Access: Read only


Figure 17-26. Endpoint Status (ENDPTSTATUS)

[Table 17-30](#) describes the endpoint status fields.

Table 17-30. ENDPTSTATUS Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–19 | — | Reserved, should be cleared |
| 18–16 | ETBR | Endpoint transmit buffer ready. One bit for each endpoint indicates status of the respective endpoint buffer. This bit is set by the hardware as a response to receiving a command from a corresponding bit in the ENDPTPRIME register. There will always be a delay between setting a bit in the ENDPTPRIME register and endpoint indicating ready. This delay time varies based upon the current USB traffic and the number of bits set in the ENDPTPRIME register. Buffer ready is cleared by USB reset, by the USB DMA system, or through the ENDPTFLUSH register. ETBR[2] (bit 18 of the register) corresponds to endpoint 2. Note that these bits are momentarily cleared by hardware during hardware endpoint re-priming operations when a dTD is retired, and the dQH is updated. |
| 15–3 | — | Reserved, should be cleared |
| 2–0 | ERBR | Endpoint receive buffer ready. One bit for each endpoint indicates status of the respective endpoint buffer. This bit is set by the hardware as a response to receiving a command from a corresponding bit in the ENDPTPRIME register. There will always be a delay between setting a bit in the ENDPTPRIME register and endpoint indicating ready. This delay time varies based upon the current USB traffic and the number of bits set in the ENDPTPRIME register. Buffer ready is cleared by USB reset, by the USB DMA system, or through the ENDPTFLUSH register. ERBR[2] corresponds to endpoint 2. Note that these bits are momentarily cleared by hardware during hardware endpoint re-priming operations when a dTD is retired, and the dQH is updated. |

17.3.2.21 Endpoint Complete Register (ENDPTCOMPLETE)—Non-EHCI


The endpoint complete register, shown in [Figure 17-27](#), is not defined in the EHCI specification. This register is only used in device mode.

Table 17-32. ENDPTCTRL0 Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 17 | — | Reserved, should be cleared. |
| 16 | TXS | TX endpoint stall. Software can write a one to this bit to force the endpoint to return a STALL handshake to the Host. It will continue returning STALL until the bit is cleared by software or it will automatically be cleared upon receipt of a new SETUP request. 1 Endpoint stalled 0 Endpoint OK |
| 15–8 | — | Reserved, should be cleared. |
| 7 | RXE | RX endpoint enable. Endpoint zero is always enabled. 0 Disabled 1 Enabled |
| 6–4 | — | Reserved, should be cleared. |
| 3–2 | RXT | RX endpoint type. Endpoint zero is always a control endpoint (00). |
| 1 | — | Reserved, should be cleared. |
| 0 | RXS | RX endpoint stall Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue returning STALL until the bit is cleared by software or it will automatically be cleared upon receipt of a new SETUP request. 1 Endpoint stalled 0 Endpoint OK |

17.3.2.23 Endpoint Control Register n (ENDPTCTRL n)—Non-EHCI

The endpoint control n registers, shown in [Figure 17-29](#), are not defined in the EHCI specification. There is an ENDPTCTRL n register of each endpoint in a device.


Figure 17-29. Endpoint Control 1 to 5 (ENDPTCTRL n)

[Table 17-33](#) describes the endpoint control n register fields.

Table 17-33. ENDPTCTRL n Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 31–24 | — | Reserved, should be cleared |
| 23 | TXE | TX endpoint enable 0 Disabled 1 Enabled |
| 22 | TXR | TX data toggle reset. Whenever a configuration event is received for this endpoint, software must write a one to this bit in order to synchronize the data PID's between the Host and device. |

Table 17-33. ENDPCTRL_n Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 21 | TXI | TX data toggle inhibit. Used only for test and should always be written as zero. Writing a one to this bit will cause this endpoint to ignore the data toggle sequence and always transmit DATA0 for a data packet. 0 PID sequencing enabled 1 PID sequencing disabled |
| 20 | — | Reserved, should be cleared |
| 19–18 | TXT | TX endpoint type 00 Control 01 Isochronous 10 Bulk 11 Interrupt Note: When only one endpoint (RX or TX, but not both) of an endpoint pair is used, the unused endpoint should be configured as a bulk type endpoint. |
| 17 | TXD | TX endpoint data source. This bit should always be written as 0, which selects the dual port memory/DMA engine as the source. |
| 16 | TXS | TX endpoint stall. This bit is set automatically upon receipt of a SETUP request if this endpoint is not configured as a control endpoint. It is cleared automatically upon receipt of a SETUP request if this endpoint is configured as a control endpoint. Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue to returning STALL until this bit is either cleared by software or automatically cleared as above. 0 Endpoint OK 1 Endpoint stalled |
| 15–8 | — | Reserved, should be cleared |
| 7 | RXE | RX endpoint enable 0 Disabled 1 Enabled |
| 6 | RXR | RX data toggle reset. Whenever a configuration event is received for this endpoint, software must write a one to this bit in order to synchronize the data PID's between the Host and device. |
| 5 | RXI | RX data toggle inhibit. This bit is only used for test and should always be written as zero. Writing a one to this bit will cause this endpoint to ignore the data toggle sequence and always accept data packets regardless of their data PID. 1 PID sequencing enabled 0 PID sequencing disabled |
| 4 | — | Reserved, should be cleared |
| 3–2 | RXT | RX endpoint type. 00 Control 01 Isochronous 10 Bulk 11 Interrupt Note: When only one endpoint (RX or TX, but not both) of an endpoint pair is used, the unused endpoint should be configured as a bulk type endpoint. |

Table 17-33. ENDPTCTRL n Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 1 | RXD | RX endpoint data sink. This bit should always be written as 0, which selects the dual port memory/DMA engine as the sink. |
| 0 | RXS | RX endpoint stall. This bit is set automatically upon receipt of a SETUP request if this endpoint is not configured as a control endpoint. It is cleared automatically upon receipt a SETUP request if this endpoint is configured as a control endpoint, Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue to returning STALL until this bit is either cleared by software or automatically cleared as above, 1 Endpoint stalled 0 Endpoint OK |

17.3.2.24 SNOOP1 and SNOOP2—Non-EHCI

Figure 17-30 shows the SNOOP1 and SNOOP2 registers. Note that these registers use big-endian byte ordering and are not defined in the EHCI specification. The SNOOP1 and SNOOP2 registers provide snooping control and address range selection function. Transactions that hit a snooping window will generate cache coherent transactions on the internal CSB bus. When the five lower bits (SNOOP n [27–31]) are equal to 00000, snooping is always disabled on the CSB for all DMA transfers. When SNOOP n [27–31] is 01011 through 11110, the twenty upper bits (SNOOP n [0–19]) provide the starting base address for which transactions are snooped. These twenty bits are compared to the twenty upper bits of the address provided by the DMA block of the USB controller. When a match occurs, the five lower bits are decoded as shown below. This provides a snooping region of 4 Kbytes to 2 Gbytes within each starting base address that is programmed by the core. The SNOOP n [20–26] are not used.


Figure 17-30. Snoop 1 and Snoop 2 (SNOOP n)

Table 17-34 describes the SNOOP n register fields.

Table 17-34. SNOOP n Register Field Descriptions

| Bits | Name | Description |
|------|---------------|---|
| 0–19 | Snoop address | The starting base address for which transactions are snooped. |

Table 17-34. SNOOP_n Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|---------------|--|
| 20–26 | — | Reserved, should be cleared |
| 27–31 | Snoop Enables | 0x00 Snooping disabled 0x0B 4-Kbyte snoop range starting at the value defined by SNOOP _n [0–19] 0x0C 8-Kbyte snoop range starting at the value defined by SNOOP _n [0–18] 0x0D 16-Kbyte snoop range starting at the value defined by SNOOP _n [0–17] 0x0E 32-Kbyte snoop range starting at the value defined by SNOOP _n [0–16] 0x0F 64-Kbyte snoop range starting at the value defined by SNOOP _n [0–15] 0x10 128-Kbyte snoop range starting at the value defined by SNOOP _n [0–14] 0x11 256-Kbyte snoop range starting at the value defined by SNOOP _n [0–13] 0x12 512-Kbyte snoop range starting at the value defined by SNOOP _n [0–12] 0x13 1-Mbyte snoop range starting at the value defined by SNOOP _n [0–11] 0x14 2-Mbyte snoop range starting at the value defined by SNOOP _n [0–10] 0x15 4-Mbyte snoop range starting at the value defined by SNOOP _n [0–9] 0x16 8-Mbyte snoop range starting at the value defined by SNOOP _n [0–8] 0x17 16-Mbyte snoop range starting at the value defined by SNOOP _n [0–7] 0x18 32-Mbyte snoop range starting at the value defined by SNOOP _n [0–6] 0x19 64-M byte snoop range starting at the value defined by SNOOP _n [0–5] 0x1A 31-Mbyte snoop range starting at the value defined by SNOOP _n [0–4] 0x1B 256-Mbyte snoop range starting at the value defined by SNOOP _n [0–3] 0x1C 512-Mbyte snoop range starting at the value defined by SNOOP _n [0–2] 0x1D 1-Gbyte snoop range starting at the value defined by SNOOP _n [0–1] 0x1E 2-Gbyte snoop range starting at the value defined by SNOOP _n [0] |

17.3.2.25 Age Count Threshold Register (AGE_CNT_THRESH)—Non-EHCI

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The age count threshold (AGE_CNT_THRESH) register provides the aging counter threshold value used to determine the priority state of the USB DR controller’s internal system interface. This is used to increase the priority state of the module’s system interface from zero to one. The actual priority level on the system bus for each state is defined by the PRI_CTRL register. See [Section 6.3.1.1, “Address Bus Arbitration with PRIORITY\[0:1\],”](#) for more details on bus priority. The threshold value is in units of *csb_clk* cycles. This register should be written during system initialization or during normal system operation when the system bus interface is idle. It can be read at any time.

If the aging counter is less than the AGE_CNT_THRESH value, priority state zero is chosen. If the aging counter is greater than or equal to the AGE_CNT_THRESH value, priority state one is chosen.

The aging counter begins to count from zero when a bus access is requested. It increments every bus cycle until the bus transaction completes. At the completion of a bus transaction, the counter is synchronously reset to zero. If there are any outstanding bus requests, the aging counter will then begin counting immediately.

The AGE_CNT_THRESH is compared against the value of the aging counter during each clock cycle of the current transaction. If AGE_CNT_THRESH is equal to zero, priority state one is always chosen. If the aging counter is less than the AGE_CNT_THRESH value, priority state zero is selected. If the aging counter is greater than or equal to the AGE_CNT_THRESH value, priority state one is selected.

The two priority states of the aging counter function each have corresponding register bits which are programmed by the CPU. Thus, when the aging counter function is at priority state zero, PRI_CTRL[30–31] are selected and used to drive bus priority levels. When the aging counter function is at priority state one, PRI_CTRL[28–29] are selected and used to drive the priority.

Figure 17-31 shows the age count threshold register.



Figure 17-31. Age Count Threshold (AGE_CNT_THRESH)

Table 17-35 describes the age count threshold register fields.

Table 17-35. AGE_CNT_THRESH Register Field Descriptions

| Bits | Name | Description |
|-------|-----------|--------------------------------|
| 0–17 | — | Reserved, should be cleared |
| 18–31 | Threshold | Aging counter threshold value. |

The setting of AGE_CNT_THRESH is highly dependent on both the mix of other controllers operating on the system bus as well as the kind of traffic moving through the USB controller. A recommended approach is first to try leaving the aging mechanism disabled and see if the USB meets performance requirements. If USB performance does not meet application requirements, try the following settings:

- Set PRI_CTRL[pri_lv10] to 0.
- Set tPRI_CTRL[pri_lv11] to 3.
- Set AGE_CNT_THRESH to 40.

This combination works for a wide variety of applications. If this combination still does not meet application requirements, try lowering AGE_CNT_THRESH by 5. On the contrary, the setting 40 may be too conservative for some applications. If USB performance is acceptable at 40, try raising the value in increments of 5. Raising AGE_CNT_THRESH benefits the other controllers on the system bus by reducing the frequency that this USB controller raises its priority to the arbiter.

17.3.2.26 Priority Control Register (PRI_CTRL)—Non-EHCI

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The priority control (PRI_CTRL) register sets the priority level for each of two priority states. The priority state is determined by the value programmed in the AGE_CNT_THRESH register and the number of *csb_clk* cycles that a particular transaction takes to complete.

Figure 17-32 shows the priority control register.



Figure 17-32. Priority Control (PRI_CTRL)

Table 17-36 describes the priority control register fields.

Table 17-36. PRI_CTRL Register Field Descriptions

| Bits | Name | Description |
|-------|---------|--------------------------------------|
| 0–27 | — | Reserved, should be cleared |
| 28–29 | pri_lv1 | Priority level for priority state 1. |
| 30–31 | pri_lv0 | Priority level for priority state 0. |

17.3.2.27 System Interface Control Register (SI_CTRL)—Non-EHCI

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The system interface control register (SI_CTRL) controls various functions pertaining to the internal system interface.

Figure 17-33 shows the system interface control register.



Figure 17-33. System Interface Control Register (SI_CTRL)

Table 17-37 describes the system interface control register fields.

Table 17-37. SI_CTRL Register Field Descriptions

| Bits | Name | Description |
|------|-------------|---|
| 0–26 | — | Reserved, should be cleared |
| 27 | err_disable | When this bit is set, it causes the controller to ignore system bus errors. If cleared the controller responds according to the values set in USBSTS[SEI] and USBINT[SEE]. 0 enable 1 disable |

Table 17-37. SI_CTRL Register Field Descriptions (continued)


| Bits | Name | Description |
|-------|-----------------|--|
| 28–30 | — | Reserved, should be cleared |
| 31 | rd_prefetch_val | Selects whether 32 bytes or 64 bytes are fetched during burst read transactions at the system interface. When this input is LOW 64 bytes are fetched and when it is HIGH 32 bytes are fetched. The setting of rd_prefetch_val must match the setting of the larger of TXPBURST and RXPBURST fields in the BURSTSIZE register. If either of these fields is 64 bytes, then rd_prefetch_val must be left cleared. Otherwise, this value should be set. 0 64-byte fetch 1 32-byte fetch |

17.3.2.28 USB General Purpose Register (CONTROL)—Non-EHCI

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The USB general purpose (CONTROL) register contains the general-purpose IP control register outputs and is shown in [Figure 17-34](#).

Offset 0x500

Access: Mixed


Figure 17-34. USB General-Purpose Register (CONTROL)

[Table 17-38](#) describes the USB general-purpose register fields.

Table 17-38. CONTROL Register Field Descriptions

| Bits | Name | Description |
|------|---------------|--|
| 0–13 | — | Reserved, must be cleared. |
| 14 | PHY_CLK_VALID | Indicates whether the PHY clock is valid (read only). When in UTMI mode, this bit reflects the value of the UTMI PHY ClkValid signal. In ULPI mode, this bit reflects the inverted ULPI DIR. In ULPI mode, this bit is not valid if the USB I/O have not been configured and after the USB_EN signal is asserted. 0 USB PHY clock is not valid 1 USB PHY clock is valid |
| 15 | WU_INT | Reflects the state of the wake up interrupt. The wake up interrupt signal is asserted when a wake-up event occurs while in a low-power suspend state. If WU_INT_EN is set, this WU_INT signal generates an interrupt to the system to indicate wake up servicing is required. WU_INT will remain set until the USB controller is exited from the low power by clearing the PORTSC[PHCD] bit. 0 Normal operation or low-power mode waiting for wakeup event 1 Low power wakeup event has occurred |

Table 17-38. CONTROL Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|----------------|---|
| 16–18 | — | Reserved, must be cleared. |
| 19–20 | CLKIN_SEL[1:0] | Select the clock source for the UTMI PHY PLL reference clock. The reference clock can be sourced from the USBDR_CLK or SYS_CLK. These bits are not relevant when in ULPI mode. 00 Reference clock is USBDR_CLK 01 Reference clock is USBDR_CLK 10 Reference clock is SYS_CLK 11 Reserved |
| 21 | PHY_CLK_SEL | Select the source of the USB link controller transceiver clock. When cleared the UTMI PHY is the source of the clock. When set, the clock is sourced from the external ULPI PHY. 0 UTMI is clock source 1 ULPI is clock source |
| 22 | UTMI_PHY_EN | Enable the UTMI PHY. The UTMI PHY is reset when placed in the disable mode. 0 UTMI PHY disabled 1 UTMI PHY enabled |
| 23 | PLL_RESET | Reset the UTMI PHY PLL. This bit is not self clearing and must be cleared to complete the reset sequence. 0 UTMI PHY in normal operating state 1 Put UTMI PHY in reset state |
| 24–25 | REFSEL[1:0] | The REFSEL signals are used to set the frequency value for the UTMI PLL reference clock. These bit fields are not relevant if not in UTMI mode. 00 Reserved 01 Reference clock frequency is 24 MHz. This is the default frequency. 10 Reference clock frequency is 48 MHz 11 Reserved |
| 26 | OTG_PORT | Enables the OTG comparators. 0 OTG comparators disabled 1 OTG comparators enabled |
| 27 | KEEP_OTG_ON | Keeps the OTG comparators on during low power suspend. 0 OTG comparators disabled during suspend 1 OTG comparators enabled during suspend |
| 28 | LSF_EN | This bit is used to enable the UTMI line state filter. When enabled the UTMI linestate[1:0] output of the UTMI PHY are filtered to account for any skew between the USB differential data lines (DP/DM). 0 Line state filter disabled 1 Line state filter enabled |
| 29 | USB_EN | UTMI mode: This bit is used to enable the USB interface. It must be set before setting RS bit in USB CMD register. 1 Enable 0 Disable ULPI mode: In safe mode, all USB interface signals are put into input mode or driven inactive, except for SUSPEND_STP, which is driven high. Also, the input signal DIR is forced to appear high to the controller. This prevents any start-up problems that otherwise could occur if the PHY and the controller take significantly different times to complete power-on reset. 1 Normal operation 0 Safe mode |

Table 17-38. CONTROL Register Field Descriptions (continued)

| Bits | Name | Description |
|------|-------------|---|
| 30 | WU_INT_EN | This bit is used to mask/unmask the system wakeup interrupt signal 0 System wakeup interrupt disabled 1 System wakeup interrupt enabled Note: PORTSC[PHCD] bit must be set for the system wakeup interrupt generation. |
| 31 | ULPI_INT_EN | Used to enable the ULPI low power wakeup interrupt from the PHY when the PHY is in low-power mode only. 0 ULPI low power wakeup interrupt disabled 1 ULPI low power wakeup interrupt enabled Note: PORTSC[PHCD] bit must be set |

17.4 Functional Description

The USB DR module can be broken down into functional sub-blocks, which are described below.

17.4.1 System Interface

The system interface block contains all the control and status registers that allow a processor to interface to the USB DR module. These registers allow the processor to control the configuration of the module, ascertain the capabilities of the module, and control the module's operation. It also has registers to control snoopability and priority of the DMA interface.

17.4.2 DMA Engine

The module contains a local DMA engine. The DMA engine interfaces internally to the CSB. It is responsible for moving all of the data to be transferred over the USB between the module and buffers in system memory. Like the system interface block, the DMA engine block uses a simple synchronous bus signaling protocol that eases connections to a number of different standard buses.

The DMA controller must access both control information and packet data from system memory. The control information is contained in link list-based queue structures. The DMA controller has state machines that are able to parse data structures defined in the EHCI specification. In host mode, the data structures are EHCI compliant and represent queues of transfers to be performed by the host controller, including the split-transaction requests that allow an EHCI controller to direct packets to FS and LS devices. In device mode, the data structures are designed to be similar to those in the EHCI specification and are used to allow device responses to be queued for each of the active pipes in the device.

17.4.3 FIFO RAM Controller

The FIFO RAM controller is used for context information and to control FIFOs between the protocol engine and the DMA controller. These FIFOs decouple the system processor/memory bus requests from the extremely tight timing required by USB.

The use of the FIFO buffers differs between host and device mode operation. In host mode, a single data channel is maintained in each direction through the buffer memory. In device mode, multiple FIFO channels are maintained for each of the active endpoints in the system.

In host mode, the USB DR module uses a 512-byte Tx buffer and a 512-byte Rx buffer. Device operation uses a single 512-byte Rx buffer and a 512-byte Tx buffer for each endpoint. The 512-byte buffers allow the module to buffer a complete HS bulk packet.

17.4.4 PHY Interface

The USB DR module interfaces to any UTMI- or ULPI-compatible PHY. The primary function of the port controller block is to isolate the rest of the module from the transceiver, and to move all of the transceiver signaling into the primary clock domain of the module. This allows the module to run synchronously with the system processor and its associated resources.

Due to pincount limitations the module only supports certain combinations of PHY interfaces and USB functionality. Refer to [Table 17-39](#) for more information.

Table 17-39. Supported PHY Interfaces

| PHY | Function |
|------|-----------------|
| UTMI | Host/Device |
| ULPI | Host/Device/OTG |

17.5 Host Data Structures

This section defines the interface data structures used to communicate control, status, and data between HCD (software) and the Enhanced Host Controller (hardware). The data structure definitions in this section support a 32-bit memory buffer address space. The interface consists of a periodic schedule, periodic frame list, asynchronous schedule, isochronous transaction descriptors, split-transaction isochronous transfer descriptors, queue heads, and queue element transfer descriptors.

The periodic frame list is the root of all periodic (isochronous and interrupt transfer type) support for the host controller interface. The asynchronous list is the root for all the bulk and control transfer type support. Isochronous data streams are managed using isochronous transaction descriptors. Isochronous split-transaction data streams are managed with split-transaction isochronous transfer descriptors. All interrupt, control, and bulk data streams are managed with queue heads and queue element transfer descriptors. These data structures are optimized to reduce the total memory footprint of the schedule and to reduce (on average) the number of memory accesses needed to execute a USB transaction.

Note that software must ensure that no interface data structure reachable by the EHCI host controller spans a 4K-page boundary.

The data structures defined in this section are (from the host controller's perspective) a mix of read-only and read/writable fields. The host controller must preserve the read-only fields on all data structure writes.

17.5.1 Periodic Frame List

Figure 17-35 shows the organization of the periodic schedule. This schedule is for all periodic transfers (isochronous and interrupt). The periodic schedule is referenced from the operational registers space using the PERIODICLISTBASE address register and the FRINDEX register. The periodic schedule is based on an array of pointers called the periodic frame list. The PERIODICLISTBASE address register is combined with the FRINDEX register to produce a memory pointer into the frame list. The periodic frame list implements a sliding window of work over time.

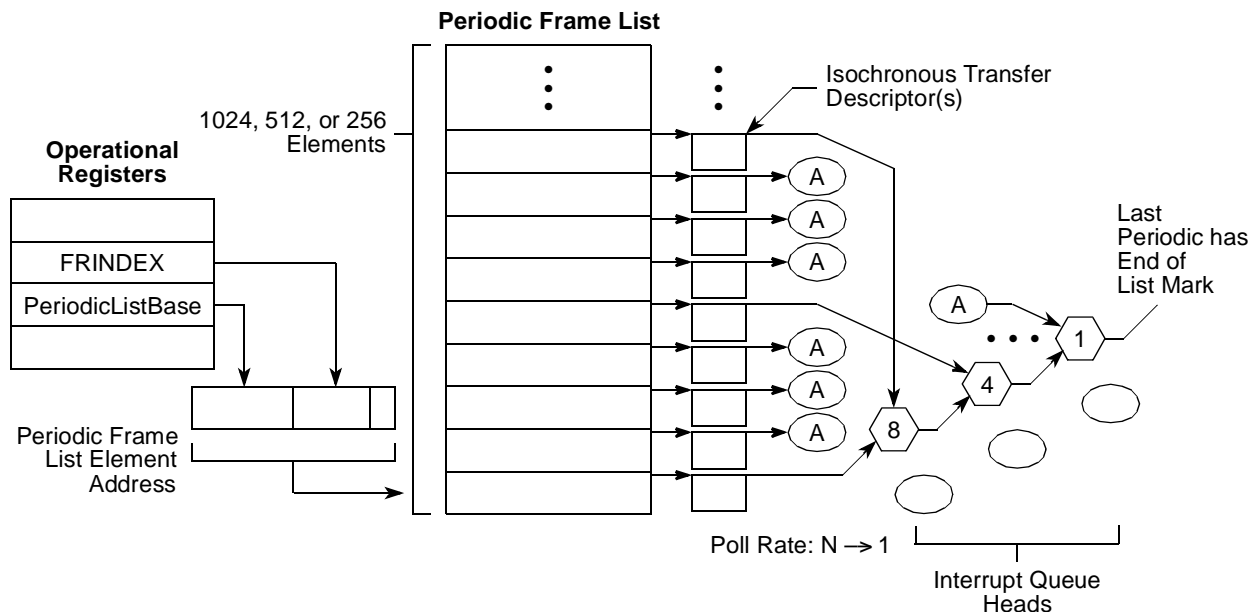


Figure 17-35. Periodic Schedule Organization

Split transaction interrupt, bulk and control are also managed using queue heads and queue element transfer descriptors.

The periodic frame list is a 4K-page aligned array of frame list link pointers. The length of the frame list may be programmable. The programmability of the periodic frame list is exported to system software through the HCCPARAMS register. If non-programmable, the length is 1024 elements. If programmable, the length can be selected by system software as one of 256, 512, or 1024 elements. An implementation must support all three sizes. Programming the size (that is, the number of elements) is accomplished by system software writing the appropriate value into frame list size field in the USBCMD register.

Frame list link pointers direct the host controller to the first work item in the frame’s periodic schedule for the current microframe. The link pointers are aligned on DWord boundaries within the frame list.

Figure 17-36 shows the format for the frame list link pointer.




Figure 17-36. Frame List Link Pointer Format

Frame list link pointers always reference memory objects that are 32-byte aligned. The referenced object may be an isochronous transfer descriptor for high-speed devices, a split-transaction isochronous transfer

descriptor (for full-speed isochronous endpoints), or a queue head (used to support high-, full- and low-speed interrupt). System software should not place non-periodic schedule items into the periodic schedule. The least-significant bits in a frame list pointer are used to key the host controller in as to the type of object the pointer is referencing.

The least-significant bit is the T bit (bit 0). When this bit is set, the host controller never uses the value of the frame list pointer as a physical memory pointer. The Typ field indicates the exact type of data structure being referenced by this pointer. The value encodings for the Typ field are given in [Table 17-40](#).

Table 17-40. Typ Field Encodings

| Typ | Description |
|-----|---|
| 00 | Isochronous transfer descriptor |
| 01 | Queue head |
| 10 | Split transaction isochronous transfer descriptor |
| 11 | Frame span traversal node |

17.5.2 Asynchronous List Queue Head Pointer

The asynchronous transfer list (based at the ASYNCLISTADDR register) is where all the control and bulk transfers are managed. Host controllers use this list only when it reaches the end of the periodic list, the periodic list is disabled, or the periodic list is empty. [Figure 17-37](#) shows the asynchronous schedule organization.

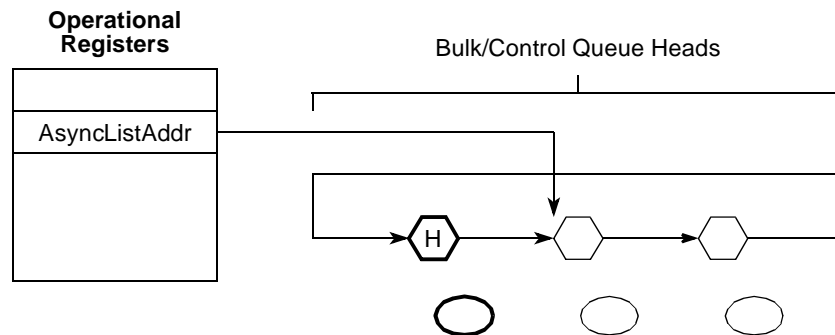


Figure 17-37. Asynchronous Schedule Organization

The asynchronous list is a simple circular list of queue heads. The ASYNCLISTADDR register is simply a pointer to the next queue head. This implements a pure round-robin service for all queue heads linked into the asynchronous list.

17.5.3 Isochronous (High-Speed) Transfer Descriptor (iTd)

[Figure 17-38](#) illustrates the format of an isochronous transfer descriptor. This structure is used only for high-speed isochronous endpoints. All other transfer types should use queue structures. Isochronous TDs must be aligned on a 32-byte boundary.

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------------------------|----|-----------------------------------|----|----|----|----|----|----|----|----|----------|---------------------|-----------------|-----------------------------------|----|----|----|----|----|----|------|------|------|------|---|---|----|-----|---|------|---|--------|
| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | offset |
| Next Link Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | 00 | Typ | T | 0x00 | | |
| Status ¹ | | Transaction 0 Length ¹ | | | | | | | | | | ioc | PG ² | Transaction 0 Offset ² | | | | | | | | | | 0x04 | | | | | | | | |
| Status ¹ | | Transaction 1 Length ¹ | | | | | | | | | | ioc | PG ² | Transaction 1 Offset ² | | | | | | | | | | 0x08 | | | | | | | | |
| Status ¹ | | Transaction 2 Length ¹ | | | | | | | | | | ioc | PG ² | Transaction 2 Offset ² | | | | | | | | | | 0x0C | | | | | | | | |
| Status ¹ | | Transaction 3 Length ¹ | | | | | | | | | | ioc | PG ² | Transaction 3 Offset ² | | | | | | | | | | 0x10 | | | | | | | | |
| Status ¹ | | Transaction 4 Length ¹ | | | | | | | | | | ioc | PG ² | Transaction 4 Offset ² | | | | | | | | | | 0x14 | | | | | | | | |
| Status ¹ | | Transaction 5 Length ¹ | | | | | | | | | | ioc | PG ² | Transaction 5 Offset ² | | | | | | | | | | 0x18 | | | | | | | | |
| Status ¹ | | Transaction 6 Length ¹ | | | | | | | | | | ioc | PG ² | Transaction 6 Offset ² | | | | | | | | | | 0x1C | | | | | | | | |
| Status ¹ | | Transaction 7 Length ¹ | | | | | | | | | | ioc | PG ² | Transaction 7 Offset ² | | | | | | | | | | 0x20 | | | | | | | | |
| Buffer Pointer (Page 0) | | | | | | | | | | | EndPt | R | Device Address | | | | | | | | | | 0x24 | | | | | | | | | |
| Buffer Pointer (Page 1) | | | | | | | | | | | I/O | Maximum Packet Size | | | | | | | | | | 0x28 | | | | | | | | | | |
| Buffer Pointer (Page 2) | | | | | | | | | | | Reserved | | | | | | | | | | Mult | 0x2C | | | | | | | | | | |
| Buffer Pointer (Page 3) | | | | | | | | | | | Reserved | | | | | | | | | | 0x30 | | | | | | | | | | | |
| Buffer Pointer (Page 4) | | | | | | | | | | | Reserved | | | | | | | | | | 0x34 | | | | | | | | | | | |
| Buffer Pointer (Page 5) | | | | | | | | | | | Reserved | | | | | | | | | | 0x38 | | | | | | | | | | | |
| Buffer Pointer (Page 6) | | | | | | | | | | | Reserved | | | | | | | | | | 0x3C | | | | | | | | | | | |

Figure 17-38. Isochronous Transaction Descriptor (iTD)

¹ Host controller read/write; all others read-only.

² These fields may be modified by the host controller if the I/O field indicates an OUT.

17.5.3.1 Next Link Pointer

The first DWord of an iTD is a pointer to the next schedule data structure, as shown in [Table 17-41](#).

Table 17-41. Next Schedule Element Pointer

| Bits | Name | Description |
|------|--------------|--|
| 31–5 | Link Pointer | Correspond to memory address signals [31:5], respectively. This field points to another isochronous transaction descriptor (iTD/siTD) or queue head (QH). |
| 4–3 | — | Reserved, should be cleared. These bits are reserved and their value has no effect on operation. Software should initialize this field to zero. |
| 2–1 | Typ | Indicates to the host controller whether the item referenced is an iTD, siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. Value encodings are: 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node) |
| 0 | T | Terminate 1 Link Pointer field is not valid. 0 Link Pointer field is valid. |

17.5.3.2 iTD Transaction Status and Control List

DWords 1–8 constitute eight slots of transaction control and status. Each transaction description includes the following:

- Status results field
- Transaction length (bytes to send for OUT transactions and bytes received for IN transactions).
- Buffer offset. The PG and Transaction n Offset fields are used with the buffer pointer list to construct the starting buffer address for the transaction.

The host controller uses the information in each transaction description, plus the endpoint information contained in the first three DWords of the buffer page pointer list, to execute a transaction on the USB.

Table 17-42 shows the iTD transaction status and control fields.

Table 17-42. iTD Transaction Status and Control

| Bits | Name | Description |
|-------|------------------------|--|
| 31–28 | Status | Records the status of the transaction executed by the host controller for this slot. This field is a bit vector with the following encoding: 31 Active. Set by software to enable the execution of an isochronous transaction by the host controller. When the transaction associated with this descriptor is completed, the host controller clears this bit indicating that a transaction for this element should not be executed when it is next encountered in the schedule. 30 Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overflow) or is unable to supply data fast enough during transmission (underflow). If an overflow condition occurs, no action is necessary. 29 Babble detected. Set by the host controller during status update when "babble" is detected during the transaction generated by this descriptor. 28 Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (Time-out, CRC, Bad PID, and so on). This bit may only be set for isochronous IN transactions. |
| 27–16 | Transaction n Length | For an OUT, this field is the number of data bytes the host controller will send during the transaction. The host controller is not required to update this field to reflect the actual number of bytes transferred during the transfer. For an IN, the initial value of the endpoint to deliver. During the status update, the host controller writes back the field is the number of bytes the host expects the number of bytes successfully received. The value in this register is the actual byte count (for example, 0 zero length data, 1 one byte, 2 two bytes, and so on). The maximum value this field may contain is 0xC00 (3072). |
| 15 | ioc | Interrupt on complete. If this bit is set, it specifies that when this transaction completes, the host controller should issue an interrupt at the next interrupt threshold. |
| 14–12 | PG | These bits are set by software to indicate which of the buffer page pointers the offset field in this slot should be concatenated to produce the starting memory address for this transaction. The valid range of values for this field is 0 to 6. |
| 11–0 | Transaction n Offset | This field is a value that is an offset, expressed in bytes, from the beginning of a buffer. This field is concatenated onto the buffer page pointer indicated in the adjacent PG field to produce the starting buffer address for this transaction. |

17.5.3.3 iTD Buffer Page Pointer List (Plus)

DWords 9–15 of an isochronous transaction descriptor are nominally page pointers (4K aligned) to the data buffer for this transfer descriptor. This data structure requires the associated data buffer to be contiguous

(relative to virtual memory), but allows the physical memory pages to be non-contiguous. Seven page pointers are provided to support the expression of eight isochronous transfers. The seven pointers allow for 3 (transactions) × 1024 (maximum packet size) × 8 (transaction records) = 24 576 bytes to be moved with this data structure, regardless of the alignment offset of the first page.

Since each pointer is a 4 K-aligned page pointer, the least-significant 12 bits in several of the page pointers are used for other purposes.

Table 17-43–Table 17-46 describe buffer pointer page *n*.

Table 17-43. Buffer Pointer Page 0 (Plus)

| Bits | Name | Description |
|-------|-------------------------|---|
| 31–12 | Buffer Pointer (Page 0) | A 4K-aligned pointer to physical memory. Corresponds to memory address bits 31–12. |
| 11–8 | EndPt | Selects the particular endpoint number on the device serving as the data source or sink. |
| 7 | — | Reserved, should be cleared. Reserved for future use and should be initialized by software to zero. |
| 6–0 | Device Address | This field selects the specific device serving as the data source or sink. |

Table 17-44. iTD Buffer Pointer Page 1 (Plus)

| Bits | Name | Description |
|-------|-------------------------|--|
| 31–12 | Buffer Pointer (Page 1) | This is a 4K aligned pointer to physical memory. Corresponds to memory address bits 31–12. |
| 11 | I/O | Direction (I/O). This field encodes whether the high-speed transaction should use an IN or OUT PID. 0 OUT 1 IN |
| 10–0 | Maximum Packet Size | This directly corresponds to the maximum packet size of the associated endpoint (<i>wMaxPacketSize</i>). This field is used for high-bandwidth endpoints where more than one transaction is issued per transaction description (for example, per microframe). This field is used with the <i>Multi</i> field to support high-bandwidth pipes. This field is also used for all IN transfers to detect packet babble. Software should not set a value larger than 1024 (0x400). Any value larger yields undefined results. |

Table 17-45. Buffer Pointer Page 2 (Plus)

| Bits | Name | Description |
|-------|-------------------------|--|
| 31–12 | Buffer Pointer (Page 2) | This is a 4K-aligned pointer to physical memory. Corresponds to memory address bits 31–12. |
| 11–2 | — | Reserved, should be cleared. This bit reserved for future use and should be cleared. |
| 1–0 | Mult | Indicates to the host controller the number of transactions that should be executed per transaction description (for example, per microframe). 00 Reserved, should be cleared. A zero in this field yields undefined results. 01 One transaction to be issued for this endpoint per microframe 10 Two transactions to be issued for this endpoint per microframe 11 Three transactions to be issued for this endpoint per microframe |

Table 17-46. Buffer Pointer Page 3–6

| Bits | Name | Description |
|-------|----------------|--|
| 31–12 | Buffer Pointer | This is a 4K aligned pointer to physical memory. Corresponds to memory address bits 31–12. |
| 11–2 | — | Reserved, should be cleared. These bits reserved for future use and should be cleared. |

17.5.4 Split Transaction Isochronous Transfer Descriptor (siTD)

All full-speed isochronous transfers through the internal transaction translator are managed using the siTD data structure. This data structure satisfies the operational requirements for managing the split transaction protocol.

Figure 17-39 shows the split-transaction isochronous transfer descriptor (siTD).

| 31 | | 30 | | 29 | | 28 | | 27 | | 26 | | 25 | | 24 | | 23 | | 22 | | 21 | | 20 | | 19 | | 18 | | 17 | | 16 | | 15 | | 14 | | 13 | | 12 | | 11 | | 10 | | 9 | | 8 | | 7 | | 6 | | 5 | | 4 | | 3 | | 2 | | 1 | | 0 | | offset |
|-------------------------|----------------|------|--|----|---|--------------------------------------|--|----|--|---------------------------------|-------|---------------|----------------|-----------------------------|--|---------------|--|-----------------|----------------------|----|--|----|--|----|--|----|--|------|------|------|------|------|------|----|------|----|--|----|--|------|--|----|--|---|--|---|--|---|--|---|--|---|--|---|--|---|--|---|--|---|--|---|--|--------|
| Next Link Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | | 00 | Typ | T | 0x00 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| I/O | Port Number | | | | 0 | Hub Address | | | | 0000 | EndPt | 0 | Device Address | | | | | | | | | | | | | | | | 0x04 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0000_0000_0000_00000 | | | | | | | | | | | | µFrame C-mask | | | | µFrame S-mask | | | | | | | | | | | | | | | | 0x08 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| ioc | P ¹ | 0000 | | | | Total Bytes to Transfer ¹ | | | | µFrame C-prog-mask ¹ | | | | Status ¹ | | | | | | | | | | | | | | | | 0x0C | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Buffer Pointer (Page 0) | | | | | | | | | | | | | | Current Offset ¹ | | | | | | | | | | | | | | | | | | | | | | | | | | 0x10 | | | | | | | | | | | | | | | | | | | | | | | | |
| Buffer Pointer (Page 1) | | | | | | | | | | | | | | 000_0000 | | | | TP ¹ | T-count ¹ | | | | | | | | | | | | | | | | 0x14 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Back Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | | 0000 | | | | T | 0x18 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

Figure 17-39. Split-Transaction Isochronous Transaction Descriptor (siTD)

¹ Host controller read/write; all others read-only.

17.5.4.1 Next Link Pointer

DWord0 of a siTD is a pointer to the next schedule data structure. Table 17-47 describes the next link pointer fields.

Table 17-47. Next Link Pointer

| Bits | Name | Description |
|------|-------------------|---|
| 31–5 | Next Link Pointer | This field contains the address of the next data object to be processed in the periodic list and corresponds to memory address signals [31:5], respectively. |
| 4–3 | — | Reserved, should be cleared. These bits must be written as zeros. |
| 2–1 | Typ | Indicates to the host controller whether the item referenced is an iTD/siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. Value encodings are: 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node) |
| 0 | T | Terminate. 0 Link pointer is valid. 1 Link pointer field is not valid. |

17.5.4.2 siTD Endpoint Capabilities/Characteristics

DWords 1 and 2 specify static information about the full-speed endpoint, the addressing of the parent Companion Controller, and microframe scheduling control.

Table 17-48 describes the endpoint and transaction translator characteristics.

Table 17-48. Endpoint and Transaction Translator Characteristics

| Bits | Name | Description |
|-------|----------------|---|
| 31 | I/O | Direction (I/O). This field encodes whether the full-speed transaction should be an IN or OUT. 0 OUT 1 IN |
| 30–24 | Port Number | This field is the port number of the recipient transaction translator. |
| 23 | — | Reserved, should be cleared. Bit reserved and should be cleared. |
| 22–16 | Hub Address | This field holds the device address of the companion controllers' hub. |
| 15–12 | — | Reserved, should be cleared. Field reserved and should be cleared. |
| 11–8 | EndPt | Endpoint Number. Selects the particular endpoint number on the device serving as the data source or sink. |
| 7 | — | Reserved, should be cleared. Bit is reserved for future use. It should be cleared. |
| 6–0 | Device Address | Selects the specific device serving as the data source or sink. |

Table 17-42 describes the microframe schedule control.

Table 17-49. Microframe Schedule Control

| Bits | Name | Description |
|-------|---------------|---|
| 31–16 | — | Reserved, should be cleared. This field reserved for future use. It should be cleared. |
| 15–8 | μFrame C-mask | Split completion mask. This field (along with the Active and SplitX- state fields in the status byte) is used to determine during which microframes the host controller should execute complete-split transactions. When the criteria for using this field is met, an all-zeros value has undefined behavior. The host controller uses the value of the three low-order bits of the FRINDEX register to index into this bit field. If the FRINDEX register value indexes to a position where the μFrame C-Mask field is a one, this siTD is a candidate for transaction execution. There may be more than one bit in this mask set. |
| 7–0 | μFrame S-mask | Split start mask. This field (along with the Active and SplitX-state fields in the Status byte) is used to determine during which microframes the host controller should execute start-split transactions. The host controller uses the value of the three low-order bits of the FRINDEX register to index into this bit field. If the FRINDEX register value indexes to a position where the μFrame S-mask field is a one, then this siTD is a candidate for transaction execution. An all zeros value in this field, in combination with existing periodic frame list has undefined results. |

17.5.4.3 siTD Transfer State

DWords 3–6 manage the state of the transfer, as described in [Table 17-50](#).

Table 17-50. siTD Transfer Status and Control

| Bits | Name | Description | |
|-------|---|--|---|
| 31 | ioc | Interrupt on complete 0 Do not interrupt when transaction is complete. 1 Do interrupt when transaction is complete. When the host controller determines that the split transaction has completed it will assert a hardware interrupt at the next interrupt threshold. | |
| 30 | P | Page select. Indicates which data page pointer should be concatenated with the CurrentOffset field to construct a data buffer pointer 0 Selects Page 0 pointer 1 Selects Page 1 pointer The host controller is not required to write this field back when the siTD is retired (Active bit transitioned from a one to a zero). | |
| 29–26 | — | Reserved, should be cleared. This field reserved for future use and should be cleared. | |
| 25–16 | Total Bytes to Transfer | This field is initialized by software to the total number of bytes expected in this transfer. Maximum value is 1023 (3FFh) | |
| 15–8 | μFrame C-prog-mask | Split complete progress mask. This field is used by the host controller to record which split-completes have been executed. | |
| 7–0 | Status | This field records the status of the transaction executed by the host controller for this slot. This field is a bit vector with the following encoding: | |
| | | Status Bits | Definition |
| | | 7 | Active. Set by software to enable the execution of an isochronous split transaction by the host controller. |
| | | 6 | ERR. Set by the host controller when an ERR response is received from the companion controller. |
| | | 5 | Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). In the case of an under run, the host controller will transmit an incorrect CRC (thus invalidating the data at the endpoint). If an overrun condition occurs, no action is necessary. |
| | | 4 | Babble detected. Set by the host controller during status update when "babble" is detected during the transaction generated by this descriptor. |
| | | 3 | Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (Time-out, CRC, Bad PID, and so on). This bit will only be set for IN transactions. |
| | | 2 | Missed microframe. The host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction. |
| | | 1 | Split transaction state (SplitXstate). The bit encodings are: 0 Do start split. This value directs the host controller to issue a Start split transaction to the endpoint when a match is encountered in the S-mask. 1 Do complete split. This value directs the host controller to issue a Complete split transaction to the endpoint when a match is encountered in the C-mask. |
| 0 | Reserved, should be cleared. Bit reserved for future use and should be cleared. | | |

17.5.4.4 siTD Buffer Pointer List (Plus)

DWords 4 and 5 are the data buffer page pointers for the transfer. This structure supports one physical page cross. The most-significant 20 bits of each DWord in this section are the 4K (page) aligned buffer pointers. The least-significant 12 bits of each DWord are used as additional transfer state.

Table 17-51 describes the siTD buffer pointer page 0.

Table 17-51. siTD Buffer Pointer Page 0 (Plus)

| Bits | Name | Description |
|-------|-------------------------|--|
| 31–12 | Buffer Pointer (Page 0) | Bits 31–12 are 4K page-aligned, physical memory addresses. These bits correspond to physical address bits 31–12 respectively. The field P specifies the current active pointer |
| 11–0 | Current Offset | The 12 least-significant bits of the Page 0 pointer is the current byte offset for the current page pointer (as selected with the page indicator bit (P field)). The host controller is not required to write this field back when the siTD is retired (Active bit transitioned from a one to a zero). |

Table 17-52 describes the siTD buffer pointer page 1.

Table 17-52. siTD Buffer Pointer Page 1 (Plus)

| Bits | Name | Description |
|-------|-------------------------|--|
| 31–12 | Buffer Pointer (Page 1) | Bits 31–12 are 4K page-aligned, physical memory addresses. These bits correspond to physical address bits 31–12 respectively. The field P specifies the current active pointer |
| 11–5 | — | Reserved, should be cleared. |
| 4–3 | TP | Transaction position. This field is used with T-count to determine whether to send all, first, middle, or last with each outbound transaction payload. System software must initialize this field with the appropriate starting value. The host controller must correctly manage this state during the lifetime of the transfer. The bit encodings are: 00 All. The entire full-speed transaction data payload is in this transaction (that is, less than or equal to 188 bytes). 01 Begin. This is the first data payload for a full-speed transaction that is greater than 188 bytes. 10 Mid. This is the middle payload for a full-speed OUT transaction that is larger than 188 bytes. 11 End. This is the last payload for a full-speed OUT transaction that was larger than 188 bytes. |
| 2–0 | T-Count | Transaction count. Software initializes this field with the number of OUT start-splits this transfer requires. Any value larger than 6 is undefined. |

17.5.4.5 siTD Back Link Pointer

DWord 6 of a siTD is simply another schedule link pointer. This pointer is always zero, or references a siTD. This pointer cannot reference any other schedule data structure.

Table 17-53 describes the siTD back link pointer.

Table 17-53. siTD Back Link Pointer

| Bits | Name | Description |
|------|--------------|--------------------------------------|
| 31–5 | Back Pointer | A physical memory pointer to an siTD |

Table 17-53. siTD Back Link Pointer (continued)

| Bits | Name | Description |
|------|------|---|
| 4–1 | — | Reserved, should be cleared. This field is reserved for future use. It should be cleared. |
| 0 | T | Terminate 0 siTD Back Pointer field is valid 1 siTD Back Pointer field is not valid |

17.5.5 Queue Element Transfer Descriptor (qTD)

This data structure is only used with a queue head. This data structure is used for one or more USB transactions. This data structure is used to transfer up to 20,480 (5×4096) bytes. The structure contains two structure pointers used for queue advancement, a DWord of transfer state, and a five-element array of data buffer pointers. This structure is 32 bytes (or one 32-byte cache line). This data structure must be physically contiguous.

The buffer associated with this transfer must be virtually contiguous. The buffer may start on any byte boundary. A separate buffer pointer list element must be used for each physical page in the buffer, regardless of whether the buffer is physically contiguous.

Host controller updates (host controller writes) to stand-alone qTDs only occur during transfer retirement. References in the following bit field definitions of updates to the qTD are to the qTD portion of a queue head.

Figure 17-40 shows the queue element transfer descriptors.

| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | offset |
|----------------------------|--------------------------------------|----|----|----|----|----|----|----|----|----|-----------------------------|---------------------|-------------------|----------|---------------------|----|------|----|----|------|----|---|---|---|---|---|------|---|------|---|---|--------|
| Next qTD Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | 0000 | T | 0x00 | | | |
| Alternate Next qTD Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | 0000 | T | 0x04 | | | |
| dt ¹ | Total Bytes to Transfer ¹ | | | | | | | | | | ioc | C_Page ¹ | Cerr ¹ | PID Code | Status ¹ | | | | | 0x08 | | | | | | | | | | | | |
| Buffer Pointer (Page 0) | | | | | | | | | | | Current Offset ¹ | | | | | | 0x0C | | | | | | | | | | | | | | | |
| Buffer Pointer (Page 1) | | | | | | | | | | | 0000_0000_0000 | | | | | | 0x10 | | | | | | | | | | | | | | | |
| Buffer Pointer (Page 2) | | | | | | | | | | | 0000_0000_0000 | | | | | | 0x14 | | | | | | | | | | | | | | | |
| Buffer Pointer (Page 3) | | | | | | | | | | | 0000_0000_0000 | | | | | | 0x18 | | | | | | | | | | | | | | | |
| Buffer Pointer (Page 4) | | | | | | | | | | | 0000_0000_0000 | | | | | | 0x1C | | | | | | | | | | | | | | | |

Figure 17-40. Queue Element Transfer Descriptor (qTD)

¹ Host controller read/write; all others read-only.

Queue element transfer descriptors must be aligned on 32-byte boundaries.

17.5.5.1 Next qTD Pointer

The first DWord of an element transfer descriptor is a pointer to another transfer element descriptor. [Table 17-54](#) describes the qTD next element transfer pointer.

Table 17-54. qTD Next Element Transfer Pointer (DWord 0)

| Bits | Name | Description |
|------|------------------|--|
| 31–5 | Next qTD Pointer | This field contains the physical memory address of the next qTD to be processed and corresponds to memory address signals [31:5], respectively. |
| 4–1 | — | Reserved, should be cleared. These bits are reserved and their value has no effect on operation. |
| 0 | T | Terminate. Indicates to the host controller that there are no more valid entries in the queue. 0 Pointer is valid (points to a valid transfer element descriptor) 1 Pointer is invalid |

17.5.5.2 Alternate Next qTD Pointer

The second DWord of a queue element transfer descriptor is used to support hardware-only advance of the data stream to the next client buffer on short packet. To be more explicit the host controller will always use this pointer when the current qTD is retired due to short packet. [Table 17-55](#) describes the alternate qTD next element transfer pointer.

Table 17-55. qTD Alternate Next Element Transfer Pointer (DWord 1)

| Bits | Name | Description |
|------|----------------------------|--|
| 31–5 | Alternate Next qTD Pointer | This field contains the physical memory address of the next qTD to be processed in the event that the current qTD execution encounters a short packet (for an IN transaction). The field corresponds to memory address signals [31:5], respectively. |
| 4–1 | — | Reserved, should be cleared. These bits are reserved and their value has no effect on operation. |
| 0 | T | Terminate. Indicates to the host controller that there are no more valid entries in the queue. 0 Pointer is valid (points to a valid transfer element descriptor) 1 Pointer is invalid |

17.5.5.3 qTD Token

The third DWord of a queue element transfer descriptor contains most of the information the host controller requires to execute a USB transaction (the remaining endpoint-addressing information is

specified in the queue head). Note that some of the field descriptions in [Table 17-56](#) reference fields are defined in the queue head. See [Section 17.5.6, “Queue Head,”](#) for more information on these fields.

Table 17-56. qTD Token (DWord 2)

| Bits | Name | Description |
|-------|-------------------------|---|
| 31 | dt | Data toggle. This is the data toggle sequence bit. The use of this bit depends on the setting of the Data Toggle Control bit in the queue head. |
| 30–16 | Total Bytes to Transfer | Total bytes to transfer. This field specifies the total number of bytes to be moved with this transfer descriptor. This field is decremented by the number of bytes actually moved during the transaction, only on the successful completion of the transaction. The maximum value software may store in this field is $5 \times 4K$ (0x5000). This is the maximum number of bytes 5 page pointers can access. If the value of this field is zero when the host controller fetches this transfer descriptor (and the active bit is set), the host controller executes a zero-length transaction and retires the transfer descriptor. It is not a requirement for OUT transfers that total bytes to transfer be an even multiple of QH[Maximum Packet Length]. If software builds such a transfer descriptor for an OUT transfer, the last transaction will always be less than QH[Maximum Packet Length]. Although it is possible to create a transfer up to 20K this assumes the page is 0. When the offset cannot be predetermined, crossing past the 5th page can be guaranteed by limiting the total bytes to 16K. Therefore, the maximum recommended transfer is 16K (0x4000). |
| 15 | ioc | Interrupt on complete. If this bit is set, the host controller should issue an interrupt at the next interrupt threshold when this qTD is completed. |
| 14–12 | C_Page | Current rage. This field is used as an index into the qTD buffer pointer list. Valid values are in the range 0x0 to 0x4. The host controller is not required to write this field back when the qTD is retired. |

Table 17-56. qTD Token (DWord 2) (continued)

| Bits | Name | Description | |
|-------|----------|---|--|
| 11–10 | Cerr | Error counter. 2-bit down counter that keeps track of the number of consecutive errors detected while executing this qTD. If this field is programmed with a non-zero value during setup, the host controller decrements the count and writes it back to the qTD if the transaction fails. If the counter counts from one to zero, the host controller marks the qTD inactive, sets the Halted bit to a one, and error status bit for the error that caused Cerr to decrement to zero. An interrupt is generated if USBINTR[UEE] is set. If the host controller driver (HCD) software programs this field to zero during setup, the host controller will not count errors for this qTD and there is no limit on the retries of this qTD. Note that write-backs of intermediate execution state are to the queue head overlay area, not the qTD. | |
| | | Error | Decrement Counter |
| | | Transaction Error | Yes |
| | | Data Buffer Error | No. Data buffer errors are host problems. They don't count against the device's retries. Note that software must not program Cerr to a value of zero when the EPS field is programmed with a value indicating a full- or low-speed device. This combination could result in undefined behavior. |
| | | Stalled | No. Detection of babble or stall automatically halts the queue head. Thus, count is not decremented |
| | | Babble Detected | No. Detection of babble or stall automatically halts the queue head. Thus, count is not decremented |
| | | No Error | No. If the EPS field indicates a HS device or the queue head is in the asynchronous schedule (and PIDCode indicates an IN or OUT) and a bus transaction completes and the host controller does not detect a transaction error, then the host controller should reset Cerr to extend the total number of errors for this transaction. For example, Cerr should be reset with maximum value (0b11) on each successful completion of a transaction. The host controller must never reset this field if the value at the start of the transaction is 0b00. |
| 9–8 | PID Code | This field is an encoding of the token, which should be used for transactions associated with this transfer descriptor. Encodings are: 00 OUT Token generates token (E1H) 01 IN Token generates token (69H) 10 SETUP Token generates token (2DH) (undefined if endpoint is an Interrupt transfer type, for example. μ Frame S-mask field in the queue head is non-zero.) 11 Reserved, should be cleared | |

Table 17-56. qTD Token (DWord 2) (continued)

| Bits | Name | Description | |
|------|--------|--|--|
| 7-0 | Status | This field is used by the host controller to communicate individual command execution states back to the host controller driver (HCD) software. This field contains the status of the last transaction performed on this qTD. The bit encodings are: | |
| | | Bits | Status Field Description |
| | | 7 | Active. Set by software to enable the execution of transactions by the host controller. |
| | | 6 | Halted. Set by the host controller during status updates to indicate that a serious error has occurred at the device/endpoint addressed by this qTD. This can be caused by babble, the error counter counting down to zero, or reception of the STALL handshake from the device during a transaction. Any time that a transaction results in the Halted bit being set, the Active bit is also cleared. |
| | | 5 | Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). If an overrun condition occurs, the host controller will force a time-out condition on the USB, invalidating the transaction at the source. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer. |
| | | 4 | Babble detected. Set by the host controller during status update when babble is detected during the transaction. In addition to setting this bit, the host controller also sets the Halted bit to a one. Since babble is considered a fatal error for the transfer, setting the Halted bit to a one insures that no more transactions occur because of this descriptor. |
| | | 3 | Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (time-out, CRC, bad PID). If the host controller sets this bit to a one, then it remains a one for the duration of the transfer. |
| | | 2 | Missed microframe. This bit is ignored unless the QH[EPS] field indicates a full- or low-speed endpoint and the queue head is in the periodic list. This bit is set when the host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer. |

Table 17-56. qTD Token (DWord 2) (continued)

| Bits | Name | Description |
|------|------|---|
| 1 | | Split transaction state (SplitXstate). This bit is ignored by the host controller unless the QH[EPS] field indicates a full- or low-speed endpoint. When a full- or low-speed device, the host controller uses this bit to track the state of the split- transaction. The functional requirements of the host controller for managing this state bit and the split transaction protocol depends on whether the endpoint is in the periodic or asynchronous schedule. The bit encodings are: 0 Do start split. This value directs the host controller to issue a start split transaction to the endpoint. 1 Do complete split. This value directs the host controller to issue a Complete split transaction to the endpoint. |
| 0 | | Ping state (P)/ERR. If the QH[EPS] field indicates a high-speed device and the PID Code indicates an OUT endpoint, then this is the state bit for the Ping protocol. The bit encodings are: 0 Do OUT. This value directs the host controller to issue an OUT PID to the endpoint. 1 Do Ping. This value directs the host controller to issue a PING PID to the endpoint. If the QH[EPS] field does not indicate a high-speed device, then this field is used as an error indicator bit. It is set by the host controller whenever a periodic split-transaction receives an ERR handshake. |

17.5.5.4 qTD Buffer Page Pointer List

The last five DWords of a queue element transfer descriptor make up an array of physical memory address pointers. These pointers reference the individual pages of a data buffer.

System software initializes the Current Offset field to the starting offset into the current page, where current page is selected with the value in the C_Page field.

Table 17-57 describes the qTD buffer pointer.

Table 17-57. qTD Buffer Pointer

| Bits | Name | Description |
|-------|--|--|
| 31–12 | Buffer Pointer (page <i>n</i>) | Each element in the list is a 4K page aligned physical memory address. The lower 12 bits in each pointer are reserved (except for the first one), as each memory pointer must reference the start of a 4K page. The field C_Page specifies the current active pointer. When the transfer element descriptor is fetched, the starting buffer address is selected using C_Page (similar to an array index to select an array element). If a transaction spans a 4K buffer boundary, the host controller must detect the page-span boundary in the data stream, increment C_Page and advance to the next buffer pointer in the list, and conclude the transaction via the new buffer pointer. |
| 11–0 | Current Offset (Page 0)/ — (Pages 1–4) | Reserved in all pointers except the first one (that is, Page 0). The host controller should ignore all reserved bits. For the page 0 current offset interpretation, this field is the byte offset into the current page (as selected by C_Page). The host controller is not required to write this field back when the qTD is retired. Software should ensure the reserved fields are initialized to zeros. |

17.5.6 Queue Head

Figure 17-41 shows the queue head structure.

| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | offset |
|---|--------------------------------------|-------------|-----------------------|----|----------|----|------------------|---------------------|-------------------|-----------------------------|---------------------|----|--------------------------|----------------|----|----|-----------------------|----|----|-------------------|----|---|---------------------|---|---|---------------------|---------------------|-----|----------------|---------------------|---|--------|
| Queue Head Horizontal Link Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | 00 | Typ | T | 0x00 | | |
| RL | | C | Maximum Packet Length | | | | H | drc | EPS | EndPt | | | I | Device Address | | | | | | | | | | | | | 0x04 ¹ | | | | | |
| Mult | | Port Number | | | Hub Addr | | | µFrame C-mask | | | | | µFrame S-mask | | | | | | | 0x08 ¹ | | | | | | | | | | | | |
| Current qTD Pointer ² | | | | | | | | | | | | | | | | | | | | | | | | | | | 00000 | | | 0x0C | | |
| Next qTD Pointer ² | | | | | | | | | | | | | | | | | | | | | | | | | | | 0000 | | T ² | 0x10 ³ | | |
| Alternate Next qTD Pointer ² | | | | | | | | | | | | | | | | | | | | | | | | | | | NakCnt ² | | T ² | 0x14 ^{3,4} | | |
| dt ¹ | Total Bytes to Transfer ² | | | | | | ioc ² | C_Page ² | Cerr ² | PID Code ² | Status ² | | | | | | | | | | | | | | | 0x18 ^{3,4} | | | | | | |
| Buffer Pointer (Page 0) ² | | | | | | | | | | Current Offset ² | | | | | | | | | | | | | | | | 0x1C ^{3,4} | | | | | | |
| Buffer Pointer (Page 1) ² | | | | | | | | | | 0000 | | | C-prog-mask ² | | | | | | | | | | | | | 0x20 ^{3,4} | | | | | | |
| Buffer Pointer (Page 2) ² | | | | | | | | | | S-bytes ² | | | | | | | FrameTag ² | | | | | | 0x24 ^{3,4} | | | | | | | | | |
| Buffer Pointer (Page 3) ² | | | | | | | | | | 0000_0000_0000 | | | | | | | | | | | | | | | | 0x28 ³ | | | | | | |
| Buffer Pointer (Page 4) ² | | | | | | | | | | 0000_0000_0000 | | | | | | | | | | | | | | | | 0x2C ³ | | | | | | |

Figure 17-41. Queue Head Layout

- ¹ Offsets 0x04 through 0x0B contain the static endpoint state.
- ² Host controller read/write; all others read-only.
- ³ Offsets 0x10 through 0x2F contain the transfer overlay.
- ⁴ Offsets 0x14 through 0x27 contain the transfer results.

17.5.6.1 Queue Head Horizontal Link Pointer

The first DWord of a queue head contains a link pointer to the next data object to be processed after any required processing in this queue has been completed, as well as the control bits defined below.

This pointer may reference a queue head or one of the isochronous transfer descriptors. It must not reference a queue element transfer descriptor.

Table 17-58 describes the queue head.

Table 17-58. Queue Head DWord 0

| Bits | Name | Description |
|------|------|--|
| 31–5 | QHLP | Queue head horizontal link pointer. This field contains the address of the next data object to be processed in the horizontal list and corresponds to memory address signals [31:5], respectively. |
| 4–3 | — | Reserved, should be cleared. These bits must be written as zeros. |

Table 17-58. Queue Head DWord 0 (continued)

| Bits | Name | Description |
|------|------|--|
| 2–1 | Typ | Indicates to the hardware whether the item referenced by the link pointer is an iTD, siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node) |
| 0 | T | Terminate. 1 Last QH (pointer is invalid). 0 Pointer is valid. If the queue head is in the context of the periodic list, a one bit in this field indicates to the host controller that this is the end of the periodic list. This bit is ignored by the host controller when the queue head is in the asynchronous schedule. Software must ensure that queue heads reachable by the host controller always have valid horizontal link pointers. |

17.5.6.2 Endpoint Capabilities/Characteristics

The second and third DWords of a queue head specify static information about the endpoint. This information does not change over the lifetime of the endpoint. There are three types of information in this region:

- Endpoint characteristics. These are the USB endpoint characteristics, which include addressing, maximum packet size, and endpoint speed.
- Endpoint capabilities. These are adjustable parameters of the endpoint. They affect how the endpoint data stream is managed by the host controller.
- Split transaction characteristics. This data structure manages full- and low-speed data streams for bulk, control, and interrupt with split transactions to USB 2.0 Hub transaction translator. Additional fields exist for addressing the hub and scheduling the protocol transactions (for periodic).

The host controller must not modify the bits in this region.

[Table 17-59](#) and [Table 17-60](#) describe the endpoint characteristics.

Table 17-59. Endpoint Characteristics: Queue Head DWord 1

| Bits | Name | Description |
|-------|-----------------------|--|
| 31–28 | RL | Nak count reload. This field contains a value, which is used by the host controller to reload the Nak Counter field. |
| 27 | C | Control endpoint flag. If the QH[EPS] field indicates the endpoint is not a high-speed device, and the endpoint is a control endpoint, then software must set this bit to a one. Otherwise, it should always set this bit to a zero. |
| 26–16 | Maximum Packet Length | This directly corresponds to the maximum packet size of the associated endpoint (wMaxPacketSize). The maximum value this field may contain is 0x400 (1024). |
| 15 | H | Head of reclamation list flag. This bit is set by system software to mark a queue head as being the head of the reclamation list. |

Table 17-59. Endpoint Characteristics: Queue Head DWord 1 (continued)

| Bits | Name | Description |
|-------|----------------|---|
| 14 | dtc | Data toggle control (DTC). Specifies where the host controller should get the initial data toggle on an overlay transition. 0 Ignore DT bit from incoming qTD. Host controller preserves DT bit in the queue head. 1 Initial data toggle comes from incoming qTD DT bit. Host controller replaces DT bit in the queue head from the DT bit in the qTD. |
| 13–12 | EPS | Endpoint speed. This is the speed of the associated endpoint. 00 Full-speed (12 Mbps) 01 Low-speed (1.5 Mbps) 10 High-speed (480 Mbps) 11 Reserved, should be cleared This field must not be modified by the host controller. |
| 11–8 | EndPt | Endpoint number. Selects the particular endpoint number on the device serving as the data source or sink. |
| 7 | I | Inactivate on next transaction. This bit is used by system software to request that the host controller set the Active bit to zero. This field is only valid when the queue head is in the periodic schedule and the EPS field indicates a full- or low-speed endpoint. Setting this bit when the queue head is in the asynchronous schedule or the EPS field indicates a high-speed device yields undefined results. |
| 6–0 | Device Address | Selects the specific device serving as the data source or sink. |

Table 17-60. Endpoint Capabilities: Queue Head DWord 2

| Bits | Name | Description |
|-------|-------------|--|
| 31–30 | Mult | High-bandwidth pipe multiplier. This field is a multiplier used to key the host controller as the number of successive packets the host controller may submit to the endpoint in the current execution. The host controller makes the simplifying assumption that software properly initializes this field (regardless of location of queue head in the schedules or other run time parameters). 00 Reserved, should be cleared. A zero in this field yields undefined results. 01 One transaction to be issued for this endpoint per microframe 10 Two transactions to be issued for this endpoint per microframe 11 Three transactions to be issued for this endpoint per microframe |
| 29–23 | Port Number | This field is ignored by the host controller unless the EPS field indicates a full- or low-speed device. The value is the port number identifier on the USB 2.0 hub (for hub at device address Hub Addr below), below which the full- or low-speed device associated with this endpoint is attached. This information is used in the split-transaction protocol. |
| 22–16 | Hub Addr | This field is ignored by the host controller unless the EPS field indicates a full- or low-speed device. The value is the USB device address of the USB 2.0 hub below which the full- or low-speed device associated with this endpoint is attached. This field is used in the split-transaction protocol. |

Table 17-60. Endpoint Capabilities: Queue Head DWord 2 (continued)

| Bits | Name | Description |
|------|---------------|--|
| 15–8 | μFrame C-mask | This field is ignored by the host controller unless the EPS field indicates this device is a low- or full-speed device and this queue head is in the periodic list. This field (along with the Active and SplitX-state fields) is used to determine during which microframes the host controller should execute a complete-split transaction. When the criteria for using this field are met, a zero value in this field has undefined behavior. This field is used by the host controller to match against the three low-order bits of the FRINDEX register. If the FRINDEX register bits decode to a position where the μFrame C- mask field is a one, then this queue head is a candidate for transaction execution. There may be more than one bit in this mask set. |
| 7–0 | μFrame S-mask | Interrupt schedule mask. This field is used for all endpoint speeds. Software should set this field to a zero when the queue head is on the asynchronous schedule. A non-zero value in this field indicates an interrupt endpoint. The host controller uses the value of the three low-order bits of the FRINDEX register as an index into a bit position in this bit vector. If the μFrame S-mask field has a one at the indexed bit position then this queue head is a candidate for transaction execution. If the EPS field indicates the endpoint is a high-speed endpoint, then the transaction executed is determined by the PID_Code field contained in the execution area. This field is also used to support split transaction types: Interrupt (IN/OUT). This condition is true when this field is non-zero and the EPS field indicates this is either a full- or low-speed device. A zero value in this field, in combination with existing in the periodic frame list has undefined results. |

17.5.6.3 Transfer Overlay

The nine DWords in this area represent a transaction working space for the host controller. The general operational model is that the host controller can detect whether the overlay area contains a description of an active transfer. If it does not contain an active transfer, then it follows the queue head horizontal link pointer to the next queue head. The host controller will never follow the next transfer queue element or alternate queue element pointers unless it is actively attempting to advance the queue. For the duration of the transfer, the host controller keeps the incremental status of the transfer in the overlay area. When the transfer is complete, the results are written back to the original queue element.

The DWord3 of a queue head contains a pointer to the source qTD currently associated with the overlay. The host controller uses this pointer to write back the overlay area into the source qTD after the transfer is complete.

Table 17-61 describes the current qTD link pointer.

Table 17-61. Current qTD Link Pointer

| Bits | Name | Description |
|------|---------------------|---|
| 31–5 | Current qTD Pointer | Current element transaction descriptor link pointer. This field contains the address Of the current transaction being processed in this queue and corresponds to memory address signals [31:5], respectively. |
| 4–0 | — | Reserved, should be cleared. These bits are ignored by the host controller when using the value as an address to write data. The actual value may vary depending on the usage. |

The DWords 4–11 of a queue head are the transaction overlay area. This area has the same base structure as a queue element transfer descriptor. The queue head utilizes the reserved fields of the page pointers to implement tracking the state of split transactions.

This area is characterized as an overlay because when the queue is advanced to the next queue element, the source queue element is merged onto this area. This area serves an execution cache for the transfer.

Table 17-62 describes the host-controller rules for bits in overlay.

Table 17-62. Host-Controller Rules for Bits in Overlay (DWords 5, 6, 8, and 9)

| DWord | QH Offset | Bits | Name | Description |
|-------|-----------|-------|-------------|--|
| 5 | 0x14 | 4–1 | NakCnt | Nak counter—RW. This field is a counter the host controller decrements whenever a transaction for the endpoint associated with this queue head results in a Nak or Nyet response. This counter is reloaded from RL before a transaction is executed during the first pass of the reclamation list (relative to an Asynchronous List Restart condition). It is also loaded from RL during an overlay. |
| 6 | 0x18 | 31 | dt | Data toggle. The Data toggle control controls whether the host controller preserves this bit when an overlay operation is performed. |
| 6 | 0x18 | 15 | ioc | Interrupt on complete. The ioc control bit is always inherited from the source qTD when the overlay operation is performed. |
| 6 | 0x18 | 11–10 | Cerr | Error counter. Copied from the qTD during the overlay and written back during queue advancement. |
| 6 | 0x18 | 0 | Status[0] | Ping state (P)/ERR. If the EPS field indicates a high-speed endpoint, then this field should be preserved during the overlay operation. |
| 8 | 0x20 | 7–0 | C-prog-mask | Split-transaction complete-split progress. Initialized to zero during any overlay. This field is used to track the progress of an interrupt split-transaction. |
| 9 | 0x24 | 11–5 | S-bytes | Software must ensure that the S-bytes field in a qTD is zero before activating the qTD. Keeps track of the number of bytes sent or received during an IN or OUT split transaction. |
| 9 | 0x24 | 4–0 | FrameTag | Split-transaction frame tag. Initialized to zero during any overlay. This field is used to track the progress of an interrupt split-transaction. |

17.5.7 Periodic Frame Span Traversal Node (FSTN)

The periodic frame span traversal node (FSTN) data structure, shown in Figure 17-42, is to be used only for managing full- and low-speed transactions that span a host-frame boundary. Software must not use an FSTN in the asynchronous schedule. An FSTN in the asynchronous schedule results in undefined behavior. Software must not use the FSTN feature with a host controller whose HCIVERSION register indicates a revision implementation under 0x0096. Note that FSTNs were not defined for EHCI implementations before Revision 0.96 of the EHCI Specification and their use may yield undefined results.

| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | offset |
|--------------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|---|---|---|---|---|----|-----|---|------|---|--------|
| Normal Path Link Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | 00 | Typ | T | 0x00 | | |
| Back Path Link Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | 00 | Typ | T | 0x04 | | |

Figure 17-42. Frame Span Traversal Node Structure

17.5.7.1 FTSN Normal Path Pointer

The first DWord of an FSTN contains a link pointer to the next schedule object. This object can be of any valid periodic schedule data type. [Table 17-63](#) describes the FTSN normal path pointer.

Table 17-63. FTSN Normal Path Pointer

| Bits | Name | Description |
|------|------|--|
| 31–5 | NPLP | Normal path link pointer. Contains the address of the next data object to be processed in the periodic list and corresponds to memory address signals [31:5], respectively. |
| 4–3 | — | Reserved, should be cleared. These bits must be written as 0s. |
| 2–1 | Typ | Indicates to the host controller whether the item referenced is a iTD/siTD, QH, or FSTN. This allows the host controller to perform the proper type of processing on the item after it is fetched. 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node) |
| 0 | T | Terminate. 0 Link pointer is valid. 1 Link pointer field is not valid. |

17.5.7.2 FSTN Back Path Link Pointer

The second DWord of an FTSN node contains a link pointer to a queue head. If the T-bit in this pointer is a zero, then this FSTN is a Save-Place indicator. Its Typ field must be set by software to indicate the target data structure is a queue head. If the T-bit in this pointer is set, then this FSTN is the Restore indicator. When the T-bit is a one, the host controller ignores the Typ field.

[Table 17-64](#) describes the FTSN back path link pointer.

Table 17-64. FSTN Back Path Link Pointer

| Bits | Name | Description |
|------|------|--|
| 31–5 | BPLP | Back path link pointer. Contains the address of a queue head. This field corresponds to memory address signals [31:5], respectively. |
| 4–3 | — | Reserved, should be cleared. These bits must be written as 0s. |
| 2–1 | Typ | Software must ensure this field is set to indicate the target data structure is a Queue Head (01). Any other value in this field yields undefined results. |
| 0 | T | Terminate. 0 Link pointer is valid (that is, the host controller may use bits 31–5 (in combination with the CTRLDSSEGMENT register if applicable) as a valid memory address). This value also indicates that this FSTN is a Save-Place indicator. 1 Link pointer field is not valid (that is, the host controller must not use bits 31–5 (in combination with the CTRLDSSEGMENT register if applicable) as a valid memory address). This value also indicates that this FSTN is a Restore indicator. |

17.6 Host Operations

The general operational model for the USB DR module in host mode is defined by the Enhanced Host Controller Interface (EHCI) Specification. The EHCI specification describes the register-level interface for a host controller for the USB Revision 2.0. It includes a description of the hardware/software interface between system software and host controller hardware. Information concerning the initialization of the USB module is included in the following section; however, the full details of the EHCI specification are beyond the scope of this document.

17.6.1 Host Controller Initialization

After initial power-on or host controller reset (hardware or through USBCMD[RST]), all of the operational registers are at their default values. After a hardware reset, only the operational registers are at their default values.

The device can be configured for either the internal UTMI PHY or an external ULPI PHY enabled with various clocking options. The configuration of each PHY interface and the various clocking options are detailed below:

To configure the internal UTMI PHY, the following initialization sequence is required:

1. After power-on reset, the UTMI PHY will be in disabled state and the PLL will be held reset.
2. Set the CONTROL[REFSEL] bits to correspond to the appropriate UTMI PLL reference clock frequency.
3. Set the CONTROL[CLKIN_SEL] bits to select the source and divider of the UTMI reference clock.
4. Set the CONTROL[PHY_CLK_SEL] bits to select the UTMI PHY as the source of USB controller PHY clock.
5. Set the CONTROL[UTMI_PHY_EN] to enable the UTMI PHY and release the PLL.
6. Wait for PHY clock to become valid. This can be determined by polling the CONTROL[PHY_CLK_VALID] status bit.

Once the PHY clock is valid the user can proceed to the host controller initialization phase.

To configure the external ULPI PHY the following initialization sequence is required:

1. The UTMI PHY should remain disabled if the ULPI is being used.
2. Set the CONTROL[PHY_CLK_SEL] bits to select the ULPI PHY as the source of USB controller PHY clock.
3. Wait for PHY clock to become valid. This can be determined by polling the CONTROL[PHY_CLK_VALID] status bit. Note that this bit is not valid once the CONTROL[USB_EN] bit is set

Once the PHY clock is valid the user can proceed to the host controller initialization phase.

In order to initialize the USB DR module, software should perform the following steps:

1. Set the controller mode to host mode. Optionally set USBMODE[SDIS] (streaming disable)

NOTE

Transitioning from device mode to host mode requires a host controller reset before modifying USBMODE.

2. Optionally modify the BURSTSIZE register.
3. Program the PTS field of the PORTSC register if using a non-ULPI PHY.
4. Set CONTROL[USB_EN].
5. Write the appropriate value to the USBINTR register to enable the appropriate interrupts.
6. Write the base address of the periodic frame list to the PERIODICLIST BASE register. If there are no work items in the periodic schedule, all elements of the periodic frame list should have their T-Bits set.
7. Write the USBCMD register to set the desired interrupt threshold, frame list size (if applicable) and turn the controller by setting the RS bit.

At this point, the USB DR module is up and running and the port registers begin reporting device connects. System software can enumerate a port through the reset process (where the port is in the enabled state). At this point, the port is active with SOFs occurring down the enabled port enabled high-speed ports, but the schedules have not yet been enabled. The EHCI host controller will not transmit SOFs to enabled Full- or Low-speed ports.

In order to communicate with devices via the asynchronous schedule, system software must write the ASYNDLISTADDR register with the address of a control or bulk queue head. Software must then enable the asynchronous schedule by writing a one to USBCMD[ASE]. In order to communicate with devices via the periodic schedule, system software must enable the periodic schedule by writing a one to USBCMD[PSE]. Note that the schedules can be turned on before the first port is reset (and enabled).

Any time the USBCMD register is written, system software must ensure the appropriate bits are preserved, depending on the intended operation.

17.6.2 Power Port

The HCSPARAMS[PPC] bit indicates whether the USB 2.0 host controller has port power control. When the PPC bit is set, the host controller supports port power switches. Each available switch has an output enable. PPE is controlled based on the state of the combination bits PPC bit, EHCI Configured (CF)-bit and individual Port Power (PP) bits.

17.6.3 Reporting Over-Current

Host ports by definition are power providers on USB. Whether the ports are considered high- or low-powered is a platform implementation issue. The EHCI PORTSC register has an over-current status and over-current change bit. The functionality of these bits is specified in the USB Specification Revision 2.0.

The over current detection and limiting logic resides outside the DR logic. The over-current condition effects the following bits in the PORTSC register on the EHCI port:

- Over-current active bit (OCA) is set. When the over-current condition goes away, the OCA will transition from a one to a zero.
- Over-current change bit (OCC) is set. On every transition of OCA, the controller will set OCC to a one. Software sets OCC to a zero by writing a one to this bit.
- Port enabled/disabled bit (PE) is cleared. When this change bit gets set, USBSTS[PCI] (the port change detect bit) is set.
- Port power (PP) bit may optionally be cleared. There is no requirement in USB that a power provider shut off power in an over current condition. It is sufficient to limit the current and leave power applied. When OCC transitions from a zero to a one, the controller also sets USBSTS[PCI] to a one. In addition, if the Port Change Interrupt Enable bit, USBINTR[PCE], is a one, the controller issues an interrupt to the system. Refer to [Table 17-65](#) for summary of behavior for over-current detection when the controller is halted (suspended from a device component point of view).

17.6.4 Suspend/Resume

The host controller provides an equivalent suspend and resume model as that defined for individual ports in a USB 2.0 hub. Control mechanisms are provided to allow system software to suspend and resume individual ports. The mechanisms allow the individual ports to be resumed completely through software initiation. Other control mechanisms are provided to parameterize the host controller's response (or sensitivity) to external resume events. In this discussion, host-initiated, or software-initiated resumes are called Resume Events/Actions; bus-initiated resume events are called wake-up events. The classes of wakeup events are:

- Remote-wakeup enabled device asserts resume signaling. In similar kind to USB 2.0 hubs, when in host mode the host controller responds to explicit device resume signaling and wake up the system (if necessary).
- Port connect and disconnect and over-current events. Sensitivity to these events can be turned on or off by using the port control bits in the PORTSC register.

Selective suspend is a feature supported by the PORTSC register. It is used to place specific ports into a suspend mode. This feature is used as a functional component for implementing the appropriate power management policy implemented in a particular operating system. When system software intends to suspend the bus, it should suspend the enabled port, then shut off the controller by setting the USBCMD[RS] to a zero.

When a wake event occurs the system will resume operation and system software must set the RS bit to a one and resume the suspended port.

17.6.4.1 Port Suspend/Resume

System software places the USB into suspend mode by writing a one into the appropriate PORTSC Suspend bit. Software must only set the Suspend bit when the port is in the enabled state (Port Enabled bit is a one).

The host controller may evaluate the Suspend bit immediately or wait until a microframe or frame boundary occurs. If evaluated immediately, the port is not suspended until the current transaction (if one is executing) completes. Therefore, there may be several microframes of activity on the port until the host controller evaluates the Suspend bit. The host controller must evaluate the Suspend bit at least every frame boundary.

System software can initiate a resume on the suspended port by writing a one to PORTSC[FPR]. Software should not attempt to resume a port unless the port reports that it is in the suspended state. If system software sets PORTSC[FPR] when the port is not in the suspended state, the resulting behavior is undefined. In order to assure proper USB device operation, software must wait for at least 10 milliseconds after a port indicates that it is suspended (Suspend bit is a one) before initiating a port resume through PORTSC[FPR]. When PORTSC[FPR] is set, the host controller sends resume signaling down the port. System software times the duration of the resume (nominally 20 milliseconds) then clears PORTSC[FPR]. When the host controller receives the write to transition PORTSC[FPR] to zero, it completes the resume sequence as defined in the USB specification, and clears both PORTSC[FPR] and PORTSC[SUSP]. Software-initiated port resumes do not affect the port change detect bit (USBSTS[PCI]) nor do they cause an interrupt if USBINTR[PCE] (port change interrupt enable) is a one. When a wake event occurs on a suspended port, the resume signaling is detected by the port and the resume is reflected downstream within 100 µsec. The port's PORTSC[FPR] bit is set and USBSTS[PCI] is set. If USBINTR[PCE] is a one, the host controller issues a hardware interrupt.

System software observes the resume event on the port, delays a port resume time (nominally 20 milliseconds), then terminates the resume sequence by clearing PORTSC[FPR] in the port. The host controller receives the write of zero to PORTSC[FPR], terminates the resume sequence and clears PORTSC[FPR] and PORTSC[SUSP]. Software can determine that the port is enabled (not suspended) by sampling the PORTSC register and observing that the SUSP and FPR bits are zero. Software must ensure that the host controller is running (that is, USBSTS[HCH] is a zero), before terminating a resume by clearing the port's PORTSC[FPR] bit. If HCH is a one when PORTSC[FPR] is cleared, then SOFs will not occur down the enabled port and the device will return to suspend mode in a maximum of 10 milliseconds.

Table 17-65 summarizes the wake-up events. Whenever a resume event is detected, USBSTS[PCI] is set. If USBINTR[PCE] (port change interrupt enable) is a one, the host controller also generates an interrupt on the resume event. Software acknowledges the resume event interrupt by clearing the USBSTS[PCI].

Table 17-65. Behavior During Wake-Up Events

| Port Status and Signaling Type | Signaled Port Response | Device State | |
|--|--|--------------|--------|
| | | D0 | not D0 |
| Port disabled, resume K-State received | No effect | N/A | N/A |
| Port suspended, Resume K-State received | Resume reflected downstream on signaled port. PORTSC[FPR] is set. USBSTS[PCI] is set. | [1], [2] | [2] |
| Port is enabled, disabled or suspended, and the port's WKDSCNNT_E bit, PORTSC[WKDS], is set. A disconnect is detected. | Depending on the initial port state, the PORTSC Connect (CCS) and Enable (PE) status bits are cleared, and the Connect Change status bit (CSC) is set. USBSTS[PCI] is set. | [1], [2] | [2] |

Table 17-65. Behavior During Wake-Up Events (continued)

| Port Status and Signaling Type | Signaled Port Response | Device State | |
|--|--|--------------|--------|
| | | D0 | not D0 |
| Port is enabled, disabled or suspended, and the port's WKDSCNNT_E bit, PORTSC[WKDS], is cleared. A disconnect is detected. | Depending on the initial port state, the PORTSC Connect (CCS) and Enable (PE) status bits are cleared, and the Connect Change status bit (CSC) is set. USBSTS[PCI] is set. | [1], [3] | [3] |
| Port is not connected and the port's WKCNTNT_E bit is a one. A connect is detected. | PORTSC Connect Status (CCS) and Connect Status Change (CSC) bits are set. USBSTS[PCI] is set. | [1], [2] | [2] |
| Port is not connected and the port's WKCNTNT_E bit is a zero. A connect is detected. | PORTSC Connect Status (CCS) and Connect Status Change (CSC) bits are set. USBSTS[PCI] is set. | [1], [3] | [3] |
| Port is connected and the port's WKOC_E bit is a one. An over-current condition occurs. | PORTSC Over-current Active (OCA), Over-current Change (OCC) bits are set. If Port Enable/Disable bit (PE) is a one, it is cleared. USBSTS[PCI] is set | [1], [2] | [2] |
| Port is connected and the port's WKOC_E bit is a zero. An over-current condition occurs. | PORTSC Over-current Active (OCA), Over-current Change (OCC) bits are set. If Port Enable/Disable bit (PE) is a one, it is cleared. USBSTS[PCI] is set. | [1], [3] | [3] |

¹ Hardware interrupt issued if USBINTR[PCE] (port change interrupt enable) is set.

² PME# asserted if enabled (Note: PME Status must always be set).

³ PME# not asserted.

17.6.5 Schedule Traversal Rules

The host controller executes transactions for devices using a simple, shared-memory schedule. The schedule is comprised of a few data structures, organized into two distinct lists. The data structures are designed to provide the maximum flexibility required by USB, minimize memory traffic and hardware/software complexity.

System software maintains two schedules for the host controller: a periodic schedule and an asynchronous schedule. The root of the periodic schedule is the PERIODICLISTBASE register. See [Section 17.3.2.6, “Periodic Frame List Base Address Register \(PERIODICLISTBASE\),”](#) for more information. The PERIODICLISTBASE register is the physical memory base address of the periodic frame list. The periodic frame list is an array of physical memory pointers. The objects referenced from the frame list must be valid schedule data structures as defined in [Section 17.5, “Host Data Structures.”](#) In each microframe, if the periodic schedule is enabled (see) then the host controller must execute from the periodic schedule before executing from the asynchronous schedule. It will only execute from the asynchronous schedule after it encounters the end of the periodic schedule. The host controller traverses the periodic schedule by constructing an array offset reference from the PERIODICLISTBASE and the FRINDEX registers (see [Figure 17-43](#)). It fetches the element and begins traversing the graph of linked schedule data structures.

The end of the periodic schedule is identified by a next link pointer of a schedule data structure having its T-bit set. When the host controller encounters a T-Bit set during a horizontal traversal of the periodic list, it interprets this as an End-Of-Periodic-List mark. This causes the host controller to cease working on the periodic schedule and transitions immediately to traversing the asynchronous schedule. Once this

transition is made, the host controller executes from the asynchronous schedule until the end of the microframe.

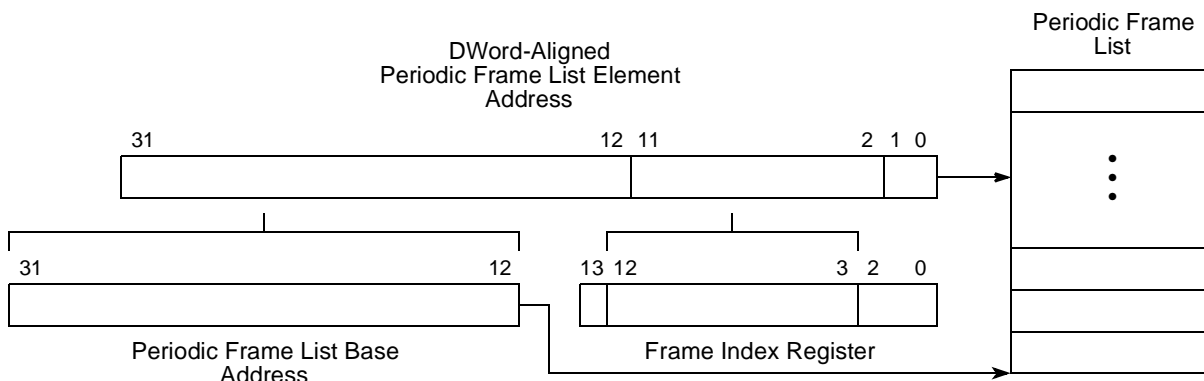


Figure 17-43. Derivation of Pointer into Frame List Array

When the host controller determines that it is time to execute from the asynchronous list, it uses the operational register ASYNCLISTADDR to access the asynchronous schedule, as shown in Figure 17-44.




Figure 17-44. General Format of Asynchronous Schedule List

The ASYNCLISTADDR register contains a physical memory pointer to the next queue head. When the host controller makes a transition to executing the asynchronous schedule, it begins by reading the queue head referenced by the ASYNCLISTADDR register. Software must set queue head horizontal pointer T-bits to a zero for queue heads in the asynchronous schedule.

17.6.6 Periodic Schedule Frame Boundaries vs. Bus Frame Boundaries

The USB Specification Revision 2.0 requires that the frame boundaries (SOF frame number changes) of the high-speed bus and the full- and low-speed bus(es) below USB 2.0 hubs be strictly aligned. Super-imposed on this requirement is that USB 2.0 hubs manage full- and low-speed transactions via a microframe pipeline (see start- (SS) and complete- (CS) splits illustrated in Figure 17-45). A simple, direct projection of the frame boundary model into the host controller interface schedule architecture creates tension (complexity for both hardware and software) between the frame boundaries and the scheduling mechanisms required to service the full- and low-speed transaction translator periodic pipelines.

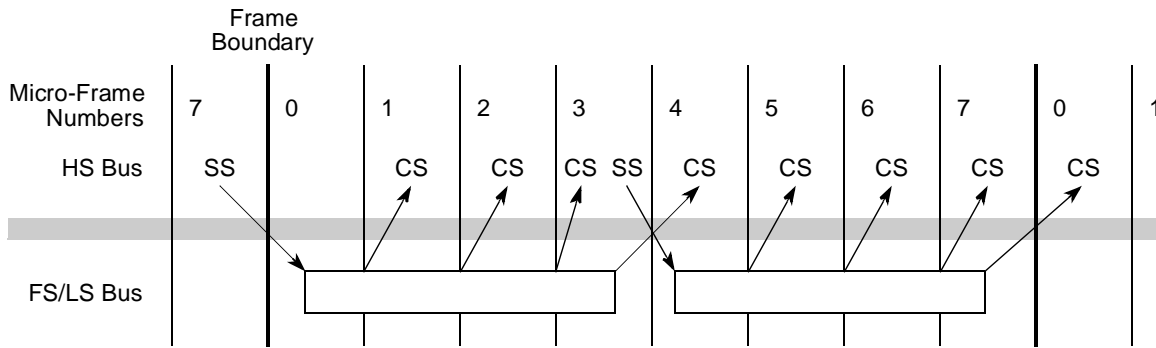


Figure 17-45. Frame Boundary Relationship Between HS Bus and FS/LS Bus

The simple projection, as [Figure 17-45](#) illustrates, introduces frame-boundary wrap conditions for scheduling on both the beginning and end of a frame. In order to reduce the complexity for hardware and software, the host controller is required to implement a one microframe phase shift for its view of frame boundaries. The phase shift eliminates the beginning of frame and frame-wrap scheduling boundary conditions.

The implementation of this phase shift requires that the host controller use one register value for accessing the periodic frame list and another value for the frame number value included in the SOF token. These two values are separate, but tightly coupled. The periodic frame list is accessed via the Frame List Index Register (FRINDEX). Bits FRINDEX[2–0], represent the microframe number. The SOF value is coupled to the value of FRINDEX[13–3]. Both FRINDEX[13–3] and the SOF value are incremented based on FRINDEX[2–0]. It is required that the SOF value be delayed from the FRINDEX value by one microframe. The one microframe delay yields a host controller periodic schedule and bus frame boundary relationship as illustrated in [Figure 17-46](#). This adjustment allows software to trivially schedule the periodic start and complete-split transactions for full-and low-speed periodic endpoints, using the natural alignment of the periodic schedule interface.

[Figure 17-46](#) illustrates how periodic schedule data structures relate to schedule frame boundaries and bus frame boundaries. To aid the presentation, two terms are defined. The host controller's view of the

1-millisecond boundaries is called H-Frames. The high-speed bus's view of the 1-millisecond boundaries is called B-Frames.




Figure 17-46. Relationship of Periodic Schedule Frame Boundaries to Bus Frame Boundaries

H-Frame boundaries for the host controller correspond to increments of FRINDEX[13–3]. Microframe numbers for the H-Frame are tracked by FRINDEX[2–0]. B-Frame boundaries are visible on the high-speed bus via changes in the SOF token's frame number. Microframe numbers on the high-speed bus are only derived from the SOF token's frame number (that is, the high-speed bus will see eight SOFs with the same frame number value). H-Frames and B-Frames have the fixed relationship (that is, B-Frames lag H-Frames by one microframe time) illustrated in Figure 17-46. The host controller's periodic schedule is naturally aligned to H-Frames. Software schedules transactions for full- and low-speed periodic endpoints relative the H-Frames. The result is these transactions execute on the high-speed bus at exactly the right time for the USB 2.0 hub periodic pipeline. As described in Section 17.3.2.4, “Frame Index Register (FRINDEX),” the SOF Value can be implemented as a shadow register (in this example, called SOFV), which lags the FRINDEX register bits [13–3] by one microframe count. Table 17-66 illustrates the required relationship between the value of FRINDEX and the value of SOFV. This lag behavior can be accomplished by incrementing FRINDEX[13–3] based on carry-out on the 7 to 0 increment of FRINDEX[2–0] and incrementing SOFV based on the transition of 0 to 1 of FRINDEX[2–0].

Software is allowed to write to FRINDEX. Section 17.3.2.4, “Frame Index Register (FRINDEX),” provides the requirements that software should adhere when writing a new value in FRINDEX.

Table 17-66. Operation of FRINDEX and SOFV (SOF Value Register)

| Current | | | Next | | |
|---------------|------|--------------|---------------|------|--------------|
| FRINDEX[13–3] | SOFV | FRINDEX[2–0] | FRINDEX[13–3] | SOFV | FRINDEX[2–0] |
| N | N | 111 | N+1 | N | 000 |
| N+1 | N | 000 | N+1 | N+1 | 001 |
| N+1 | N+1 | 001 | N+1 | N+1 | 010 |
| N+1 | N+1 | 010 | N+1 | N+1 | 011 |
| N+1 | N+1 | 011 | N+1 | N+1 | 100 |
| N+1 | N+1 | 100 | N+1 | N+1 | 101 |
| N+1 | N+1 | 101 | N+1 | N+1 | 110 |
| N+1 | N+1 | 110 | N+1 | N+1 | 111 |

17.6.7 Periodic Schedule

The periodic schedule traversal is enabled or disabled through USBCMD[PSE] (periodic schedule enable). If USBCMD[PSE] is cleared, then the host controller simply does not try to access the periodic frame list via the PERIODICLISTBASE register. Likewise, when USBCMD[PSE] is a one, then the host controller does use the PERIODICLISTBASE register to traverse the periodic schedule. The host controller will not react to modifications to USBCMD[PSE] immediately. In order to eliminate conflicts with split transactions, the host controller evaluates USBCMD[PSE] only when FRINDEX[2–0] is zero. System software must not disable the periodic schedule if the schedule contains an active split transaction work item that spans the 0b000 microframe. These work items must be removed from the schedule before USBCMD[PSE] is cleared. USBSTS[PS] (periodic schedule status) indicates status of the periodic schedule. System software enables (or disables) the periodic schedule by setting (or clearing) USBCMD[PSE]. Software then can poll USBSTS[PS] to determine when the periodic schedule has made the desired transition. Software must not modify USBCMD[PSE] unless the value of USBCMD[PSE] equals that of USBSTS[PS].

The periodic schedule is used to manage all isochronous and interrupt transfer streams. The base of the periodic schedule is the periodic frame list. Software links schedule data structures to the periodic frame list to produce a graph of scheduled data structures. The graph represents an appropriate sequence of transactions on the USB. Figure 17-47 illustrates isochronous transfers (using iTDs and siTDs) with a period of one are linked directly to the periodic frame list. Interrupt transfers (are managed with queue heads) and isochronous streams with periods other than one are linked following the period-one iTD/siTDs. Interrupt queue heads are linked into the frame list ordered by poll rate. Longer poll rates are

linked first (for example, closest to the periodic frame list), followed by shorter poll rates, with queue heads with a poll rate of one, on the very end.

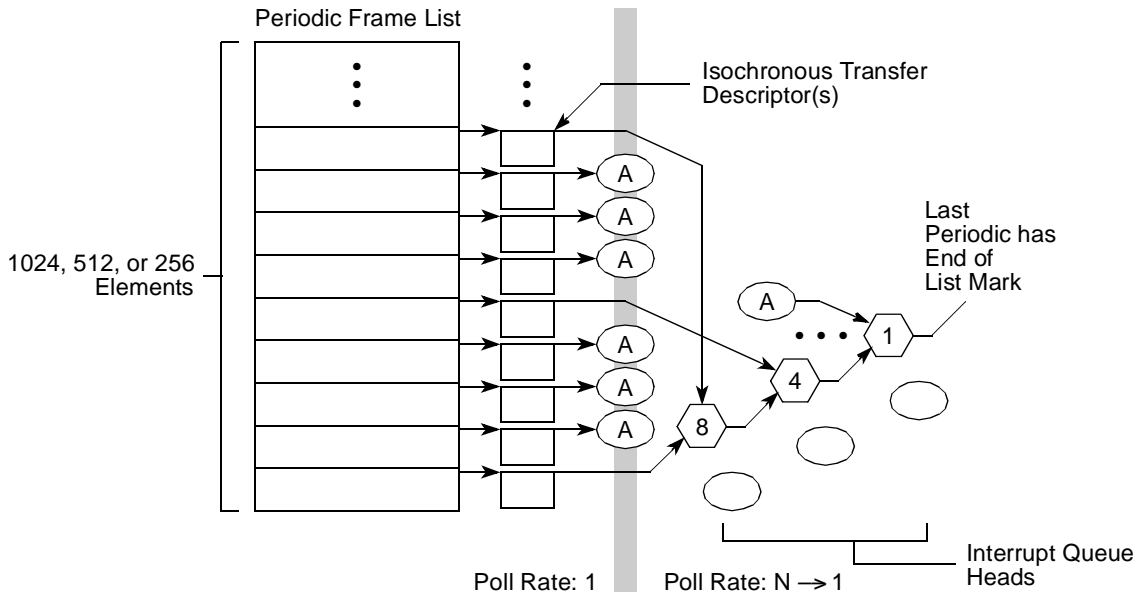


Figure 17-47. Example Periodic Schedule

17.6.8 Managing Isochronous Transfers Using iTDs

The structure of an iTD is presented in isochronous (high-speed) transfer descriptor (iTID). There are four distinct sections to an iTD:

- Next link pointer
 - This is the first field.
 - This field is for schedule linkage purposes only.
- Transaction description array
 - This is an eight-element array.
 - Each element represents control and status information for one microframe's worth of transactions for a single high-speed isochronous endpoint.
- Buffer page pointer array
 - This is a 7-element array of physical memory pointers to data buffers.
 - These are 4K aligned pointers to physical memory.
- Endpoint capabilities
 - This area utilizes the unused low-order 12 bits of the buffer page pointer array.
 - Its fields are used across all transactions executed for this iTD, including endpoint addressing, transfer direction, maximum packet size and high-bandwidth multiplier.

17.6.8.1 Host Controller Operational Model for iTDs

The host controller uses FRINDEX register bits 12–3 to index into the periodic frame list. This means that the host controller visits each frame list element eight consecutive times before incrementing to the next periodic frame list element. Each iTD contains eight transaction descriptions, which map directly to FRINDEX register bits 2–0. Each iTD can span 8 microframes worth of transactions. When the host controller fetches an iTD, it uses FRINDEX register bits 2–0 to index into the transaction description array. When the first iTD in the periodic list is traversed after periodic schedule is enabled, the value of FRINDEX[2:0] may be other than 0, so the first transaction issued by the controller may be any of the eight available active transactions. If the active bit in the Status field of the indexed transaction description is cleared, the host controller ignores the iTD and follows the Next pointer to the next schedule data structure.

When the indexed active bit is a one the host controller continues to parse the iTD. It stores the indexed transaction description and the general endpoint information (device address, endpoint number, maximum packet size, and so on). It also uses the Page Select (PG) field to index the buffer pointer array, storing the selected buffer pointer and the next sequential buffer pointer. For example, if PG field is a 0, then the host controller will store Page 0 and Page 1.

The host controller constructs a physical data buffer address by concatenating the current buffer pointer (as selected using the current transaction description's PG field) and the transaction description's Transaction Offset field. The host controller uses the endpoint addressing information and I/O-bit to execute a transaction to the appropriate endpoint. When the transaction is complete, the host controller clears the active bit and writes back any additional status information to the Status field in the currently selected transaction description.

The data buffer associated with the iTD must be virtually contiguous memory. Seven page pointers are provided to support eight high-bandwidth transactions regardless of the starting packet's offset alignment into the first page. A starting buffer pointer (physical memory address) is constructed by concatenating the page pointer (example: page 0 pointer) selected by the active transaction descriptions' PG (example value: 0b00) field with the transaction offset field. As the transaction moves data, the host controller must detect when an increment of the current buffer pointer will cross a page boundary. When this occurs the host controller simply replaces the current buffer pointer's page portion with the next page pointer (example: page 1 pointer) and continues to move data. The size of each bus transaction is determined by the value in the Maximum Packet Size field. An iTD supports high-bandwidth pipes via the Mult (multiplier) field. When the Mult field is 1, 2, or 3, the host controller executes the specified number of Maximum Packet sized bus transactions for the endpoint in the current microframe. In other words, the Mult field represents a transaction count for the endpoint in the current microframe. If the Mult field is zero, the operation of the host controller is undefined. The transfer description is used to service all transactions indicated by the Mult field.

For OUT transfers, the value of the Transaction n Length field represents the total bytes to be sent during the microframe. The Mult field must be set by software to be consistent with Transaction n Length and Maximum Packet Size. The host controller will send the bytes in Maximum Packet Sized portions. After each transaction, the host controller decrements it's local copy of Transaction n Length by Maximum Packet Size. The number of bytes the host controller sends is always Maximum Packet Size or Transaction n Length, whichever is less. The host controller advances the transfer state in the transfer description,

updates the appropriate record in the iTD and moves to the next schedule data structure. The maximum sized transaction supported is 3×1024 bytes.

For IN transfers, the host controller issues Mult transactions. It is assumed that software has properly initialized the iTD to accommodate all possible data. During each IN transaction, the host controller must use Maximum Packet Size to detect packet babble errors. The host controller keeps the sum of bytes received in the Transaction n Length field.

After all transactions for the endpoint have completed for the microframe, Transaction n Length contains the total bytes received. The following actions can occur:

- If the final value of Transaction n Length is less than the value of Maximum Packet Size, less data was received than was allowed for from the associated endpoint. This short packet condition does not set USBSTS[UI] (USB interrupt). The host controller does not detect this condition.
- If the device sends more than Transaction n Length or Maximum Packet Size bytes (whichever is less), the host controller sets the Babble Detected bit and clears the Active bit. Note, that the host controller is not required to update the iTD field Transaction n Length in this error scenario.
- If the Mult field is greater than one, the host controller automatically executes the value of Mult transactions. The host controller does not execute all Mult transactions in the following cases:
 - The endpoint is an OUT and Transaction n Length goes to zero before all the Mult transactions have executed (ran out of data).
 - The endpoint is an IN and the endpoint delivers a short packet, or an error occurs on a transaction before Mult transactions have been executed. The end of microframe may occur before all of the transaction opportunities have been executed. When this happens, the transfer state of the transfer description is advanced to reflect the progress that was made; the result is written back to the iTD; and the host controller proceeds to processing the next microframe.

17.6.8.2 Software Operational Model for iTDs

A client buffer request to an isochronous endpoint may span 1 to N microframes. When N is larger than one, system software may have to use multiple iTDs to read or write data with the buffer (if N is larger than eight, it must use more than one iTD).

Figure 17-48 illustrates the simple model of how a client buffer is mapped by system software to the periodic schedule (that is, the periodic frame list and a set of iTDs).




Figure 17-48. Example Association of iTDs to Client Request Buffer

On the right is the client description of its request. The description includes a buffer base address plus additional annotations to identify which portions of the buffer should be used with each bus transaction. In the middle is the iTD data structures used by the system software to service the client request. Each iTD can be initialized to service up to 24 transactions, organized into eight groups of up to three transactions each. Each group maps to one microframe's worth of transactions. The EHCI controller does not provide per transaction results within a microframe. It treats the per microframe transactions as a single logical transfer.

On the left is the host controller's frame list. System software establishes references from the appropriate locations in the frame list to each of the appropriate iTDs. If the buffer is large, system software can use a small set of iTDs to service the entire buffer. System software can activate the transaction description records (contained in each iTD) in any pattern required for the particular data stream.

As noted above, the client request includes a pointer to the base of the buffer and offsets into the buffer to annotate which buffer sections are to be used on each bus transaction that occurs on this endpoint. System software must initialize each transaction description in an iTD to ensure it uses the correct portion of the client buffer. For example, for each transaction description, the PG field is set to index the correct physical buffer page pointer and the Transaction Offset field is set relative to the correct buffer pointer page (for example, the same one referenced by the PG field). When the host controller executes a transaction it selects a transaction description record based on FRINDEX[2–0]. It then uses the current Page Buffer Pointer (as selected by the PG field) and concatenates to the transaction offset field. The result is a starting buffer address for the transaction. As the host controller moves data for the transaction, it must watch for a page wrap condition and properly advance to the next available Page Buffer Pointer. System software must not use the Page 6 buffer pointer in a transaction description where the length of the transfer will wrap

a page boundary. Doing so yields undefined behavior. The host controller hardware is not required to alias the page selector to page zero. USB 2.0 isochronous endpoints can specify a period greater than one. Software can achieve the appropriate scheduling by linking iTDs into the appropriate frames (relative to the frame list) and by setting appropriate transaction description elements active bits to a one.

17.6.8.2.1 Periodic Scheduling Threshold

The Isochronous Scheduling Threshold field in the HCCPARAMS capability register is an indicator to system software as to how the host controller pre-fetches and effectively caches schedule data structures. It is used by system software when adding isochronous work items to the periodic schedule. The value of this field indicates to system software the minimum distance it can update isochronous data (relative to the current location of the host controller execution in the periodic list) and still have the host controller process them.

The iTD and siTD data structures each describe 8 microframes worth of transactions. The host controller is allowed to cache one (or more) of these data structures in order to reduce memory traffic. There are three basic caching models that account for the fact the isochronous data structures span 8 microframes. The three caching models are: no caching, microframe caching and frame caching.

When software is adding new isochronous transactions to the schedule, it always performs a read of the FRINDEX register to determine the current frame and microframe the host controller is currently executing. Of course, there is no information about where in the microframe the host controller is, so a constant uncertainty factor of one microframe has to be assumed. Combining the knowledge of where the host controller is executing with the knowledge of the caching model allows the definition of simple algorithms for how closely software can reliably work to the executing host controller.

No caching is indicated with a value of zero in the Isochronous Scheduling Threshold field. The host controller may pre-fetch data structures during a periodic schedule traversal (per microframe) but will always dump any accumulated schedule state at the end of the microframe. At the appropriate time relative to the beginning of every microframe, the host controller always begins schedule traversal from the frame list. Software can use the value of the FRINDEX register (plus the constant 1 uncertainty-factor) to determine the approximate position of the executing host controller. When no caching is selected, software can add an isochronous transaction as near as 2 microframes in front of the current executing position of the host controller.

Frame caching is indicated with a non-zero value in bit [7] of the Isochronous Scheduling Threshold field. In the frame-caching model, system software assumes that the host controller caches one (or more) isochronous data structures for an entire frame (8 microframes). Software uses the value of the FRINDEX register (plus the constant 1 uncertainty) to determine the current microframe/frame (assume modulo 8 arithmetic in adding the constant 1 to the microframe number). For any current frame N , if the current microframe is 0 to 6, then software can safely add isochronous transactions to Frame $N + 1$. If the current microframe is 7, then software can add isochronous transactions to Frame $N + 2$.

Microframe caching is indicated with a non-zero value in the least-significant 3 bits of the Isochronous Scheduling Threshold field. System software assumes the host controller caches one or more periodic data structures for the number of microframes indicated in the Isochronous Scheduling Threshold field. For example, if the count value were 2, then the host controller keeps a window of two microframes worth of

state (current microframe, plus the next) on chip. On each microframe boundary, the host controller releases the current microframe state and begins accumulating the next microframe state.

17.6.9 Asynchronous Schedule

The asynchronous schedule traversal is enabled or disabled through USBCMD[ASE] (asynchronous schedule enable). If USBCMD[ASE] is cleared, then the host controller simply does not try to access the asynchronous schedule via the ASYNCLISTADDR register. Likewise, if USBCMD[ASE] is set, the host controller does use the ASYNCLISTADDR register to traverse the asynchronous schedule. Modifications to USBCMD[ASE] are not necessarily immediate. Rather the new value of the bit will only be taken into consideration the next time the host controller needs to use the value of the ASYNCLISTADDR register to get the next queue head.

USBSTS[AS] indicates status of the asynchronous schedule. System software enables (or disables) the asynchronous schedule by writing a one (or zero) to USBCMD[ASE]. Software then can poll USBSTS[AS] to determine when the asynchronous schedule has made the desired transition. Software must not modify USBCMD[ASE] unless the value of USBCMD[ASE] equals that of the USBSTS[AS] (asynchronous schedule status).

The asynchronous schedule is used to manage all Control and Bulk transfers. Control and Bulk transfers are managed using queue head data structures. The asynchronous schedule is based at the ASYNCLISTADDR register. The default value of the ASYNCLISTADDR register after reset is undefined and the schedule is disabled when USBCMD[ASE] is cleared.

Software may only write this register with defined results when the schedule is disabled, for example, USBCMD[ASE] and the USBSTS[AS] are cleared. System software enables execution from the asynchronous schedule by writing a valid memory address (of a queue head) into this register. Then software enables the asynchronous schedule by setting USBCMD[ASE]. The asynchronous schedule is actually enabled when USBSTS[AS] is set.

When the host controller begins servicing the asynchronous schedule, it begins by using the value of the ASYNCLISTADDR register. It reads the first referenced data structure and begins executing transactions and traversing the linked list as appropriate. When the host controller completes processing the asynchronous schedule, it retains the value of the last accessed queue head's horizontal pointer in the ASYNCLISTADDR register. Next time the asynchronous schedule is accessed, this is the first data structure that is serviced. This provides round-robin fairness for processing the asynchronous schedule.

A host controller completes processing the asynchronous schedule when one of the following events occur:

- The end of a microframe occurs.
- The host controller detects an empty list condition
- The schedule has been disabled through USBCMD[ASE].

The queue heads in the asynchronous list are linked into a simple circular list as shown in [Figure 17-44](#). Queue head data structures are the only valid data structures that may be linked into the asynchronous schedule. An isochronous transfer descriptor (iT_D or si_{TD}) in the asynchronous schedule yields undefined results.

The maximum packet size field in a queue head is sized to accommodate the use of this data structure for all non-isochronous transfer types. The USB Specification, Revision 2.0 specifies the maximum packet sizes for all transfer types and transfer speeds. System software should always parameterize the queue head data structures according to the core specification requirements.

17.6.9.1 Adding Queue Heads to Asynchronous Schedule

This is a software requirement section. There are two independent events for adding queue heads to the asynchronous schedule. The first is the initial activation of the asynchronous list. The second is inserting a new queue head into an activated asynchronous list.

Activation of the list is simple. System software writes the physical memory address of a queue head into the ASYNCLISTADDR register, then enables the list by setting USBCMD[ASE] to a one.

When inserting a queue head into an active list, software must ensure that the schedule is always coherent from the host controllers' point of view. This means that the system software must ensure that all queue head pointer fields are valid. For example qTD pointers have T-Bits set or reference valid qTDs and the Horizontal Pointer references a valid queue head data structure. The following algorithm represents the functional requirements:

```

InsertQueueHead (pQHeadCurrent, pQueueHeadNew)
--
-- Requirement: all inputs must be properly initialized.
--
-- pQHeadCurrent is a pointer to a queue head that is
-- already in the active list
-- pQHeadNew is a pointer to the queue head to be added
--
-- This algorithm links a new queue head into a existing
-- list
--
pQueueHeadNew.HorizontalPointer = pQueueHeadCurrent.HorizontalPointer
pQueueHeadCurrent.HorizontalPointer = physicalAddressOf(pQueueHeadNew)
End InsertQueueHead

```

17.6.9.2 Removing Queue Heads from Asynchronous Schedule

This is a software requirement section. There are two independent events for removing queue heads from the asynchronous schedule. The first is shutting down (deactivating) the asynchronous list. The second is extracting a single queue head from an activated list. Software deactivates the asynchronous schedule by setting USBCMD[ASE] to a zero. Software can determine when the list is idle when USBSTS[AS] is cleared. The normal mode of operation is that software removes queue heads from the asynchronous schedule without shutting it down. Software must not remove an active queue head from the schedule. Software should first deactivate all active qTDs, wait for the queue head to go inactive, then remove the queue head from the asynchronous list. Software removes a queue head from the asynchronous list using the following algorithm. Software merely must ensure all of the link pointers reachable by the host controller are kept consistent.

```

UnlinkQueueHead (pQHeadPrevious, pQueueHeadToUnlink, pQHeadNext)
--
-- Requirement: all inputs must be properly initialized.
--

```

```

-- pQHeadPrevious is a pointer to a queue head that
-- references the queue head to remove
-- pQHeadToUnlink is a pointer to the queue head to be
-- removed
-- pQheadNext is a pointer to a queue head still in the
-- schedule. Software provides this pointer with the
-- following strict rules:
-- if the host software is one queue head, then
-- pQHeadNext must be the same as
-- QueueheadToUnlink.HorizontalPointer. If the host
-- software is unlinking a consecutive series of
-- queue heads, QHeadNext must be set by software to
-- the queue head remaining in the schedule.
--
-- This algorithm unlinks a queue head from a circular list
--
pQueueHeadPrevious.HorizontalPointer = pQueueHeadToUnlink.HorizontalPointer
pQueueHeadToUnlink.HorizontalPointer = pQHeadNext
End UnlinkQueueHead
    
```

If software removes the queue head with the H-bit set, it must select another queue head still linked into the schedule and set its H-bit. This should be completed before removing the queue head. The requirement is that software keep one queue head in the asynchronous schedule, with its H-bit set. At the point software has removed one or more queue heads from the asynchronous schedule, it is unknown whether the host controller has a cached pointer to them. Similarly, it is unknown how long the host controller might retain the cached information, as it is implementation dependent and may be affected by the actual dynamics of the schedule load. Therefore, once software has removed a queue head from the asynchronous list, it must retain the coherency of the queue head (link pointers). It cannot disturb the removed queue heads until it knows that the host controller does not have a local copy of a pointer to any of the removed data structures.

The method software uses to determine when it is safe to modify a removed queue head is to handshake with the host controller. The handshake mechanism allows software to remove items from the asynchronous schedule, then execute a simple, lightweight handshake that is used by software as a key that it can free (or reuse) the memory associated the data structures it has removed from the asynchronous schedule.

The handshake is implemented with three bits in the host controller. The first bit is a command bit (USBCMD[IAA]—interrupt on async advance doorbell) that allows software to inform the host controller that something has been removed from its asynchronous schedule. The second bit is a status bit (USBSTS[AAI]—interrupt on async advance) that the host controller sets after it has released all on-chip state that may potentially reference one of the data structures just removed. When the host controller sets this status bit, it also clears the command bit. The third bit is an interrupt enable (USBINTR[AAE]—interrupt on async advance enable) that is matched with the status bit. If the status bit is set and the interrupt enable bit is set, then the host controller asserts a hardware interrupt.

Figure 17-49 illustrates a general example where consecutive queue heads (B and C) are unlinked from the schedule using the algorithm above. Before the unlink operation, the host controller has a copy of queue head A.

The unlink algorithm requires that as software unlinks each queue head, the unlinked queue head is loaded with the address of a queue head that will remain in the asynchronous schedule.

When the host controller observes that doorbell bit being set, it makes a note of the local reachable schedule information. In this example, the local reachable schedule information includes both queue heads (A & B). It is sufficient that the host controller can set the status bit (and clear the doorbell bit) as soon as it has traversed beyond current reachable schedule information (that is, traversed beyond queue head (B) in this example).

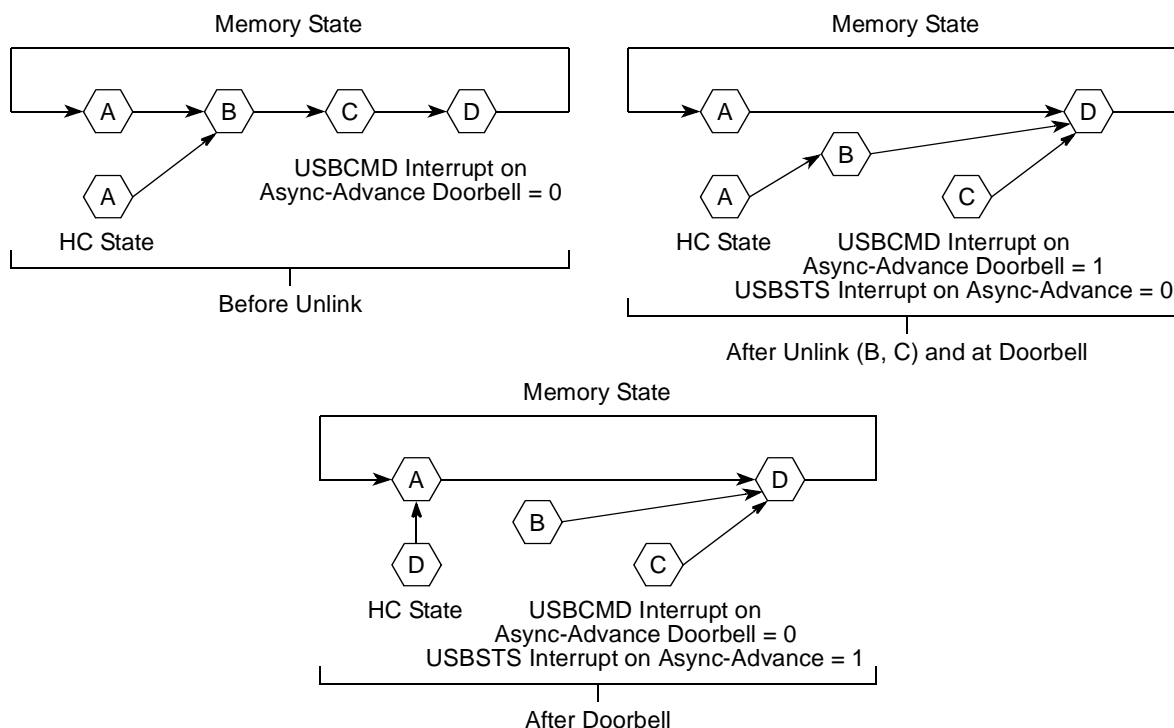


Figure 17-49. Generic Queue Head Unlink Scenario

Alternatively, a host controller implementation is allowed to traverse the entire asynchronous schedule list (for example, observed the head of the queue (twice)) before setting USBSTS[AAI].

Software may re-use the memory associated with the removed queue heads after it observes USBSTS[AAI] is set, following assertion of the doorbell. Software should acknowledge the interrupt on async advance status as indicated in the USBSTS register, before using the doorbell handshake again

17.6.9.3 Empty Asynchronous Schedule Detection

EHCI uses two bits to detect when the asynchronous schedule is empty. The queue head data structure (see [Figure 17-41](#)) defines an H-bit in the queue head, which allows software to mark a queue head as being the head of the reclaim list. host controller also keeps a 1-bit flag in the USBSTS register (Reclamation) that is cleared when the host controller observes a queue head with the H-bit set. The reclamation flag in the status register is set when any USB transaction from the asynchronous schedule is executed (or whenever the asynchronous schedule starts, see [Section 17.6.9.4, “Asynchronous Schedule Traversal: Start Event.”](#))

If the controller ever encounters an H-bit of one and a Reclamation bit of zero, the controller simply stops traversal of the asynchronous schedule.

Figure 17-50 shows an example illustrating the H-bit in a schedule.

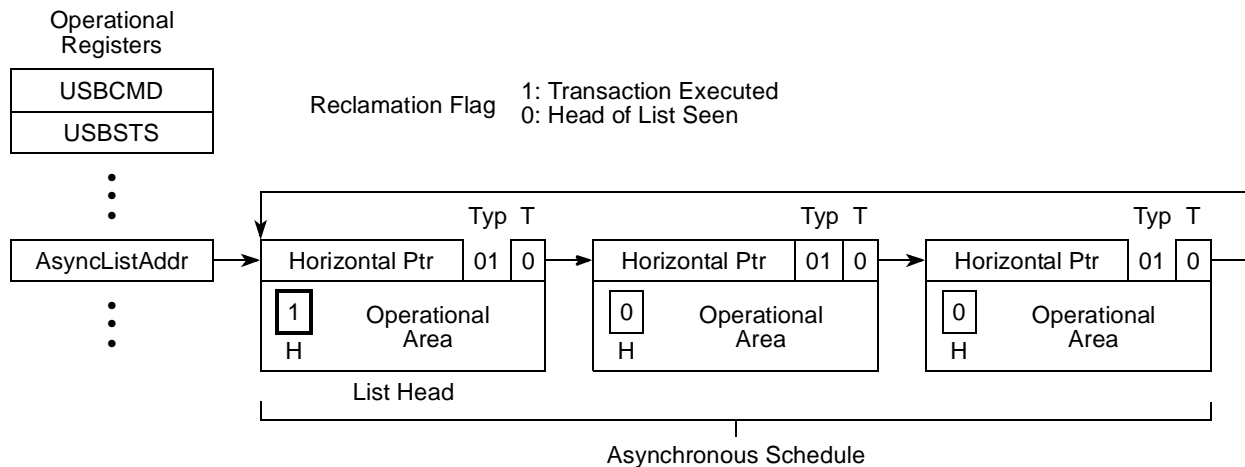


Figure 17-50. Asynchronous Schedule List with Annotation to Mark Head of List

17.6.9.4 Asynchronous Schedule Traversal: Start Event

Once the host controller has idled itself using the empty schedule detection, it naturally activates and begins processing from the Periodic Schedule at the beginning of each microframe. In addition, it may have idled itself early in a microframe. When this occurs (idles early in the microframe) the host controller must occasionally reactivate during the microframe and traverse the asynchronous schedule to determine whether any progress can be made. Asynchronous schedule Start Events are defined to be:

- Whenever the host controller transitions from the periodic schedule to the asynchronous schedule. If the periodic schedule is disabled and the asynchronous schedule is enabled, then the beginning of the microframe is equivalent to the transition from the periodic schedule, or
- The asynchronous schedule traversal restarts from a sleeping state.

17.6.9.5 Reclamation Status Bit (USBSTS Register)

The operation of the empty asynchronous schedule detection feature depends on the proper management of the Reclamation bit (RCL) in the USBSTS register. The host controller tests for an empty schedule just after it fetches a new queue head while traversing the asynchronous schedule. The host controller sets USBSTS[RCL] whenever an asynchronous schedule traversal Start Event occurs. USBSTS[RCL] is also set whenever the host controller executes a transaction while traversing the asynchronous schedule. The host controller clears USBSTS[RCL] whenever it finds a queue head with its H-bit set. Software should only set a queue head's H-bit if the queue head is in the asynchronous schedule. If software sets the H-bit in an interrupt queue head, the resulting behavior is undefined. The host controller may clear USBSTS[RCL] when executing from the periodic schedule.

17.6.10 Managing Control/Bulk/Interrupt Transfers via Queue Heads

This section presents an overview of how the host controller interacts with queuing data structures.

Queue heads use the Queue Element Transfer Descriptor (qTD) structure defined in [Section 17.5.5, “Queue Element Transfer Descriptor \(qTD\).”](#)

One queue head is used to manage the data stream for one endpoint. The queue head structure contains static endpoint characteristics and capabilities. It also contains a working area from where individual bus transactions for an endpoint are executed. Each qTD represents one or more bus transactions, which is defined in the context of the EHCI specification as a transfer.

The general processing model for the host controller's use of a queue head is simple:

- Read a queue head,
- Execute a transaction from the overlay area,
- Write back the results of the transaction to the overlay area
- Move to the next queue head.

If the host controller encounters errors during a transaction, the host controller will set one of the error reporting bits in the queue head's Status field. The Status field accumulates all errors encountered during the execution of a qTD (that is, the error bits in the queue head Status field are sticky until the transfer (qTD) has completed). This state is always written back to the source qTD when the transfer is complete. On transfer (for example, buffer or halt conditions) boundaries, the host controller must auto-advance (without software intervention) to the next qTD. Additionally, the hardware must be able to halt the queue so no additional bus transactions will occur for the endpoint and the host controller will not advance the queue.

17.6.10.1 Buffer Pointer List Use for Data Streaming with qTDs

A qTD has an array of buffer pointers, which is used to reference the data buffer for a transfer. The EHCI specification requires that the buffer associated with the transfer be virtually contiguous. This means that if the buffer spans more than one physical page, it must obey the following rules:

- The first portion of the buffer must begin at some offset in a page and extend through the end of the page.
- The remaining buffer cannot be allocated in small chunks scattered around memory. For each 4K chunk beyond the first page, each buffer portion matches to a full 4K page. The final portion, which may only be large enough to occupy a portion of a page, must start at the top of the page and be contiguous within that page.

Figure 17-51 illustrates these requirements.

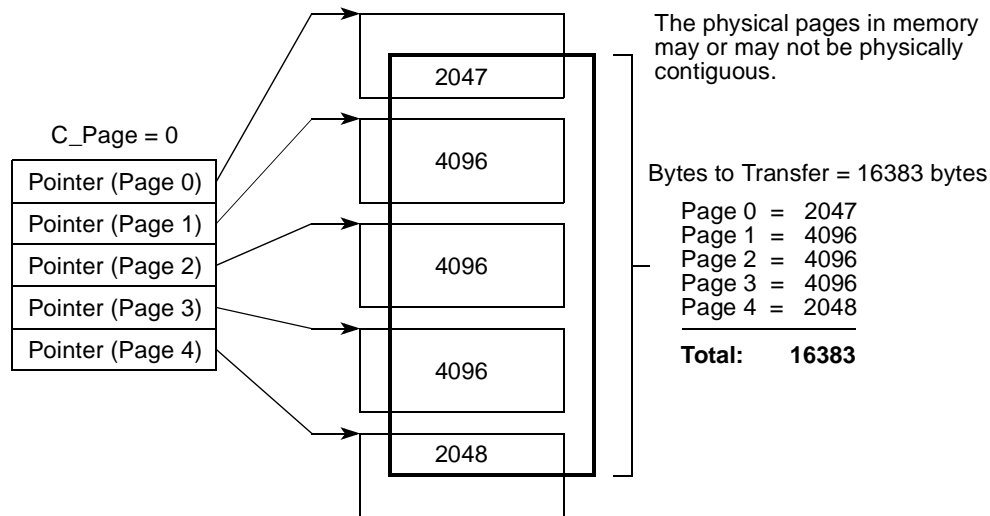


Figure 17-51. Example Mapping of qTD Buffer Pointers to Buffer Pages

The buffer pointer list in the qTD is long enough to support a maximum transfer size of 20K bytes. This case occurs when all five buffer pointers are used and the first offset is zero. A qTD handles a 16Kbyte buffer with any starting buffer alignment.

The host controller uses the `C_Page` field as an index value to determine which buffer pointer in the list should be used to start the current transaction. The host controller uses a different buffer pointer for each physical page of the buffer. This is always true, even if the buffer is physically contiguous.

The host controller must detect when the current transaction spans a page boundary and automatically move to the next available buffer pointer in the page pointer list. The next available pointer is reached by incrementing `C_Page` and pulling the next page pointer from the list. Software must ensure there are sufficient buffer pointers to move the amount of data specified in the Bytes to Transfer field.

Figure 17-51 illustrates a nominal example of how System software would initialize the buffer pointers list and the `C_Page` field for a transfer size of 16383 bytes. `C_Page` is cleared. The upper 20-bits of Page 0 references the start of the physical page. Current Offset (the lower 12-bits of queue head Dword 7) holds the offset in the page for example, 2049 (for example, 4096-2047). The remaining page pointers are set to reference the beginning of each subsequent 4K page.

For the first transaction on the qTD (assuming a 512-byte transaction), the host controller uses the first buffer pointer (page 0 because `C_Page` is cleared) and concatenates the Current Offset field. The 512 bytes are moved during the transaction, the Current Offset and Total Bytes to Transfer are adjusted by 512 and written back to the queue head working area.

During the 4th transaction, the host controller needs 511 bytes in page 0 and one byte in page 1. The host controller will increment `C_Page` (to 1) and use the page 1 pointer to move the final byte of the transaction. After the 4th transaction, the active page pointer is the page 1 pointer and Current Offset has rolled to one, and both are written back to the overlay area. The transactions continue for the rest of the buffer, with the

host controller automatically moving to the next page pointer (that is, C_Page) when necessary. There are three conditions for how the host controller handles C_Page.

- The current transaction does not span a page boundary. The value of C_Page is not adjusted by the host controller.
- The current transaction does span a page boundary. The host controller must detect the page cross condition and advance to the next buffer while streaming data to/from the USB.
- The current transaction completes on a page boundary (that is, the last byte moved for the current transaction is the last byte in the page for the current page pointer). The host controller must increment C_Page before writing back status for the transaction.

Note that the only valid adjustment the host controller may make to C_Page is to increment by one.

17.6.10.2 Adding Interrupt Queue Heads to the Periodic Schedule

The link path(s) from the periodic frame list to a queue head establishes in which frames a transaction can be executed for the queue head. Queue heads are linked into the periodic schedule so they are polled at the appropriate rate. System software sets a bit in a queue head's S-Mask to indicate which microframe within a 1 millisecond period a transaction should be executed for the queue head. Software must ensure that all queue heads in the periodic schedule have S-Mask set to a non-zero value. An S-mask with a zero value in the context of the periodic schedule yields undefined results.

If the desired poll rate is greater than one frame, system software can use a combination of queue head linking and S-Mask values to spread interrupts of equal poll rates through the schedule so that the periodic bandwidth is allocated and managed in the most efficient manner possible. Some examples are illustrated in [Table 17-67](#).

Table 17-67. Example Periodic Reference Patterns for Interrupt Transfers

| Frame # Reference Sequence | Description |
|-------------------------------------|--|
| 0, 2, 4, 6, 8, ... S-Mask = 0x01 | A queue head for the bInterval of 2 milliseconds (16 microframes) is linked into the periodic schedule so that it is reachable from the periodic frame list locations indicated in the previous column. In addition, the S-Mask field in the queue head is set to 0x01, indicating that the transaction for the endpoint should be executed on the bus during microframe 0 of the frame. |
| 0, 2, 4, 6, 8, ... S-Mask = 0x02 | Another example of a queue head with a bInterval of 2 milliseconds is linked into the periodic frame list at exactly the same interval as the previous example. However, the S-Mask is set to 0x02 indicating that the transaction for the endpoint should be executed on the bus during microframe 1 of the frame. |

17.6.10.3 Managing Transfer Complete Interrupts from Queue Heads

The host controller sets an interrupt to be signaled at the next interrupt threshold when the completed transfer (qTD) has an Interrupt on Complete (IOC) bit set, or whenever a transfer (qTD) completes with a short packet. If system software needs multiple qTDs to complete a client request (that is, like a control transfer) the intermediate qTDs do not require interrupts. System software may only need a single interrupt to notify it that the complete buffer has been transferred. System software may set IOC's to occur more frequently. A motivation for this may be that it wants early notification so that interface data structures can be re-used in a timely manner.

17.6.11 Ping Control

USB 2.0 defines an addition to the protocol for high-speed devices called Ping. Ping is required for all USB 2.0 High-speed bulk and control endpoints. Ping is not allowed for a split-transaction stream. This extension to the protocol eliminates the bad side-effects of Naking OUT endpoints. The Status field has a Ping State bit, which the host controller uses to determine the next actual PID it will use in the next transaction to the endpoint (see [Table 17-56](#)). The Ping State bit is only managed by the host controller for queue heads that meet all of the following criteria:

- The queue head is not an interrupt
- The EPS field equals High-Speed
- The PIDCode field equals OUT

[Table 17-68](#) illustrates the state transition table for the host controller's responsibility for maintaining the PING protocol. Refer to Chapter 8 in the *USB Specification, Revision 2.0* for detailed description on the Ping protocol.

Table 17-68. Ping Control State Transition Table

| Current | Event | | Next |
|---------|-------|----------------------|----------------------|
| | Host | Device | |
| Do Ping | PING | Nak | Do Ping |
| Do Ping | PING | Ack | Do OUT |
| Do Ping | PING | XactErr ¹ | Do Ping |
| Do Ping | PING | Stall | N/C ² |
| Do OUT | OUT | Nak | Do Ping |
| Do OUT | OUT | Nyet | Do Ping ³ |
| Do OUT | OUT | Ack | Do OUT |
| Do OUT | OUT | XactErr ¹ | Do Ping |
| Do OUT | OUT | Stall | N/C ² |

¹ Transaction Error (XactErr) is any time the host misses the handshake.

² No transition change required for the Ping State bit. The Stall handshake results in the endpoint being halted (for example, Active cleared and Halt set). Software intervention is required to restart queue.

³ A Nyet response to an OUT means that the device has accepted the data, but cannot receive any more at this time. Host must advance the transfer state and additionally, transition the Ping State bit to Do Ping.

The Ping State bit is described in [Table 17-56](#). The defined ping protocol allows the host to be imprecise on the initialization of the ping protocol (that is, start in Do OUT when we don't know whether there is space on the device or not). The host controller manages the Ping State bit. System software sets the initial value in the queue head when it initializes a queue head. The host controller preserves the Ping State bit across all queue advancements. This means that when a new qTD is written into the queue head overlay area, the previous value of the Ping State bit is preserved.

17.6.12 Split Transactions

USB 2.0 defines extensions to the bus protocol for managing USB 1.x data streams through USB 2.0 hubs. This section describes how the host controller uses the interface data structures to manage data streams with full- and low-speed devices, connected below a USB 2.0 hub, utilizing the split transaction protocol. Refer to the USB 2.0 Specification for the complete definition of the split transaction protocol. Full- and low-speed devices are enumerated identically as high-speed devices, but the transactions to the full- and low-speed endpoints use the split-transaction protocol on the high-speed bus. The split transaction protocol is an encapsulation of (or wrapper around) the full- or low-speed transaction. The high-speed wrapper portion of the protocol is addressed to the USB 2.0 hub and transaction translator below which the full- or low-speed device is attached.

EHCI uses dedicated data structures for managing full-speed isochronous data streams. Control, Bulk and Interrupt are managed using the queuing data structures. The interface data structures need to be programmed with the device address and the transaction translator number of the USB 2.0 hub operating as the low-/full-speed host controller for this link. The following sections describe the details of how the host controller processes and manages the split transaction protocol.

17.6.12.1 Split Transactions for Asynchronous Transfers

A queue head in the asynchronous schedule with an EPS field indicating a full-or low-speed device indicates to the host controller that it must use split transactions to stream data for this queue head. All full-speed bulk and full-, low-speed control are managed via queue heads in the asynchronous schedule.

Software must initialize the queue head with the appropriate device address and port number for the transaction translator that is serving as the full-/low-speed host controller for the links connecting the endpoint. Software must also initialize the split transaction state bit (SplitXState) to Do-Start-Split. Finally, if the endpoint is a control endpoint, then system software must set the Control Transfer Type (C) bit in the queue head to a one. If this is not a control transfer type endpoint, the C bit must be initialized by software to be a zero. This information is used by the host controller to properly set the Endpoint Type (ET) field in the split transaction bus token. When the C bit is a zero, the split transaction token's ET field is set to indicate a bulk endpoint. When the C bit is a one, the split transaction token's ET field is set to indicate a control endpoint. Refer to Chapter 8 of *USB Specification, Revision 2.0* for details.




Figure 17-52. Host Controller Asynchronous Schedule Split-Transaction State Machine

17.6.12.1.1 Asynchronous—Do-Start-Split

Do-Start-Split is the state which software must initialize a full- or low-speed asynchronous queue head. This state is entered from the Do-Complete-Split state only after a complete-split transaction receives a valid response from the transaction translator that is not a Nyet handshake.

For queue heads in this state, the host controller executes a start-split transaction to the transaction translator. If the bus transaction completes without an error and PID Code indicates an IN or OUT transaction, then the host controller reloads the error counter (Cerr). If it is a successful bus transaction and the PID Code indicates a SETUP, the host controller will not reload the error counter. If the transaction translator responds with a Nak, the queue head is left in this state, and the host controller proceeds to the next queue head in the asynchronous schedule.

If the host controller times out the transaction (no response, or bad response) the host controller decrements Cerr and proceeds to the next queue head in the asynchronous schedule.

17.6.12.1.2 Asynchronous—Do-Complete-Split

This state is entered from the Do-Start-Split state only after a start-split transaction receives an Ack handshake from the transaction translator.

For queue heads in this state, the host controller executes a complete-split transaction to the transaction translator. If the transaction translator responds with a Nyet handshake, the queue head is left in this state, the error counter is reset and the host controller proceeds to the next queue head in the asynchronous schedule. When a Nyet handshake is received for a bus transaction where the queue head's PID Code indicates an IN or OUT, the host controller reloads the error counter (Cerr). When a Nyet handshake is received for a complete-split bus transaction where the queue head's PID Code indicates a SETUP, the host controller must not adjust the value of Cerr.

Independent of PID Code, the following responses have the indicated effects:

- Transaction Error (XactErr). Timeout/data CRC failure. The error counter (Cerr) is decremented by one and the complete split transaction is immediately retried (if possible). If there is not enough time in the microframe to execute the retry, the host controller ensures that the next time the host controller begins executing from the Asynchronous schedule, it must begin executing from this queue head. If another start-split (for some other endpoint) is sent to the transaction translator before the complete-split is really completed, the transaction translator could dump the results (which were never delivered to the host). This is why the core specification states the retries must be immediate. When the host controller returns to the asynchronous schedule in the next microframe, the first transaction from the schedule will be the retry for this endpoint. If Cerr went to zero, the host controller halts the queue.
- NAK. The target endpoint Nak'd the full- or low-speed transaction. The state of the transfer is not advanced and the state is exited. If the PID Code is a SETUP, then the Nak response is a protocol error. The XactErr status bit is set and the Cerr field is decremented.
- STALL. The target endpoint responded with a STALL handshake. The host controller sets the halt bit in the status byte, retires the qTD but does not attempt to advance the queue.

If the PID Code indicates an IN, then any of following responses are expected:

- **DATA0/1.** On reception of data, the host controller ensures the PID matches the expected data toggle and checks CRC. If the packet is good, the host controller advances the state of the transfer (for example, moves the data pointer by the number of bytes received, decrements the BytesToTransfer field by the number of bytes received, and toggles the dt bit). The host controller then exits this state. The response and advancement of transfer may trigger other processing events, such as retirement of the qTD and advancement of the queue.

If the data sequence PID does not match the expected, the data is ignored, the transfer state is not advanced and this state is exited.

If the PID Code indicates an OUT/SETUP, then any of following responses are expected:

- **ACK.** The target endpoint accepted the data, so the host controller must advance the state of the transfer. The Current Offset field is incremented by Maximum Packet Length or Bytes to Transfer, whichever is less. The Bytes To Transfer field is decremented by the same amount and the data toggle bit (dt) is toggled. The host controller then exits this state.

Advancing the transfer state may cause other processing events such as retirement of the qTD and advancement of the queue.

17.6.12.2 Split Transaction Interrupt

Split-transaction Interrupt-IN/OUT endpoints are managed using the same data structures used for high-speed interrupt endpoints. They both co-exist in the periodic schedule. Queue heads/qTDs offer the set of features required for reliable data delivery, which is characteristic to interrupt transfer types. The split-transaction protocol is managed completely within this defined functional transfer framework. For example, for a high-speed endpoint, the host controller will visit a queue head, execute a high-speed transaction (if criteria are met) and advance the transfer state (or not) depending on the results of the entire transaction. For low- and full-speed endpoints, the details of the execution phase are different (that is, takes more than one bus transaction to complete), but the remainder of the operational framework is intact.

17.6.12.2.1 Split Transaction Scheduling Mechanisms for Interrupt

Full- and low-speed Interrupt queue heads have an EPS field indicating full- or low-speed and have a non-zero S-mask field. The host controller can detect this combination of parameters and assume the endpoint is a periodic endpoint. Low- and full-speed interrupt queue heads require the use of the split transaction protocol. The host controller sets the Endpoint Type (ET) field in the split token to indicate the transaction is an interrupt. These transactions are managed through a transaction translator's periodic pipeline. Software should not set these fields to indicate the queue head is an interrupt unless the queue head is used in the periodic schedule.

System software manages the per/transaction translator periodic pipeline by budgeting and scheduling exactly during which microframes the start-splits and complete-splits for each endpoint will occur. The characteristics of the transaction translator are such that the high-speed transaction protocol must execute during explicit microframes, or the data or response information in the pipeline is lost. [Figure 17-53](#) illustrates the general scheduling boundary conditions that are supported by the EHCI periodic schedule

and queue head data structure. The S and C_n labels indicate microframes where software can schedule start-splits and complete splits (respectively).




Figure 17-53. Split Transaction, Interrupt Scheduling Boundary Conditions

The scheduling cases are:

- Case 1: The normal scheduling case is where the entire split transaction is completely bounded by a frame (H-Frame in this case).
- Case 2a through Case 2c: The USB 2.0 hub pipeline rules states clearly, when and how many complete-splits must be scheduled to account for earliest to latest execution on the full/low-speed link. The complete-splits may span the H-Frame boundary when the start-split is in microframe 4 or later. When this occurs, the H-Frame to B-Frame alignment requires that the queue head be reachable from consecutive periodic frame list locations. System software cannot build an efficient schedule that satisfies this requirement unless it uses FSTNs.

Figure 17-54 illustrates the general layout of the periodic schedule.

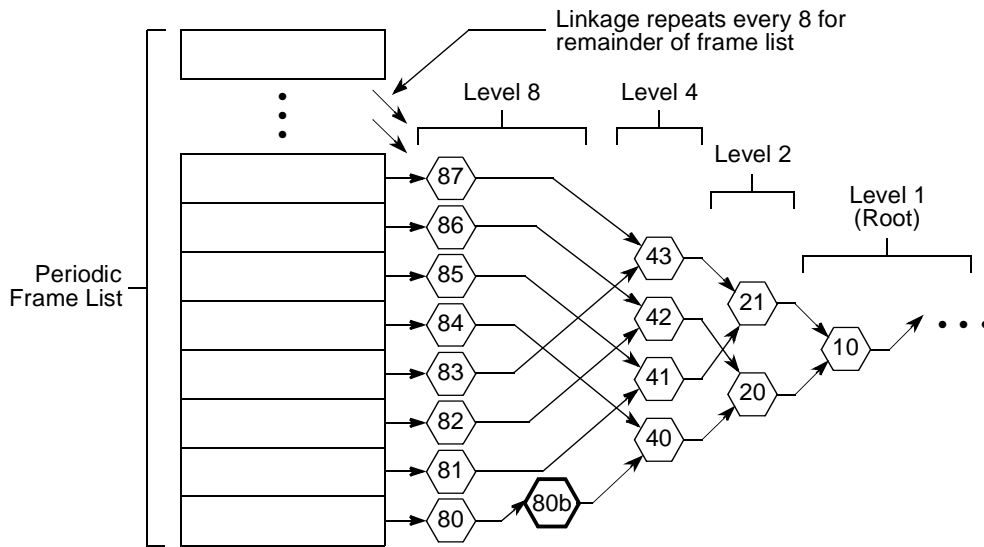


Figure 17-54. General Structure of EHCI Periodic Schedule Utilizing Interrupt Spreading

The periodic frame list is effectively the leaf level a binary tree, which is always traversed leaf to root. Each level in the tree corresponds to a 2^N poll rate. Software can efficiently manage periodic bandwidth on the USB by spreading interrupt queue heads that have the same poll rate requirement across all the available paths from the frame list. For example, system software can schedule eight poll rate 8 queue heads and account for them once in the high-speed bus bandwidth allocation.

When an endpoint is allocated an execution footprint that spans a frame boundary, the queue head for the endpoint must be reachable from consecutive locations in the frame list. An example would be if 8_{0b} were such an endpoint. Without additional support on the interface, to get 8_{0b} reachable at the correct time, software would have to link 8_1 to 8_{0b} . It would then have to move 4_1 and everything linked after into the same path as 4_0 . This upsets the integrity of the binary tree and disallows the use of the spreading technique.

FSTN data structures are used to preserve the integrity of the binary-tree structure and enable the use of the spreading technique. Section 17.5.7, “Periodic Frame Span Traversal Node (FSTN),” defines the hardware and software operational model requirements for using FSTNs.

The following queue head fields are initialized by system software to instruct the host controller when to execute portions of the split-transaction protocol.

- SplitXState. This is a single bit residing in the Status field of a queue head (Table 17-56). This bit is used to track the current state of the split transaction.
- Frame S-mask. This is a bit-field where-in system software sets a bit corresponding to the microframe (within an H-Frame) that the host controller should execute a start-split transaction. This is always qualified by the value of the SplitXState bit in the Status field of the queue head. For example, referring to Figure 17-53, case one, the S-mask would have a value of 0b0000_0001 indicating that if the queue head is traversed by the host controller, and the SplitXState indicates Do_Start, and the current microframe as indicated by FRINDEX[2-0] is 0, then execute a start-split transaction.

- **Frame C-mask.** This is a bit-field where system software sets one or more bits corresponding to the microframes (within an H-Frame) that the host controller should execute complete-split transactions. The interpretation of this field is always qualified by the value of the SplitXState bit in the Status field of the queue head. For example, referring to [Figure 17-53](#), case one, the C-mask would have a value of 0b0001_1100 indicating that if the queue head is traversed by the host controller, and the SplitXState indicates Do_Complete, and the current microframe as indicated by FRINDEX[2-0] is 2, 3, or 4, then execute a complete-split transaction. It is software's responsibility to ensure that the translation between H-Frames and B-Frames is correctly performed when setting bits in S-mask and C-mask.

17.6.12.2.2 Host Controller Operational Model for FSTNs

The FSTN data structure is used to manage Low/Full-speed interrupt queue heads that need to be reached from consecutive frame list locations (that is, boundary cases 2a through 2c). An FSTN is essentially a back pointer, similar in intent to the back pointer field in the siTD data structure.

This feature provides software a simple primitive to save a schedule position, redirect the host controller to traverse the necessary queue heads in the previous frame, then restore the original schedule position and complete normal traversal.

There are four components to the use of FSTNs:

- FSTN data structure, defined in [Section 17.5.7, “Periodic Frame Span Traversal Node \(FSTN\).”](#)
- A Save Place indicator; this is always an FSTN with its Back Path Link Pointer[T] bit cleared.
- A Restore indicator; this is always an FSTN with its Back Path Link Pointer[T] bit set.
- Host controller FSTN traversal rules.

When the host controller encounters an FSTN during microframes 2 through 7 it simply follows the node's Normal Path Link Pointer to access the next schedule data structure. Note that the FSTN's Normal Path Link Pointer[T] bit may set, which the host controller must interpret as the end of periodic list mark.

When the host controller encounters a Save-Place FSTN in microframes 0 or 1, it saves the value of the Normal Path Link Pointer and sets an internal flag indicating that it is executing in Recovery Path mode. Recovery Path mode modifies the host controller's rules for how it traverses the schedule and limits which data structures are considered for execution of bus transactions. The host controller continues executing in Recovery Path mode until it encounters a Restore FSTN or it determines that it has reached the end of the microframe.

The rules for schedule traversal and limited execution while in Recovery Path mode are:

- Always follow the Normal Path Link Pointer when it encounters an FSTN that is a Save-Place indicator. The host controller must not recursively follow Save-Place FSTNs. Therefore, while executing in Recovery Path mode, it must never follow an FSTN's Back Path Link Pointer.
- Do not process an siTD or iTD data structure; simply follow its Next Link Pointer.
- Do not process a QH (Queue Head) whose EPS field indicates a high-speed device; simply follow its Horizontal Link Pointer.
- When a QH's EPS field indicates a Full/Low-speed device, the host controller only considers it for execution if its SplitXState is DoComplete (note: this applies whether the PID Code indicates an

IN or an OUT). Refer to the *EHCI Specification* for a complete list of additional conditions that must be met in general for the host controller to issue a bus transaction. Note that the host controller must not execute a Start-split transaction while executing in Recovery Path mode. Refer to the *EHCI Specification* for special handling when in Recovery Path mode.

- Stop traversing the recovery path when it encounters an FSTN that is a Restore indicator. The host controller unconditionally uses the saved value of the Save-Place FSTN's Normal Path Link Pointer when returning to the normal path traversal. The host controller must clear the context of executing a Recovery Path when it restores schedule traversal to the Save-Place FSTN's Normal Path Link Pointer.

If the host controller determines that there is not enough time left in the microframe to complete processing of the periodic schedule, it abandons traversal of the recovery path, and clears the context of executing a recovery path. The result is that at the start of the next consecutive microframe, the host controller starts traversal at the frame list.

An example traversal of a periodic schedule that includes FSTNs is illustrated in [Figure 17-55](#).

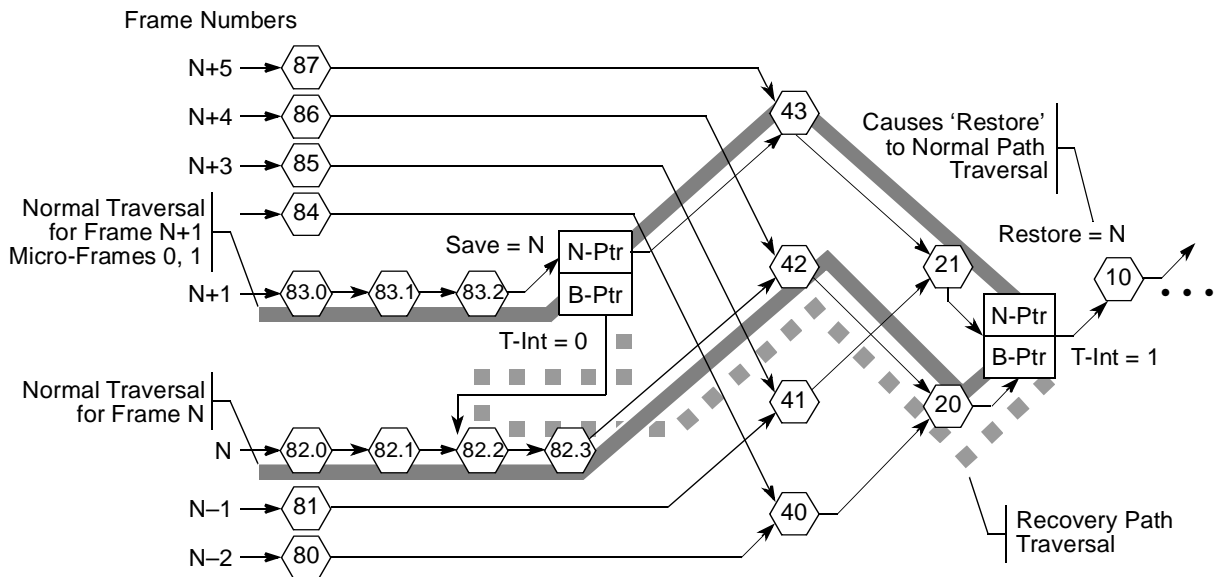


Figure 17-55. Example Host Controller Traversal of Recovery Path via FSTNs

In frame N (microframes 0-7), for this example, the host controller traverses all of the schedule data structures utilizing the Normal Path Link Pointers in any FSTNs it encounters. This is because the host controller has not yet encountered a Save-Place FSTN so it is not executing in Recovery Path mode. When it encounters the Restore FSTN, (Restore-N), during microframes 0 and 1, it uses Restore-N. Normal Path Link Pointer to traverse to the next data structure (that is, normal schedule traversal). This is because the host controller must use a Restore FSTN's Normal Path Link Pointer when not executing in a Recovery-Path mode. The nodes traversed during frame N include: {82.0, 82.1, 82.2, 82.3, 42, 20, Restore-N, 10 ... }.

In frame N+1 (microframes 0 and 1), when the host controller encounters Save-Path FSTN (Save-N), it observes that Save-N.Back Path Link Pointer.T-bit is zero (definition of a Save-Path indicator). The host controller saves the value of Save-N. Normal Path Link Pointer and follows Save-N.Back Path Link

Pointer. At the same time, it sets an internal flag indicating that it is now in Recovery Path mode (the recovery path is annotated in Figure 17-55 with a large dashed line). The host controller continues traversing data structures on the recovery path and executing only those bus transactions as noted above, on the recovery path until it reaches Restore FSTN (Restore-N). Restore-N.Back Path Link Pointer.T-bit is set (definition of a Restore indicator), so the host controller exits Recovery Path mode by clearing the internal Recovery Path mode flag and commences (restores) schedule traversal using the saved value of the Save-Place FSTN's Normal Path Link Pointer (for example, Save-N.Normal Path Link Pointer). The nodes traversed during these microframes include: {8_{3,0}, 8_{3,1}, 8_{3,2}, Save-A, 8_{2,2}, 8_{2,3}, 4₂, 2₀, Restore-N, 4₃, 2₁, Restore-N, 10 ...}.

In frame N+1 (microframes 2-7), when the host controller encounters Save-Path FSTN Save-N, it unconditionally follows Save-N.Normal Path Link Pointer. The nodes traversed during these microframes include: {8_{3,0}, 8_{3,1}, 8_{3,2}, Save-A, 4₃, 2₁, Restore-N, 1₀ ...}.

17.6.12.2.3 Software Operational Model for FSTNs

Software must create a consistent, coherent schedule for the host controller to traverse. When using FSTNs, system software must adhere to the following rules:

- Each Save-Place indicator requires a matching Restore indicator.
The Save-Place indicator is an FSTN with a valid Back Path Link Pointer and T-bit equal to zero. Note that Back Path Link Pointer[Typ] field must be set to indicate the referenced data structure is a queue head. The Restore indicator is an FSTN with its Back Path Link Pointer[T] bit set.
A Restore FSTN may be matched to one or more Save-Place FSTNs. For example, if the schedule includes a poll-rate 1 level, then system software only needs to place a Restore FSTN at the beginning of this list in order to match all possible Save-Place FSTNs.
- If the schedule does not have elements linked at a poll-rate level of one, and one or more Save-Place FSTNs are used, then System Software must ensure the Restore FSTN's Normal Path Link Pointer's T-bit is set, as this is used to mark the end of the periodic list.
- When the schedule does have elements linked at a poll rate level of one, a Restore FSTN must be the first data structure on the poll rate one list. All traversal paths from the frame list converge on the poll-rate one list. System software must ensure that Recovery Path mode is exited before the host controller is allowed to traverse the poll rate level one list.
- A Save-Place FSTN's Back Path Link Pointer must reference a queue head data structure. The referenced queue head must be reachable from the previous frame list location. In other words, if the Save-Place FSTN is reachable from frame list offset N, then the FSTN's Back Path Link Pointer must reference a queue head that is reachable from frame list offset N-1.

Software should make the schedule as efficient as possible. What this means in this context is that software should have no more than one Save-Place FSTN reachable in any single frame. Note there are times when two (or more, depending on the implementation) could exist as full-/low-speed footprints change with bandwidth adjustments. This could occur, for example when a bandwidth rebalance causes system software to move the Save-Place FSTN from one poll rate level to another. During the transition, software must preserve the integrity of the previous schedule until the new schedule is in place.

17.6.12.2.4 Tracking Split Transaction Progress for Interrupt Transfers

To correctly maintain the data stream, the host controller must be able to detect and report errors where data is lost. For interrupt-IN transfers, data is lost when it makes it into the USB 2.0 hub, but the USB 2.0 host system is unable to get it from the USB 2.0 hub and into the system before it expires from the transaction translator pipeline. When a lost data condition is detected, the queue is halted, thus signaling system software to recover from the error. A data-loss condition exists whenever a start-split is issued, accepted and successfully executed by the USB 2.0 hub, but the complete-splits get unrecoverable errors on the high-speed link, or the complete-splits do not occur at the correct times. One reason complete-splits might not occur at the right time would be due to host-induced system hold-offs that cause the host controller to miss bus transactions because it cannot get timely access to the schedule in system memory.

The same condition can occur for an interrupt-OUT, but the result is not an endpoint halt condition, but rather effects only the progress of the transfer. The queue head has the following fields to track the progress of each split transaction. These fields are used to keep incremental state about which (and when) portions have been executed.

- **C-prog-mask.** This is an eight-bit bit-vector where the host controller keeps track of which complete-splits have been executed. Due to the nature of the transaction translator periodic pipeline, the complete-splits need to be executed in-order. The host controller needs to detect when the complete-splits have not been executed in order. This can only occur due to system hold-offs where the host controller cannot get to the memory-based schedule. C-prog-mask is a simple bit-vector that the host controller sets one of the C-prog-mask bits for each complete-split executed. The bit position is determined by the microframe number in which the complete-split was executed. The host controller always checks C-prog-mask before executing a complete-split transaction. If the previous complete-splits have not been executed then it means one (or more) have been skipped and data has potentially been lost.
- **FrameTag.** This field is used by the host controller during the complete-split portion of the split transaction to tag the queue head with the frame number (H-Frame number) when the next complete split must be executed.
- **S-bytes.** This field can be used to store the number of data payload bytes sent during the start-split (if the transaction was an OUT). The S-bytes field must be used to accumulate the data payload bytes received during the complete-splits (for an IN).

17.6.12.2.5 Split Transaction Execution State Machine for Interrupt

In the following section, all references to microframe are in the context of a microframe within an H-Frame.

As with asynchronous Full- and Low-speed endpoints, a split-transaction state machine is used to manage the split transaction sequence. Aside from the fields defined in the queue head for scheduling and tracking the split transaction, the host controller calculates one internal mechanism that is also used to manage the split transaction. The internal calculated mechanism is:

- **cMicroFrameBit.** This is a single-bit encoding of the current microframe number. It is an eight-bit value calculated by the host controller at the beginning of every microframe. It is calculated from the three least significant bits of the FRINDEX register (that is, $cMicroFrameBit = (1 \text{ shifted-left}(FRINDEX[2-0]))$). The cMicroFrameBit has at most one bit asserted, which always

corresponds to the current microframe number. For example, if the current microframe is 0, then `cMicroFrameBit` will equal `0b0000_0001`.

The variable `cMicroFrameBit` is used to compare against the S-mask and C-mask fields to determine whether the queue head is marked for a start- or complete-split transaction for the current microframe.

Figure 17-56 illustrates how a complete interrupt split transaction is managed. There are two phases to each split transaction. The first is a single start-split transaction, which occurs when the `SplitXState` is at `Do_Start` and the single bit in `cMicroFrameBit` has a corresponding bit active in `QH[S-mask]`. The transaction translator does not acknowledge the receipt of the periodic start-split, so the host controller unconditionally transitions the state to `Do_Complete`. Due to the available jitter in the transaction translator pipeline, there is more than one complete-split transaction scheduled by software for the `Do_Complete` state. This translates simply to the fact that there are multiple bits set in the `QH[C-mask]` field.

The host controller keeps the queue head in the `Do_Complete` state until the split transaction is complete (see definition below), or an error condition triggers the three-strikes-rule (for example, after the host tries the same transaction three times, and each encounters an error, the host controller stops retrying the bus transaction and halts the endpoint, thus requiring system software to detect the condition and perform system-dependent recovery).




Figure 17-56. Split Transaction State Machine for Interrupt

17.6.12.2.6 Periodic Interrupt—Do-Start-Split

This is the state software must initialize a full- or low-speed interrupt queue head `StartXState` bit. This state is entered from the `Do_Complete Split` state only after the split transaction is complete. This occurs when

one of the following events occur: The transaction translator responds to a complete-split transaction with one of the following:

- **NAK.** A NAK response is a propagation of the full- or low-speed endpoint's NAK response.
- **ACK.** An ACK response is a propagation of the full- or low-speed endpoint's ACK response. Only occurs on an OUT endpoint.
- **DATA 0/1.** Only occurs for INs. Indicates that this is the last of the data from the endpoint for this split transaction.
- **ERR.** The transaction on the low-/full-speed link below the transaction translator had a failure (for example, timeout, bad CRC, and so on).
- **NYET (and Last).** The host controller issued the last complete-split and the transaction translator responded with a NYET handshake. This means that the start-split was not correctly received by the transaction translator, so it never executed a transaction to the full- or low-speed endpoint, see [Section 17.6.12.2.7, “Periodic Interrupt—Do-Complete-Split,”](#) for the definition of 'Last'.

Each time the host controller visits a queue head in this state (once within the Execute Transaction state), bit-wise ANDs QH[S-mask] with cMicroFrameBit to determine whether to execute a start-split. If the result is non-zero, then the host controller issues a start-split transaction. If the PID Code field indicates an IN transaction, the host controller must zero-out the QH[S-bytes] field. After the split-transaction has been executed, the host controller sets up state in the queue head to track the progress of the complete-split phase of the split transaction. Specifically, it records the expected frame number into QH[FrameTag] field, sets C-prog-mask to zero (0x00), and exits this state. Note that the host controller must not adjust the value of Cerr as a result of completion of a start-split transaction.

17.6.12.2.7 Periodic Interrupt—Do-Complete-Split

This state is entered unconditionally from the Do Start Split state after a start-split transaction is executed on the bus. Each time the host controller visits a queue head in this state (once within the Execute Transaction state), it checks to determine whether a complete-split transaction should be executed now.

There are four tests to determine whether a complete-split transaction should be executed.

- **Test A.** cMicroFrameBit is bit-wise ANDed with QH[C-mask] field. A non-zero result indicates that software scheduled a complete-split for this endpoint, during this microframe.
- **Test B.** QH[FrameTag] is compared with the current contents of FRINDEX[7–3]. An equal indicates a match.
- **Test C.** The complete-split progress bit vector is checked to determine whether the previous bit is set, indicating that the previous complete-split was appropriately executed. An example algorithm for this test is provided below:

```

Algorithm Boolean CheckPreviousBit(QH.C-prog-mask, QH.C-mask, cMicroFrameBit)
Begin
-- Return values:
-- TRUE - no error
-- FALSE - error
--
Boolean rvalue = TRUE;
previousBit = cMicroframeBit logical-rotate-right(1)
-- Bit-wise anding previousBit with C-mask indicates
-- whether there was an intent

```



```

-- to send a complete split in the previous microframe. So,
-- if the
-- 'previous bit' is set in C-mask, check C-prog-mask to
-- make sure it
-- happened.
If (previousBit bitAND QH.C-mask)then
    If not(previousBit bitAND QH.C-prog-mask) then
        rvalue = FALSE;
    End if
End If
-- If the C-prog-mask already has a one in this bit position,
-- then an aliasing
-- error has occurred. It will probably get caught by the
-- FrameTag Test, but
-- at any rate it is an error condition that as detectable here
-- should not allow
-- a transaction to be executed.
If (cMicroFrameBit bitAND QH.C-prog-mask) then
    rvalue = FALSE;
End if
return (rvalue)
End Algorithm
    
```

- Test D. Check to see if a start-split should be executed in this microframe. Note this is the same test performed in the Do Start Split state. Whenever it evaluates to TRUE and the controller is NOT processing in the context of a Recovery Path mode, it means a start-split should occur in this microframe. Test D and Test A evaluating to TRUE at the same time is a system software error. Behavior is undefined.

If (A .and. B .and. C .and. not(D)) then the host controller will execute a complete-split transaction. When the host controller commits to executing the complete-split transaction, it updates QH[C-prog-mask] by bit-ORing with cMicroFrameBit. On completion of the complete-split transaction, the host controller records the result of the transaction in the queue head and sets QH[FrameTag] to the expected H-Frame number. The effect to the state of the queue head and thus the state of the transfer depends on the response by the transaction translator to the complete-split transaction. The following responses have the effects (note that any responses that result in decrementing of the Cerr will result in the queue head being halted by the host controller if the result of the decrement is zero):

- NYET (and Last)

On each NYET response, the host controller checks to determine whether this is the last complete-split for this split transaction. Last is defined in this context as the condition where all of the scheduled complete-splits have been executed. If it is the last complete-split (with a NYET response), then the transfer state of the queue head is not advanced (never received any data) and this state exited. The transaction translator must have responded to all the complete-splits with NYETs, meaning that the start-split issued by the host controller was not received. The start-split should be retried at the next poll period.

The test for whether this is the Last complete split can be performed by XOR QH[C-mask] with QH[C-prog-mask]. If the result is all zeros then all complete-splits have been executed. When this condition occurs, the XactErr status bit is set and the Cerr field is decremented.
- NYET (and not Last)

See above description for testing for Last. The complete-split transaction received a NYET response from the transaction translator. Do not update any transfer state (except for C-prog-mask and FrameTag) and stay in this state. The host controller must not adjust Cerr on this response.

- Transaction Error (XactErr). Timeout, data CRC failure, and so on.
The Cerr field is decremented and the XactErr bit in the Status field is set. The complete split transaction is immediately retried (if Cerr is non-zero). If there is not enough time in the microframe to complete the retry and the endpoint is an IN, or Cerr is decremented to a zero from a one, the queue is halted. If there is not enough time in the microframe to complete the retry and the endpoint is an OUT and Cerr is not zero, then this state is exited (that is, return to Do Start Split). This results in a retry of the entire OUT split transaction, at the next poll period. Refer to Chapter 11 Hubs (specifically the section on full- and low-speed interrupts) in the *USB Specification Revision 2.0* for detailed requirements on why these errors must be immediately retried.
- ACK
This can only occur if the target endpoint is an OUT. The target endpoint ACK'd the data and this response is a propagation of the endpoint ACK up to the host controller. The host controller must advance the state of the transfer. The Current Offset field is incremented by Maximum Packet Length or Bytes to Transfer, whichever is less. The field Bytes To Transfer is decremented by the same amount. And the data toggle bit (dt) is toggled. The host controller will then exit this state for this queue head. The host controller must reload Cerr with maximum value on this response. Advancing the transfer state may cause other process events such as retirement of the qTD and advancement of the queue.
- MDATA
This response will only occur for an IN endpoint. The transaction translator responded with zero or more bytes of data and an MDATA PID. The incremental number of bytes received is accumulated in QH[S-bytes]. The host controller must not adjust Cerr on this response.
- DATA0/1
This response may only occur for an IN endpoint. The number of bytes received is added to the accumulated byte count in QH[S-bytes]. The state of the transfer is advanced by the result and the host controller exits this state for this queue head.
Advancing the transfer state may cause other processing events such as retirement of the qTD and advancement of the queue.
If the data sequence PID does not match the expected, the entirety of the data received in this split transaction is ignored, the transfer state is not advanced and this state is exited.
- NAK
The target endpoint Nak'd the full- or low-speed transaction. The state of the transfer is not advanced, and this state is exited. The host controller must reload Cerr with maximum value on this response.
- ERR
There was an error during the full- or low-speed transaction. The ERR status bit is set, Cerr is decremented, the state of the transfer is not advanced, and this state is exited.
- STALL

The queue is halted (an exit condition of the Execute Transaction state). The status field bits: Active bit is cleared and the Halted bit is set and the qTD is retired. Responses which are not enumerated in the list or which are received out of sequence are illegal and may result in undefined host controller behavior. The other possible combinations of tests A, B, C, and D may indicate that data or response was lost. [Table 17-69](#) lists the possible combinations and the appropriate action.

Table 17-69. Interrupt IN/OUT Do Complete Split State Execution Criteria

| Condition | Action | Description |
|-----------------------|--|--|
| not(A) not(D) | Ignore QHD | Neither a start nor complete-split is scheduled for the current microframe. Host controller should continue walking the schedule. |
| A not(C) | If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split | Progress bit check failed. This means a complete-split has been missed. There is the possibility of lost data. If PID Code is an IN, then the queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Microframe bit in the status field to a one. |
| A not(B) C | If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split | QH.FrameTag test failed. This means that exactly one or more H-Frames have been skipped. This means complete-splits and have missed. There is the possibility of lost data. If PID Code is an IN, then the queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Microframe bit in the status field to a one. |
| A B C not(D) | Execute complete-split | This is the non-error case where the host controller executes a complete-split transaction. |
| D | If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split | This is a degenerate case where the start-split was issued, but all of the complete-splits were skipped and all possible intervening opportunities to detect the missed data failed to fire. If PID Code is an IN, then the Queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Microframe bit in the status field to a one. Note that when executing in the context of a Recovery Path mode, the host controller is allowed to process the queue head and take the actions indicated above, or it may wait until the queue head is visited in the normal processing mode. Regardless, the host controller must not execute a start-split in the context of a executing in a Recovery Path mode. |

17.6.12.2.8 Managing the QH[FrameTag] Field

The QH[FrameTag] field in a queue head is completely managed by the host controller. The rules for setting QH[FrameTag] are simple:

- Rule 1: If transitioning from Do Start Split to Do Complete Split and the current value of FRINDEX[2–0] is 6, QH[FrameTag] is set to $FRINDEX[7–3] + 1$. This accommodates split transactions whose start-split and complete-splits are in different H-Frames (case 2a, see [Figure 17-53](#)).
- Rule 2: If the current value of FRINDEX[2–0] is 7, QH[FrameTag] is set to $FRINDEX[7–3] + 1$. This accommodates staying in Do Complete Split for cases 2a, 2b, and 2c in [Figure 17-53](#).

- Rule 3: If transitioning from Do_Start Split to Do Complete Split and the current value of FRINDEX[2–0] is not 6, or currently in Do Complete Split and the current value of (FRINDEX[2–0]) is not 7, FrameTag is set to FRINDEX[7–3]. This accommodates all other cases in [Figure 17-53](#).

17.6.12.2.9 Rebalancing the Periodic Schedule

System software must occasionally adjust a periodic queue head's S-mask and C-mask fields during operation. This need occurs when adjustments to the periodic schedule create a new bandwidth budget and one or more queue head's are assigned new execution footprints (that is, new S-mask and C-mask values).

It is imperative that system software must not update these masks to new values in the midst of a split transaction. In order to avoid any race conditions with the update, the host controller provides a simple assist to system software. System software sets the Inactivate-on-next-Transaction (I) bit to signal the host controller that it intends to update the S-mask and C-mask on this queue head. System software then waits for the host controller to observe the I-bit is set and transitions the Active bit to a zero. The rules for how and when the host controller clears the Active bit are:

- If the Active bit is cleared, no action is taken. The host controller does not attempt to advance the queue when the I-bit is set.
- If the Active bit is set and the SplitXState is DoStart (regardless of the value of S-mask), the host controller simply clears the Active bit. The host controller is not required to write the transfer state back to the current qTD. Note that if the S-mask indicates that a start-split is scheduled for the current microframe, the host controller must not issue the start-split bus transaction; it must clear the Active bit.

System software must save transfer state before setting the I-bit. This is required so that it can correctly determine what transfer progress (if any) occurred after the I-bit was set and the host controller executed its final bus-transaction and cleared the Active bit.

After system software has updated the S-mask and C-mask, it must then reactivate the queue head. Since the Active bit and the I-bit cannot be updated with the same write, system software needs to use the following algorithm to coherently re-activate a queue head that has been stopped using the I-bit.

1. Set the Halted bit, then
2. Clear the I-bit, then
3. Set the Active bit and clear the Halted bit in the same write.

Setting the Halted bit inhibits the host controller from attempting to advance the queue between the time the I-bit is cleared and the Active bit is set.

17.6.12.3 Split Transaction Isochronous

Full-speed isochronous transfers are managed using the split-transaction protocol through a USB 2.0 transaction translator in a USB 2.0 hub. The host controller utilizes siTD data structure to support the special requirements of isochronous split-transactions. This data structure uses the scheduling model of isochronous TDs (see [Section 17.6.8, “Managing Isochronous Transfers Using iTDs,”](#) for the operational model of iTDs) with the contiguous data feature provided by queue heads. This simple arrangement allows

a single isochronous scheduling model and adds the additional feature that all data received from the endpoint (per split transaction) must land into a contiguous buffer.

17.6.12.3.1 Split Transaction Scheduling Mechanisms for Isochronous

Full-speed isochronous transactions are managed through a transaction translator's periodic pipeline. As with full- and low-speed interrupt, system software manages each transaction translator's periodic pipeline by budgeting and scheduling exactly during which microframes the start-splits and complete-splits for each full-speed isochronous endpoint occur. The requirements described in [Section 17.6.12.2.1, “Split Transaction Scheduling Mechanisms for Interrupt,”](#) apply.

[Figure 17-57](#) illustrates the general scheduling boundary conditions that are supported by the EHCI periodic schedule. The S_n and C_n labels indicate microframes where software can schedule start- and complete-splits (respectively). The H-Frame boundaries are marked with a large, solid bold vertical line.

The B-Frame boundaries are marked with a large, bold, dashed line. The bottom of Figure 17-57 illustrates the relationship of an siTD to the H-Frame.




Figure 17-57. Split Transaction, Isochronous Scheduling Boundary Conditions

When the endpoint is an isochronous OUT, there are only start-splits, and no complete-splits. When the endpoint is an isochronous IN, there is at most one start-split and one to N complete-splits. The scheduling boundary cases are:

- Case 1: The entire split transaction is completely bounded by an H-Frame. For example, the start-splits and complete-splits are all scheduled to occur in the same H-Frame.
- Case 2a: This boundary case is where one or more (at most two) complete-splits of a split transaction IN are scheduled across an H-Frame boundary. This can only occur when the split transaction has the possibility of moving data in B-Frame, microframes 6 or 7 (H-Frame microframe 7 or 0). When an H-Frame boundary wrap condition occurs, the scheduling of the split

transaction spans more than one location in the periodic list.(for example, it takes two siTDs in adjacent periodic frame list locations to fully describe the scheduling for the split transaction).

Although the scheduling of the split transaction may take two data structures, all of the complete-splits for each full-speed IN isochronous transaction must use only one data pointer. For this reason, siTDs contain a back pointer.

Software must never schedule full-speed isochronous OUTs across an H-Frame boundary.

- Case 2b: This case can only occur for a very large isochronous IN. It is the only allowed scenario where a start-split and complete-split for the same endpoint can occur in the same microframe. Software must enforce this rule by scheduling the large transaction first. Large is defined to be anything larger than 579 byte maximum packet size.

A subset of the same mechanisms employed by full- and low-speed interrupt queue heads are employed in siTDs to schedule and track the portions of isochronous split transactions. The following fields are initialized by system software to instruct the host controller when to execute portions of the split transaction protocol:

- SplitXState
This is a single bit residing in the Status field of an siTD (see [Table 17-50](#)). This bit is used to track the current state of the split transaction. The rules for managing this bit are described in [Section 17.6.12.3.3, “Split Transaction Execution State Machine for Isochronous.”](#)
- Frame S-mask
This is a bit-field wherein system software sets a bit corresponding to the microframe (within an H-Frame) that the host controller should execute a start-split transaction. This is always qualified by the value of the SplitXState bit. For example, referring to the IN example in [Figure 17-57](#), case 1, the S-mask would have a value of 0b0000_0001 indicating that if the siTD is traversed by the host controller, and the SplitXState indicates Do Start Split, and the current microframe as indicated by FRINDEX[2–0] is 0, then execute a start-split transaction.
- Frame C-mask
This is a bit-field where system software sets one or more bits corresponding to the microframes (within an H-Frame) that the host controller should execute complete-split transactions. The interpretation of this field is always qualified by the value of the SplitXState bit. For example, referring to the IN example in [Figure 17-57](#), case 1, the C-mask would have a value of 0b 0011_1100 indicating that if the siTD is traversed by the host controller, and the SplitXState indicates Do Complete Split, and the current microframe as indicated by FRINDEX[2–0] is 2, 3, 4, or 5, then execute a complete-split transaction.
- Back Pointer
This field in a siTD is used to complete an IN split-transaction using the previous H-Frame's siTD. This is only used when the scheduling of the complete-splits span an H-Frame boundary.

There exists a one-to-one relationship between a high-speed isochronous split transaction (including all start- and complete-splits) and one full-speed isochronous transaction. An siTD contains (amongst other things) buffer state and split transaction scheduling information. An siTD's buffer state always maps to one full-speed isochronous data payload. This means that for any full-speed transaction payload, a single siTD's data buffer must be used. This rule applies to both IN and OUTs. An siTD's scheduling information

usually also maps to one high-speed isochronous split transaction. The exception to this rule is the H-Frame boundary wrap cases mentioned above.

The siTD data structure describes at most, one frame's worth of high-speed transactions and that description is strictly bounded within a frame boundary. Figure 17-58 illustrates some examples. On the top are examples of the full-speed transaction footprints for the boundary scheduling cases described above. In the middle are time-frame references for both the B-Frames (HS/FS/LS Bus) and the H-Frames. On the bottom is illustrated the relationship between the scope of an siTD description and the time references. Each H-Frame corresponds to a single location in the periodic frame list. The implication is that each siTD is reachable from a single periodic frame list location at a time.

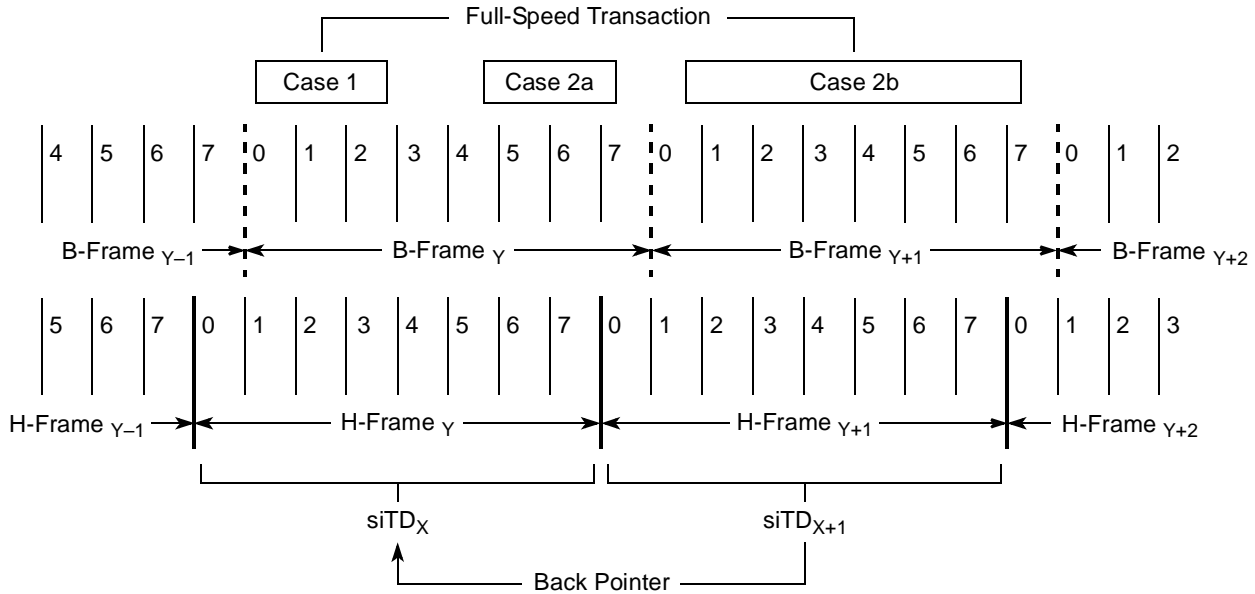


Figure 17-58. siTD Scheduling Boundary Examples

Each case is described below:

- Case 1: One siTD is sufficient to describe and complete the isochronous split transaction because the whole isochronous split transaction is tightly contained within a single H-Frame.
- Case 2a, 2b: Although both INs and OUTs can have these footprints, OUTs always take only one siTD to schedule. However, INs (for these boundary cases) require two siTDs to complete the scheduling of the isochronous split transaction. siTD_X is used to always issue the start-split and the first N complete-splits. The full-speed transaction (for these cases) can deliver data on the full-speed bus segment during microframe 7 of H-Frame _{$\gamma+1$} , or microframe 0 of H-Frame _{$\gamma+2$} . The complete splits are scheduled using siTD _{$X+2$} (not shown). The complete-splits to extract this data must use the buffer pointer from siTD _{$X+1$} . The only way for the host controller to reach siTD _{$X+1$} from H-Frame _{$\gamma+2$} is to use siTD _{$X+2$} 's back pointer.

Software must apply the following rules when calculating the schedule and linking the schedule data structures into the periodic schedule:

- Software must ensure that an isochronous split-transaction is started so that it will complete before the end of the B-Frame.

- Software must ensure that for a single full-speed isochronous endpoint, there is never a start-split and complete-split in H-Frame, microframe 1. This is mandated as a rule so that case 2a and case 2b can be discriminated. According to the core USB specification, the long isochronous transaction illustrated in Case 2b, could be scheduled so that the start-split was in microframe 1 of H-Frame N and the last complete-split would need to occur in microframe 1 of H-Frame N+1. However, it is impossible to discriminate between cases 2a and case 2b, which has significant impact on the complexity of the host controller.

17.6.12.3.2 Tracking Split Transaction Progress for Isochronous Transfers

Isochronous endpoints do not employ the concept of a halt on error, however the host controller does identify and report per-packet errors observed in the data stream. This includes schedule traversal problems (skipped microframes), timeouts and corrupted data received.

In similar kind to interrupt split-transactions, the portions of the split transaction protocol must execute in the microframes they are scheduled. The queue head data structure used to manage full- and low-speed interrupt has several mechanisms for tracking when portions of a transaction have occurred. Isochronous transfers use siTDs for their transfers and the data structures are only reachable using the schedule in the exact microframe in which they are required (so all the mechanism employed for tracking in queue heads is not required for siTDs). Software has the option of reusing siTD several times in the complete periodic schedule. However, it must ensure that the results of split transaction N are consumed and the siTD re-initialized (activated) before the host controller gets back to the siTD (in a future microframe).

Split-transaction isochronous OUTs utilize a low-level protocol to indicate which portions of the split transaction data have arrived. Control over the low-level protocol is exposed in an siTD using the fields Transaction Position (TP) and Transaction Count (T-count). If the entire data payload for the OUT split transaction is larger than 188 bytes, there will be more than one start-split transaction, each of which require proper annotation. If host hold-offs occur, then the sequence of annotations received from the host will not be complete, which is detected and handled by the transaction translator. See [Section 17.6.12.3.1, “Split Transaction Scheduling Mechanisms for Isochronous,”](#) for a description on how these fields are used during a sequence of start-split transactions.

The fields siTD[T-Count] and siTD[TP] are used by the host controller to drive and sequence the transaction position annotations. It is the responsibility of system software to properly initialize these fields in each siTD. Once the budget for a split-transaction isochronous endpoint is established, S-mask, T-Count, and TP initialization values for all the siTD associated with the endpoint are constant. They remain constant until the budget for the endpoint is recalculated by software and the periodic schedule adjusted.

For IN-endpoints, the transaction translator simply annotates the response data packets with enough information to allow the host controller to identify the last data. As with split transaction Interrupt, it is the host controller's responsibility to detect when it has missed an opportunity to execute a complete-split. The following field in the siTD is used to track and detect errors in the execution of a split transaction for an IN isochronous endpoint.

- C-prog-mask. This is an eight-bit bit-vector where the host controller keeps track of which complete-splits have been executed. Due to the nature of the transaction translator periodic pipeline, the complete-splits need to be executed in-order. The host controller needs to detect when

the complete-splits have not been executed in order. This can only occur due to system hold-offs where the host controller cannot get to the memory-based schedule. C-prog-mask is a simple bit-vector that the host controller sets a bit for each complete-split executed. The bit position is determined by the microframe (FRINDEX[2-0]) number in which the complete-split was executed. The host controller always checks C-prog-mask before executing a complete-split transaction. If the previous complete-splits have not been executed, then it means one (or more) have been skipped and data has potentially been lost. System software is required to initialize this field to zero before setting an siTD's Active bit to a one.

If a transaction translator returns with the final data before all of the complete-splits have been executed, the state of the transfer is advanced so that the remaining complete-splits are not executed. It is important to note that an IN siTD is retired based solely on the responses from the transaction translator to the complete-split transactions. This means, for example, that it is possible for a transaction translator to respond to a complete-split with an MDATA PID. The number of bytes in the MDATA's data payload could cause the siTD[Total Bytes to Transfer] field to decrement to zero. This response can occur, before all of the scheduled complete-splits have been executed. In other interface, data structures (for example, high-speed data streams through queue heads), the transition of Total Bytes to Transfer to zero signals the end of the transfer and results in clearing the Active bit. However, in this case, the result has not been delivered by the transaction translator and the host must continue with the next complete-split transaction to extract the residual transaction state. This scenario occurs because of the pipeline rules for a transaction translator. In summary, the periodic pipeline rules require that on a microframe boundary, the transaction translator holds the final two bytes received (if it has not seen an End Of Packet (EOP)) in the full-speed bus pipe stage and gives the remaining bytes to the high-speed pipeline stage. At the microframe boundary, the transaction translator could have received the entire packet (including both CRC bytes) but not received the packet EOP. In the next microframe, the transaction translator responds with an MDATA and sends all of the data bytes (with the two CRC bytes being held in the full-speed pipeline stage). This could cause the siTD to decrement its Total Bytes to Transfer field to zero, indicating it has received all expected data. The host must still execute one more (scheduled) complete-split transaction in order to extract the results of the full-speed transaction from the transaction translator (for example, the transaction translator may have detected a CRC failure, and this result must be forwarded to the host).

If the host experiences hold-offs that cause the host controller to skip one or more (but not all) scheduled split transactions for an isochronous OUT, then the protocol to the transaction translator is not consistent and the transaction translator detects and reacts to the problem. Likewise, for host hold-offs that cause the host controller to skip one or more (but not all) scheduled split transactions for an isochronous IN, the C-prog-mask is used by the host controller to detect errors. However, if the host experiences a hold-off that causes it to skip all of an siTD, or an siTD expires during a host hold off (for example, a hold-off occurs and the siTD is no longer reachable by the host controller in order for it to report the hold-off event), then system software must detect that the siTDs have not been processed by the host controller (for example, state not advanced) and report the appropriate error to the client driver.

17.6.12.3.3 Split Transaction Execution State Machine for Isochronous

In this section, all references to microframe are in the context of a microframe within an H-Frame.

If the Active bit in the Status byte is a zero, the host controller ignores the siTD and continues traversing the periodic schedule. Otherwise the host controller processes the siTD as specified below. A split

transaction state machine is used to manage the split-transaction protocol sequence. The host controller uses the fields defined in Section 17.6.12.3.2, “Tracking Split Transaction Progress for Isochronous Transfers,” plus the variable `cMicroFrameBit` defined in Section 17.6.12.2.5, “Split Transaction Execution State Machine for Interrupt,” to track the progress of an isochronous split transaction. Figure 17-59 illustrates the state machine for managing an `siTD` through an isochronous split transaction. Bold, dotted circles denote the state of the Active bit in the Status field of a `siTD`. The Bold, dotted arcs denote the transitions between these states. Solid circles denote the states of the split transaction state machine and the solid arcs denote the transitions between these states. Dotted arcs and boxes reference actions that take place either as a result of a transition or from being in a state.




Figure 17-59. Split Transaction State Machine for Isochronous

17.6.12.3.4 Periodic Isochronous—Do-Start-Split

Isochronous split transaction OUTs use only this state. An `siTD` for a split-transaction isochronous IN is either initialized to this state, or the `siTD` transitions to this state from Do Complete Split when a case 2a (IN) or 2b scheduling boundary isochronous split-transaction completes.

Each time the host controller reaches an active `siTD` in this state, it checks the `siTD[S-mask]` against `cMicroFrameBit`. If there is a one in the appropriate position, the `siTD` executes a start-split transaction. By definition, the host controller cannot reach an `siTD` at the wrong time. If the I/O field indicates an IN, then the start-split transaction includes only the extended token plus the full-speed token. Software must initialize the `siTD[Total Bytes To Transfer]` field to the number of bytes expected. This is usually the maximum packet size for the full-speed endpoint. The host controller exits this state when the start-split transaction is complete.

The remainder of this section is specific to an isochronous OUT endpoint (that is, the I/O field indicates an OUT). When the host controller executes a start-split transaction for an isochronous OUT it includes a data payload in the start-split transaction. The memory buffer address for the data payload is constructed by concatenating siTD[Current Offset] with the page pointer indicated by the page select field (siTD[P]). A zero in this field selects Page 0 and a 1 selects Page 1. During the start-split for an OUT, if the data transfer crosses a page boundary during the transaction, the host controller must detect the page cross, update the siTD[P] bit from a zero to a one, and begin using the siTD Page 1 with siTD[Current Offset] as the memory address pointer. The field siTD[TP] is used to annotate each start-split transaction with the indication of which part of the split-transaction data the current payload represents (ALL, BEGIN, MID, END). In all cases, the host controller simply uses the value in siTD[TP] to mark the start-split with the correct transaction position code.

T-count is always initialized to the number of start-splits for the current frame. TP is always initialized to the first required transaction position identifier. The scheduling boundary case (see [Figure 17-58](#)) is used to determine the initial value of TP. The initial cases are summarized in [Table 17-70](#).

Table 17-70. Initial Conditions for OUT siTD TP and T-Count Fields

| Case | T-Count | TP | Description |
|-------|---------|-------|--|
| 1, 2a | =1 | ALL | When the OUT data payload is less than (or equal to) 188 bytes, only one start-split is required to move the data. The one start-split must be marked with an ALL. |
| 1, 2a | !=1 | BEGIN | When the OUT data payload is greater than 188 bytes more than one start-split must be used to move the data. The initial start-split must be marked with a BEGIN. |

After each start-split transaction is complete, the host controller updates T-count and TP appropriately so that the next start-split is correctly annotated. [Table 17-71](#) illustrates all of the TP and T-count transitions, which must be accomplished by the host controller.

Table 17-71. Transaction Position (TP)/Transaction Count (T-Count) Transition Table

| TP | T-Count Next | TP Next | Description |
|-------|--------------|---------|--|
| ALL | 0 | N/A | Transition from ALL, to done. |
| BEGIN | 1 | END | Transition from BEGIN to END. Occurs when T-count starts at 2. |
| BEGIN | !=1 | MID | Transition from BEGIN to MID. Occurs when T-count starts at greater than 2. |
| MID | !=1 | MID | TP stays at MID while T-count is not equal to 1 (for example, greater than 1). This case can occur for any of the scheduling boundary cases where the T-count starts greater than 3. |
| MID | 1 | END | Transition from MID to END. This case can occur for any of the scheduling boundary cases where the T-count starts greater than 2. |

The start-split transactions do not receive a handshake from the transaction translator, so the host controller always advances the transfer state in the siTD after the bus transaction is complete. To advance the transfer state the following operations take place:

- The siTD[Total Bytes To Transfer] and the siTD[Current Offset] fields are adjusted to reflect the number of bytes transferred.
- The siTD[P] (page select) bit is updated appropriately.

- The siTD[TP] and siTD[T-count] fields are updated appropriately as defined in [Table 17-71](#).

These fields are then written back to the memory based siTD. The S-mask is fixed for the life of the current budget. As mentioned above, TP and T-count are set specifically in each siTD to reflect the data to be sent from this siTD. Therefore, regardless of the value of S-mask, the actual number of start-split transactions depends on T-count (or equivalently, Total Bytes to Transfer). The host controller must clear the Active bit when it detects that all of the schedule data has been sent to the bus. The preferred method is to detect when T-Count decrements to zero as a result of a start-split bus transaction. Equivalently, the host controller can detect when Total Bytes to Transfer decrements to zero. Either implementation must ensure that if the initial condition is Total Bytes to Transfer is equal to zero and T-count is equal to a one, the host controller issues a single start-split, with a zero-length data payload. Software must ensure that TP, T-count and Total Bytes to Transfer are set to deliver the appropriate number of bus transactions from each siTD. An inconsistent combination yields undefined behavior.

If the host experiences hold-offs that cause the host controller to skip start-split transactions for an OUT transfer, the state of the transfer does not progress appropriately. The transaction translator observes protocol violations in the arrival of the start-splits for the OUT endpoint (that is, the transaction position annotation is incorrect as received by the transaction translator).

Example scenarios are described in [Section 17.6.12.3.7, “Split Transaction for Isochronous—Processing Example.”](#)

The host controller can optionally track the progress of an OUT split transaction by setting appropriate bits in the siTD[C-prog-mask] as it executes each scheduled start-split. The checkPreviousBit() algorithm defined in [Section 17.6.12.3.5, “Periodic Isochronous—Do Complete Split,”](#) can be used prior to executing each start-split to determine whether start-splits were skipped. The host controller can use this mechanism to detect missed microframes. It can then clear the siTD's Active bit and stop execution of this siTD. This saves on both memory and high-speed bus bandwidth.

17.6.12.3.5 Periodic Isochronous—Do Complete Split

This state is only used by a split-transaction isochronous IN endpoint. This state is entered unconditionally from the Do Start State after a start-split transaction is executed for an IN endpoint. Each time the host controller visits an siTD in this state, it conducts a number of tests to determine whether it should execute a complete-split transaction. The sequence in which they are applied depends on which microframe the host controller is currently executing, which means that the tests might not be applied until after the siTD referenced from the back pointer has been fetched. The individual tests are as follows.

- Test A
cMicroFrameBit is bit-wise ANDed with the siTD[C-mask] field. A non-zero result indicates that software scheduled a complete-split for this endpoint, during this microframe. This test is always applied to a newly fetched siTD that is in this state.
- Test B
The siTD[C-prog-mask] bit vector is checked to determine whether the previous complete splits have been executed. An example algorithm is given below (this is slightly different than the algorithm used in [Section 17.6.12.2.7, “Periodic Interrupt—Do-Complete-Split”](#)). The sequence in which this test is applied depends on the current value of FRINDEX[2–0]. If FRINDEX[2–0] is 0 or 1, it is not applied until the back pointer has been used. Otherwise it is applied immediately.

```

Algorithm Boolean CheckPreviousBit(siTD.C-prog-mask, siTD.C-mask, cMicroFrameBit)
Begin
    Boolean rvalue = TRUE;
    previousBit = cMicroFrameBit rotate-right(1)
    -- Bit-wise anding previousBit with C-mask indicates whether there
    -- was an intent to send a complete split in the previous micro-
    -- frame. So, if the 'previous bit' is set in C-mask, check
    -- C-prog-mask to make sure it happened.
    if previousBit bitAND siTD.C-mask then
        if not (previousBit bitAND siTD.C-prog-mask) then
            rvalue = FALSE
        End if
    End if
    Return rvalue
End Algorithm

```

If Test A is true and FRINDEX[2–0] is zero or one, this is a case 2a or 2b scheduling boundary (see [Figure 17-57](#)). See [Section 17.6.12.3.6, “Complete-Split for Scheduling Boundary Cases 2a, 2b,”](#) for details in handling this condition.

If Test A and Test B evaluate to true, the host controller executes a complete-split transaction using the transfer state of the current siTD. When the host controller commits to executing the complete-split transaction, it updates QH[C-prog-mask] by bit-ORing with cMicroFrameBit. The transfer state is advanced based on the completion status of the complete-split transaction. To advance the transfer state of an IN siTD, the host controller must perform the following actions:

1. Decrement the number of bytes received from siTD[Total Bytes To Transfer]
2. Adjust siTD[Current Offset] by the number of bytes received
3. Adjust the siTD[P] (page select) field if the transfer caused the host controller to use the next page pointer
4. Set any appropriate bits in the siTD[Status] field, depending on the results of the transaction.

Note that if the host controller encounters a condition where siTD[Total Bytes To Transfer] is zero, and it receives more data, the host controller must not write the additional data to memory. The siTD[Status-Active] bit must be cleared and the siTD[Status-Babble Detected] bit must be set. The fields siTD[Total Bytes To Transfer], siTD[Current Offset], and siTD[P] are not required to be updated as a result of this transaction attempt.

The host controller accepts (assuming good data packet CRC and sufficient room in the buffer as indicated by the value of siTD[Total Bytes To Transfer]) MDATA and DATA0/1 data payloads up to and including 192 bytes. The host controller may optionally clear siTD[Status-Active] and set siTD[Status-Babble Detected] when it receives MDATA or DATA0/1 with a data payload of more than 192 bytes. The following responses have the noted effects:

- **ERR**
The full-speed transaction completed with a time-out or bad CRC and this is a reflection of that error to the host. The host controller sets the ERR bit in the siTD[Status] field and clears the Active bit.
- **Transaction Error (XactErr)**

The complete-split transaction encounters a Timeout, CRC16 failure, and so on. The siTD[Status] field XactErr field is set and the complete-split transaction must be retried immediately. The host controller must use an internal error counter to count the number of retries as a counter field is not provided in the siTD data structure. The host controller will not retry more than two times. If the host controller exhausts the retries or the end of the microframe occurs, the Active bit is cleared.

- DATA_x (0 or 1)

This response signals that the final data for the split transaction has arrived. The transfer state of the siTD is advanced and the Active bit is cleared. If the Bytes To Transfer field has not decremented to zero (including the reception of the data payload in the DATA_x response), then less data than was expected, or allowed for was actually received. This short packet event does not set the USB interrupt status bit (USBSTS[UI]) to a one. The host controller will not detect this condition.

- NYET (and Last)

On each NYET response, the host controller also checks to determine whether this is the last complete-split for this split transaction. Last was defined in [Section 17.6.12.2.7, “Periodic Interrupt—Do-Complete-Split.”](#) If it is the last complete-split (with a NYET response), then the transfer state of the siTD is not advanced (never received any data) and the Active bit is cleared. No bits are set in the Status field because this is essentially a skipped transaction. The transaction translator must have responded to all the scheduled complete-splits with NYETs, meaning that the start-split issued by the host controller was not received. This result should be interpreted by system software as if the transaction was completely skipped. The test for whether this is the last complete split can be performed by XORing C-mask with C-prog-mask. A zero result indicates that all complete-splits have been executed.

- MDATA (and Last)

See above description for testing for Last. This can only occur when there is an error condition. Either there has been a babble condition on the full-speed link, which delayed the completion of the full-speed transaction, or software set up the S-mask and/or C-masks incorrectly. The host controller must set the XactErr bit and clear the Active bit.

- NYET (and not Last)

See above description for testing for Last. The complete-split transaction received a NYET response from the transaction translator. Do not update any transfer state (except for C-prog-mask) and stay in this state.

- MDATA (and not Last)

The transaction translator responds with an MDATA when it has partial data for the split transaction. For example, the full-speed transaction data payload spans from microframe X to X+1 and during microframe X, the transaction translator responds with an MDATA and the data accumulated up to the end of microframe X. The host controller advances the transfer state to reflect the number of bytes received.

If Test A succeeds, but Test B fails, it means that one or more of the complete-splits have been skipped. The host controller sets the Missed Micro-Frame status bit and clears the Active bit.

17.6.12.3.6 Complete-Split for Scheduling Boundary Cases 2a, 2b

Boundary cases 2a and 2b (INs only) (see [Figure 17-57](#)) require that the host controller use the transaction state context of the previous siTD to finish the split transaction. [Table 17-72](#) enumerates the transaction state fields.

Table 17-72. Summary siTD Split Transaction State

| Buffer State | Status | Execution Progress |
|--|------------------------------|--------------------|
| Total Bytes To Transfer P (page select) Current Offset TP (transaction position) T-count (transaction count) | All bits in the status field | C-prog-mask |

NOTE

TP and T-count are used only for Host to Device (OUT) endpoints.

If software has budgeted the schedule of this data stream with a frame wrap case, then it must initialize the siTD[Back Pointer] field to reference a valid siTD and have the T bit in the siTD[Back Pointer] field cleared. Otherwise, software must set the T bit in siTD[Back Pointer]. The host controller's rules for interpreting when to use the siTD[Back Pointer] field are listed below. These rules apply only when the siTD's Active bit is a one and the SplitXState is Do Complete Split.

- When cMicroFrameBit is a 0x1 and the siTD_X[Back Pointer] T-bit is zero, or
- If cMicroFrameBit is a 0x2 and siTD_X[S-mask[0]] is zero

When either of these conditions apply, then the host controller must use the transaction state from siTD_{X-1}.

In order to access siTD_{X-1}, the host controller reads on-chip the siTD referenced from siTD_X[Back Pointer].

The host controller must save the entire state from siTD_X while processing siTD_{X-1}. This is to accommodate for case 2b processing. The host controller must not recursively walk the list of siTD[Back Pointers].

If siTD_{X-1} is active (Active bit is set and SplitXStat is Do Complete Split), then both Test A and Test B are applied as described above. If these criteria to execute a complete-split are met, the host controller executes the complete split and evaluates the results as described above. The transaction state (see [Table 17-72](#)) of siTD_{X-1} is appropriately advanced based on the results and written back to memory. If the resultant state of siTD_{X-1}'s Active bit is a one, then the host controller returns to the context of siTD_X, and follows its next pointer to the next schedule item. No updates to siTD_X are necessary.

If siTD_{X-1} is active (Active bit is set and SplitXStat is Do Start Split), then the host controller must clear the Active bit and set the Missed Micro-Frame status bit and the resultant status is written back to memory.

If siTD_{X-1}'s Active bit is cleared, (because it was cleared when the host controller first visited siTD_{X-1} via siTD_X's back pointer, it transitioned to zero as a result of a detected error, or the results of siTD_{X-1}'s complete-split transaction cleared it), then the host controller returns to the context of siTD_X and transitions its SplitXState to Do Start Split. The host controller then determines whether the case 2b start split boundary condition exists (that is, if cMicroframeBit is 1 and siTD_X[S-mask[0]] is 1). If this criterion

is met the host controller immediately executes a start-split transaction and appropriately advances the transaction state of $siTD_X$, then follows $siTD_X[Next\ Pointer]$ to the next schedule item. If the criterion is not met, the host controller simply follows $siTD_X[Next\ Pointer]$ to the next schedule item. Note that in the case of a 2b boundary case, the split-transaction of $siTD_{X-1}$ will have its Active bit cleared when the host controller returns to the context of $siTD_X$. Also, note that software should not initialize an $siTD$ with C-mask bits 0 and 1 set and an S-mask with bit 0 set. This scheduling combination is not supported and the behavior of the host controller is undefined.

17.6.12.3.7 Split Transaction for Isochronous—Processing Example

There is an important difference between how the hardware/software manages the isochronous split transaction state machine and how it manages the asynchronous and interrupt split transaction state machines. The asynchronous and interrupt split transaction state machines are encapsulated within a single queue head. The progress of the data stream depends on the progress of each split transaction. In some respects, the split-transaction state machine is sequenced using the Execute Transaction queue head traversal state machine.

Isochronous is a pure time-oriented transaction/data stream. The interface data structures are optimized to efficiently describe transactions that need to occur at specific times. The isochronous split-transaction state machine must be managed across these time-oriented data structures. This means that system software must correctly describe the scheduling of split-transactions across more than one data structure.

Then the host controller must make the appropriate state transitions at the appropriate times, in the correct data structures.

For example, [Table 17-73](#) illustrates a few frames worth of scheduling required to schedule a case 2a full-speed isochronous data stream.

Table 17-73. Example Case 2a—Software Scheduling $siTD$ s for an IN Endpoint

| $siTD_X$ | | Micro-Frames | | | | | | | | InitialSplitXState |
|----------|--------|--------------------------|---|---|---|---|---|---|---|--------------------|
| # | Masks | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | |
| X | S-Mask | | | | | 1 | | | | Do Start Split |
| | C-Mask | 1 | 1 | | | | | 1 | 1 | |
| X+1 | S-Mask | | | | | 1 | | | | Do Complete Split |
| | C-Mask | 1 | 1 | | | | | 1 | 1 | |
| X+2 | S-Mask | | | | | 1 | | | | Do Complete Split |
| | C-Mask | 1 | 1 | | | | | 1 | 1 | |
| X+3 | S-Mask | Repeats previous pattern | | | | | | | | Do Complete Split |
| | C-Mask | | | | | | | | | |

This example shows the first three $siTD$ s for the transaction stream. Since this is the case-2a frame-wrap case, S-masks of all $siTD$ s for this endpoint have a value of 0x10 (a one bit in microframe 4) and C-mask value of 0xC3 (one-bits in microframes 0,1, 6 and 7). Additionally, software ensures that the Back Pointer field of each $siTD$ references the appropriate $siTD$ data structure (and the Back Pointer T-bits are cleared).

The initial SplitXState of the first siTD is Do Start Split. The host controller will visit the first siTD eight times during frame X. The C-mask bits in microframes 0 and 1 are ignored because the state is Do Start Split. During microframe 4, the host controller determines that it can run a start-split (and does) and changes SplitXState to Do Complete Split. During microframes 6 and 7, the host controller executes complete-splits. Notice the siTD for frame X+1 has its SplitXState initialized to Do Complete Split. As the host controller continues to traverse the schedule during H-Frame X+1, it will visit the second siTD eight times. During microframes 0 and 1 it will detect that it must execute complete-splits.

During H-Frame X+1, microframe 0, the host controller detects that siTD_{X+1}'s Back Pointer[T] bit is a zero, saves the state of siTD_{X+1} and fetches siTD_X. It executes the complete split transaction using the transaction state of siTD_X. If the siTD_X split transaction is complete, siTD's Active bit is cleared and results written back to siTD_X. The host controller retains the fact that siTD_X is retired and transitions the SplitXState in siTD_{X+1} to Do Start Split. At this point, the host controller is prepared to execute the start-split for siTD_{X+1} when it reaches microframe 4. If the split-transaction completes early (transaction-complete is defined in [Section 17.6.12.3.5, “Periodic Isochronous—Do Complete Split”](#)), that is, before all the scheduled complete-splits have been executed, the host controller changes siTD_X[SplitXState] to Do Start Split early and naturally skips the remaining scheduled complete-split transactions. For this example, siTD_{X+1} does not receive a DATA0 response until H-Frame X+2, microframe 1.

During H-Frame X+2, microframe 0, the host controller detects that siTD_{X+2}'s Back Pointer[T] bit is zero, saves the state of siTD_{X+2} and fetches siTD_{X+1}. As described above, it executes another split transaction, receives an MDATA response, updates the transfer state, but does not modify the Active bit. The host controller returns to the context of siTD_{X+2}, and traverses its next pointer without any state change updates to siTD_{X+2}.

During H-Frame X+2, microframe 1, the host controller detects siTD_{X+2}'s S-mask[0] bit is zero, saves the state of siTD_{X+2} and fetches siTD_{X+1}. It executes another complete-split transaction, receives a DATA0 response, updates the transfer state and clears the Active bit. It returns to the state of siTD_{X+2} and changes its SplitXState to Do Start Split. At this point, the host controller is prepared to execute start-splits for siTD_{X+2} when it reaches microframe 4.

17.6.13 Port Test Modes

EHCI host controllers implement the port test modes Test J_State, Test K_State, Test_Packet, Test Force_Enable, and Test SEO_NAK as described in the *USB Specification Revision 2.0*. The required, port test sequence, assuming the CF-bit in the CONFIGFLAG register is set, is as follows:

1. Disable the periodic and asynchronous schedules by clearing the USBCMD[ASE] and USBCMD[PSE].
2. Place all enabled root ports into the suspended state by setting the Suspend bit in the PORTSC register (PORTSC[SUSP]).
3. Clear USBCMD[RS] (run/stop) and wait for USBSTS[HCH] to transition to a one. Note that an EHCI host controller implementation may optionally allow port testing with RS set. However, all host controllers must support port testing with RS cleared and HCH set.

4. Set the Port Test Control field in the port under test PORTSC register to the value corresponding to the desired test mode. If the selected test is Test_Force_Enable, then USBCMD[RS] must then be transitioned back to one, in order to enable transmission of SOFs out of the port under test.
5. When the test is complete, system software must ensure the host controller is halted (HCH bit is a one) then it terminates and exits test mode by setting USBCMD[RST].

17.6.14 Interrupts

The EHCI host controller hardware provides interrupt capability based on a number of sources. The following list describes the general groups of interrupt sources:

- Interrupts as a result of executing transactions from the schedule (success and error conditions),
- Host controller events (Port change events, and so on.)
- Host controller error events

All transaction-based sources are maskable through the host controller's Interrupt Enable register (USBINTR). Additionally, individual transfer descriptors can be marked to generate an interrupt on completion. This section describes each interrupt source and the processing that occurs in response to the interrupt.

During normal operation, interrupts may be immediate or deferred until the next interrupt threshold occurs. The interrupt threshold is a tunable parameter via the interrupt threshold control field in the USBCMD register. The value of this register controls when the host controller generates an interrupt on behalf of normal transaction execution. When a transaction completes during an interrupt interval period, the interrupt signaling the completion of the transfer will not occur until the interrupt threshold occurs. For example, the default value is eight microframes. This means that the host controller will not generate interrupts any more frequently than once every eight microframes.

[Section 17.6.14.2.4, “Host System Error”](#) details effects of a host system error.

If an interrupt has been scheduled to be generated for the current interrupt threshold interval, the interrupt is not signaled until after the status for the last complete transaction in the interval has been written back to system memory. This may sometimes result in the interrupt not being signaled until the next interrupt threshold.

Initial interrupt processing is the same, regardless of the reason for the interrupt. When an interrupt is signaled by the hardware, CPU control is transferred to host controller's USB interrupt handler. The precise mechanism to accomplish the transfer is OS specific. For this discussion it is just assumed that control is received. When the interrupt handler receives control, its first action is to read the USBSTS. It then acknowledges the interrupt by clearing all of the interrupt status bits by writing ones to these bit positions. The handler then determines whether the interrupt is due to schedule processing or some other event. After acknowledging the interrupt, the handler (via an OS-specific mechanism), schedules a deferred procedure call (DPC) which will execute later. The DPC routine processes the results of the schedule execution. The precise mechanisms used are beyond the scope of this document.

NOTE

The only method software should use for acknowledging an interrupt is by transitioning the appropriate status bits in the USBSTS register from a one to a zero.

17.6.14.1 Transfer/Transaction Based Interrupts

These interrupt sources are associated with transfer and transaction progress. They are all dependent on the next interrupt threshold.

17.6.14.1.1 Transaction Error

A transaction error is any error that caused the host controller to think that the transfer did not complete successfully. [Table 17-74](#) lists the events/responses that the host can observe as a result of a transaction. The effects of the error counter and interrupt status are summarized in the following paragraphs. Most of these errors set the XactErr status bit in the appropriate interface data structure.

Table 17-74. Summary of Transaction Errors

| Event/ Result | Queue Head/qTD/iTD/siTD Side Effects | | USBSTS[USBERRINT] |
|----------------------|--------------------------------------|--|-------------------|
| | Cerr | Status Field | |
| CRC | -1 | XactErr set | 1 ¹ |
| Timeout | -1 | XactErr set | 1 ¹ |
| Bad PID ² | -1 | XactErr set | 1 ¹ |
| Babble | N/A | See Section 17.6.14.1.2, "Serial Bus Babble" | 1 |
| Buffer Error | N/A | See Section 17.6.14.1.3, "Data Buffer Error" | |

¹ If occurs in a queue head, then USBERRINT is asserted only when Cerr counts down from a one to a zero. In addition the queue is halted.

² The host controller received a response from the device, but it could not recognize the PID as a valid PID.

There is a small set of protocol errors that relate only when executing a queue head and fit under the umbrella of a WRONG PID error that are significant to explicitly identify. When these errors occur, the XactErr status bit in the queue head is set and the Cerr field is decremented. When the PID Code indicates a SETUP, the following responses are protocol errors and result in XactErr bit being set and the Cerr field being decremented.

- EPS field indicates a high-speed device and it returns a Nak handshake to a SETUP.
- EPS field indicates a high-speed device and it returns a Nyet handshake to a SETUP.
- EPS field indicates a low- or full-speed device and the complete-split receives a Nak handshake.

17.6.14.1.2 Serial Bus Babble

When a device transmits more data on the USB than the host controller is expecting for this transaction, it is defined to be babbling. In general, this is called a packet babble. When a device sends more data than the maximum length number of bytes, the host controller sets the babble detected bit to a one and halts the

endpoint if it is using a queue head. Maximum length is defined as the minimum of total bytes to transfer and maximum packet size. The Cerr field is not decremented for a packet babble condition (only applies to queue heads).

A babble condition also exists if IN transaction is in progress at High-speed EOF2 point. This is called a frame babble. A frame babble condition is recorded into the appropriate schedule data structure. In addition, the host controller must disable the port to which the frame babble is detected.

USBSTS[UEI] (USB error interrupt) is set and if the USBINTR[UEE] (USB error interrupt enable) is set, then a hardware interrupt is signaled to the system at the next interrupt threshold. The host controller must never start an OUT transaction that babbles across a microframe EOF.

NOTE

When a host controller detects a data PID mismatch, it must either: disable the packet babble checking for the duration of the bus transaction or do packet babble checking based solely on Maximum Packet Size. The USB core specification defines the requirements on a data receiver when it receives a data PID mismatch (for example, expects a DATA0 and gets a DATA1 or visa-versa). In summary, it must ignore the received data and respond with an ACK handshake, in order to advance the transmitter's data sequence. The EHCI interface allows system software to provide buffers for a Control, Bulk or Interrupt IN endpoint that are not an even multiple of the maximum packet size specified by the device. Whenever a device misses an ACK for an IN endpoint, the host and device are out of synchronization with respect to the progress of the data transfer. The host controller may have advanced the transfer to a buffer that is less than maximum packet size. The device re-sends its maximum packet size data packet, with the original data PID, in response to the next IN token. In order to properly manage the bus protocol, the host controller must disable the packet babble check when it observes the data PID mismatch.

17.6.14.1.3 Data Buffer Error

This event indicates that an overrun of incoming data or a underrun of outgoing data has occurred for this transaction. This would generally be caused by the host controller not being able to access required data buffers in memory within necessary latency requirements. These conditions are not considered transaction errors, and do not effect the error count in the queue head. When these errors do occur, the host controller records the fact the error occurred by setting the Data Buffer Error bit in the queue head, iTD or siTD.

If the data buffer error occurs on a non-isochronous IN, the host controller will not issue a handshake to the endpoint. This forces the endpoint to resend the same data (and data toggle) in response to the next IN to the endpoint.

If the data buffer error occurs on an OUT, the host controller must corrupt the end of the packet so that it cannot be interpreted by the device as a good data packet. Simply truncating the packet is not considered acceptable. An acceptable implementation option is to 1's complement the CRC bytes and send them. There are other options suggested in the transaction translator section of the *USB Specification Revision 2.0*.

17.6.14.1.4 USB Interrupt (Interrupt on Completion (IOC))

Transfer Descriptors (iTDs, siTDs, and queue heads (qTDs)) contain a bit that can be set to cause an interrupt on their completion. The completion of the transfer associated with that schedule item causes USBSTS[UI] (USB interrupt) to be set. In addition, if a short packet is encountered on an IN transaction associated with a queue head, then this event also causes USBINT to be set. If USBINTR[UE] (USB interrupt enable) is set, a hardware interrupt is signaled to the system at the next interrupt threshold. If the completion is because of errors, USBSTS[UEI] (USB error interrupt) is also set.

17.6.14.1.5 Short Packet

Reception of a data packet that is less than the endpoint's Max Packet size during Control, Bulk or Interrupt transfers signals the completion of the transfer. Whenever a short packet completion occurs during a queue head execution, USBSTS[UI] (USB interrupt bit) is set. If the USB interrupt enable bit is set (USBINTR[UE]), a hardware interrupt is signaled to the system at the next interrupt threshold.

17.6.14.2 Host Controller Event Interrupts

These interrupt sources are independent of the interrupt threshold (with the one exception being the Interrupt on Async Advance).

17.6.14.2.1 Port Change Events

Port registers contain status and status change bits. When the status change bits are set, the host controller sets the USBSTS[PCI]. If the port change interrupt enable bit (PCE) in the USBINTR register is set, the host controller issues a hardware interrupt. The port status change bits in PORTSC include:

- Connect change status (CSC)
- Port enable/disable change (PEC)
- Over-current change (OCC)
- Force port resume (FPR)

17.6.14.2.2 Frame List Rollover

This event indicates that the host controller has wrapped the frame list. The current programmed size of the frame list effects how often this interrupt occurs. If the frame list size is 1024, then the interrupt occurs every 1024 milliseconds, if it is 512, then it occurs every 512 milliseconds, and so on. When a frame list rollover is detected, the host controller sets the frame list rollover bit, USBSTS[FRI]. If USBINTR[FRE] is set (frame list rollover enable), the host controller issues a hardware interrupt. This interrupt is not delayed to the next interrupt threshold.

17.6.14.2.3 Interrupt on Async Advance

This event is used for deterministic removal of queue heads from the asynchronous schedule. Whenever the host controller advances the on-chip context of the asynchronous schedule, it evaluates the value of USBCMD[IAA]. If it is set, it sets USBSTS[AAI]. If USBINTR[AAE] is set, the host controller issues a hardware interrupt at the next interrupt threshold. A detailed explanation of this feature is described in [Section 17.6.9.2, “Removing Queue Heads from Asynchronous Schedule.”](#)

17.6.14.2.4 Host System Error

The host controller is a bus master and any interaction between the host controller and the system may experience errors. The type of host error may be catastrophic to the host controller making it impossible for the host controller to continue in a coherent fashion. Behavior for these types of errors is to halt the host controller. Host-based error must result in the following actions:

- USBCMD[RS] is cleared.
- USBSTS[SEI] and USBSTS[HCH] register are set
- If the host system error enable bit, USBINTR[SEE] is set, the host controller issues a hardware interrupt. This interrupt is not delayed to the next interrupt threshold.

Table 17-75 summarizes the required actions taken on the various host errors.

Table 17-75. Summary Behavior on Host System Errors

| Cycle Type | Master Abort | Target Abort | Data Phase Parity |
|---------------------------------|--------------|--------------|-------------------|
| Frame list pointer fetch (read) | Fatal | Fatal | Fatal |
| siTD fetch (read) | Fatal | Fatal | Fatal |
| siTD status write-back (write) | Fatal | Fatal | Fatal |
| iTD fetch (read) | Fatal | Fatal | Fatal |
| iTD status write-back (write) | Fatal | Fatal | Fatal |
| qTD fetch (read) | Fatal | Fatal | Fatal |
| qHD status write-back (write) | Fatal | Fatal | Fatal |
| Data write | Fatal | Fatal | Fatal |
| Data read | Fatal | Fatal | Fatal |

NOTE

After a host system error, software must reset the host controller using USBCMD[RST] before re-initializing and restarting the host controller.

17.7 Device Data Structures

This section defines the interface data structures used to communicate control, status, and data between device controller driver (DCD) software and the device controller. The data structure definitions in this chapter support a 32-bit memory buffer address space. The interface consists of device queue heads and transfer descriptors.

NOTE

Software must ensure that no interface data structure reachable by the device controller spans a 4K-page boundary.

The data structures defined in the section are (from the device controller's perspective) a mix of read-only and read/writable fields. The device controller must preserve the read-only fields on all data structure writes.

The USB DR module includes DCD software called the USB 2.0 Device API. The device API provides an easy to use Application Program Interface for developing device (peripheral) applications. The device API incorporates and abstracts for the application developer all of the elements of the program interface.

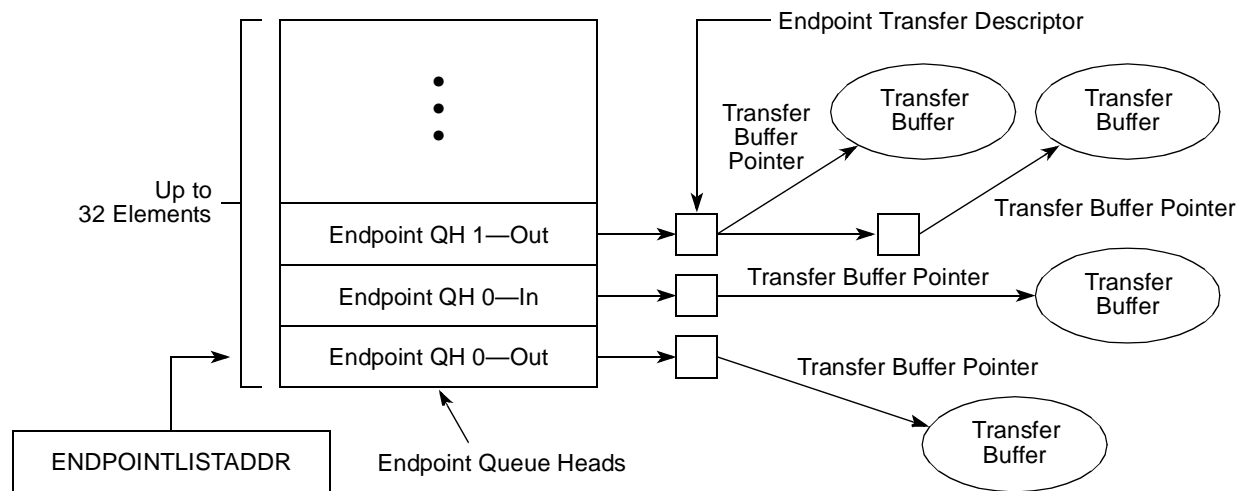


Figure 17-60. Endpoint Queue Head Organization

17.7.1 Endpoint Queue Head

The device Endpoint Queue Head (dQH) is where all transfers are managed. The dQH is a 48-byte data structure, but must be aligned on 64-byte boundaries. During priming of an endpoint, the dTD (device transfer descriptor) is copied into the overlay area of the dQH, which starts at the nextTD pointer DWord and continues through the end of the buffer pointers DWords. After a transfer is complete, the dTD status DWord is updated in the dTD pointed to by the currentTD pointer. While a packet is in progress, the overlay area of the dQH is used as a staging area for the dTD so that the Device Controller can access needed information with little minimal latency.

Figure 17-61 shows the Endpoint Queue Head structure.

| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | offset |
|--------------------------------------|----|--------------------------|----|----|----|-----------------------|----|----|----|-----------------------------|----|------------------|----|-----|----|--------------------|----|--------------------|-------------------|---------------------|----|-------------------|---|---|---|---|---|------|---|---|---|--------|
| Mult | | zlt | | 00 | | Maximum Packet Length | | | | | | | | | | ios | | 000_0000_0000_0000 | | | | | | | | | | 0x00 | | | | |
| Current dTD Pointer ¹ | | | | | | | | | | | | | | | | 0_0000 | | 0x04 | | | | | | | | | | | | | | |
| Next dTD Pointer ¹ | | | | | | | | | | | | | | | | 0000 | | T ¹ | 0x08 ² | | | | | | | | | | | | | |
| 0 | | Total Bytes ¹ | | | | | | | | | | ioc ¹ | | 000 | | MultO ¹ | | 00 | | Status ¹ | | 0x0C ² | | | | | | | | | | |
| Buffer Pointer (Page 0) ¹ | | | | | | | | | | Current Offset ¹ | | | | | | | | | | 0x10 ² | | | | | | | | | | | | |
| Buffer Pointer (Page 1) ¹ | | | | | | | | | | Reserved | | | | | | | | | | 0x14 ² | | | | | | | | | | | | |
| Buffer Pointer (Page 2) ¹ | | | | | | | | | | Reserved | | | | | | | | | | 0x18 ² | | | | | | | | | | | | |
| Buffer Pointer (Page 3) ¹ | | | | | | | | | | Reserved | | | | | | | | | | 0x1C ² | | | | | | | | | | | | |
| Buffer Pointer (Page 4) ¹ | | | | | | | | | | Reserved | | | | | | | | | | 0x20 ² | | | | | | | | | | | | |
| Reserved | | | | | | | | | | | | | | | | 0x24 | | | | | | | | | | | | | | | | |
| Setup Buffer Bytes 3–0 ¹ | | | | | | | | | | | | | | | | 0x28 | | | | | | | | | | | | | | | | |
| Setup Buffer Bytes 7–4 ¹ | | | | | | | | | | | | | | | | 0x2C | | | | | | | | | | | | | | | | |

Figure 17-61. Endpoint Queue Head Layout

¹ Device controller read/write; all others read-only.

² Offsets 0x08 through 0x20 contain the transfer overlay.

17.7.1.1 Endpoint Capabilities/Characteristics

This DWord specifies static information about the endpoint, in other words, this information does not change over the lifetime of the endpoint. Device controller software should not attempt to modify this information while the corresponding endpoint is enabled.

Table 17-76 describes the endpoint capabilities and characteristics fields.

Table 17-76. Endpoint Capabilities/Characteristics

| Bits | Name | Description |
|-------|------|--|
| 31–30 | Mult | Mult. This field is used to indicate the number of packets executed per transaction description as given by the following: 00 - Execute N Transactions as demonstrated by the USB variable length packet protocol where N is computed using the Maximum Packet Length (dQH) and the Total Bytes field (dTD) 01 Execute 1 Transaction. 10 Execute 2 Transactions. 11 Execute 3 Transactions. Note: Non-ISO endpoints must set Mult = 00. Note: ISO endpoints must set Mult = 01, 10, or 11 as needed. |
| 29 | zlt | Zero length termination select. This bit is used to indicate when a zero length packet is used to terminate transfers where to total transfer length is a multiple. This bit is not relevant for Isochronous transfers. 0 Enable zero length packet to terminate transfers equal to a multiple of the Maximum Packet Length. (default). 1 Disable the zero length packet on transfers that are equal in length to a multiple Maximum Packet Length. |

Table 17-76. Endpoint Capabilities/Characteristics (continued)

| Bits | Name | Description |
|-------|-----------------------|--|
| 28–27 | — | Reserved, should be cleared. These bit reserved for future use and should be cleared. |
| 26–16 | Maximum Packet Length | Maximum packet length. This directly corresponds to the maximum packet size of the associated endpoint (wMaxPacketSize). The maximum value this field may contain is 0x400 (1024). |
| 15 | ios | Interrupt on setup (IOS). This bit is used on control type endpoints to indicate if USBINT is set in response to a setup being received. |
| 14–0 | | Reserved, should be cleared. Bits reserved for future use and should be cleared. |

17.7.1.2 Transfer Overlay

The seven DWords in the overlay area represent a transaction working space for the device controller. The general operational model is that the device controller can detect whether the overlay area contains a description of an active transfer. If it does not contain an active transfer, then it will not read the associated endpoint.

After an endpoint is readied, the dTD is copied into this queue head overlay area by the device controller. Until a transfer is expired, software must not write the queue head overlay area or the associated transfer descriptor. When the transfer is complete, the device controller will write the results back to the original transfer descriptor and advance the queue.

See dTD for a description of the overlay fields.

17.7.1.3 Current dTD Pointer

The current dTD pointer is used by the device controller to locate the transfer in progress. This word is for USB_DR (hardware) use only and should not be modified by DCD software.

[Table 17-77](#) describes the current dTD pointer fields.

Table 17-77. Current dTD Pointer

| Bits | Description |
|------|---|
| 31–5 | Current dtd. This field is a pointer to the dTD that is represented in the transfer overlay area. This field is modified by the Device Controller to next dTD pointer during endpoint priming or queue advance. |
| 4–0 | Reserved, should be cleared. Bit reserved for future use and should be cleared. |

17.7.1.4 Setup Buffer

The setup buffer is dedicated storage for the 8-byte data that follows a setup PID.

NOTE

Each endpoint has a TX and an RX dQH associated with it, and only the RX queue head is used for receiving setup data packets.

Table 17-78 describes the multiple mode control fields.

Table 17-78. Multiple Mode Control

| DWord | Bits | Description |
|-------|------|--|
| 1 | 31–0 | Setup Buffer 0. This buffer contains bytes 3 to 0 of an incoming setup buffer packet and is written by the device controller to be read by software. |
| 2 | 31–0 | Setup Buffer 1. This buffer contains bytes 7 to 4 of an incoming setup buffer packet and is written by the device controller to be read by software. |

17.7.2 Endpoint Transfer Descriptor (dTD)

The dTD describes the location and quantity of data to be sent/received for given transfer to the device controller. The DCD should not attempt to modify any field in an active dTD except the Next Link Pointer, which should only be modified as described in Section 17.8.5, “Managing Transfers with Transfer Descriptors.”

Figure 17-62 shows the endpoint transfer descriptor.

| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | offset |
|--------------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----------------------------|---------------------------|----|-------|----|----|---------------------|---|---|---|---|------|------|---|------|---|--------|
| Next Link Pointer | | | | | | | | | | | | | | | | | | | | | | | | | | | | 0000 | T | 0x00 | | |
| Total Bytes ¹ | | | | | | | | | | | | | | | | ioc | 000 | | MultO | 00 | | Status ¹ | | | | | | | | | | 0x04 |
| Buffer Pointer (Page 0) | | | | | | | | | | | | | | | | Current Offset ¹ | | | | | | | | | | | | 0x08 | | | | |
| Buffer Pointer (Page 1) | | | | | | | | | | | | | | | | 0 | Frame Number ¹ | | | | | | | | | | 0x0C | | | | | |
| Buffer Pointer (Page 2) | | | | | | | | | | | | | | | | 0000_0000_0000 | | | | | | | | | | | | 0x10 | | | | |
| Buffer Pointer (Page 3) | | | | | | | | | | | | | | | | 0000_0000_0000 | | | | | | | | | | | | 0x14 | | | | |
| Buffer Pointer (Page 4) | | | | | | | | | | | | | | | | 0000_0000_0000 | | | | | | | | | | | | 0x18 | | | | |

Figure 17-62. Endpoint Transfer Descriptor (dTD)

¹ Device controller read/write; all others read-only.

Table 17-79 describes the next dTD pointer fields.

Table 17-79. Next dTD Pointer

| Bits | Description |
|------|---|
| 31–5 | Next transfer element pointer. This field contains the physical memory address of the next dTD to be processed. The field corresponds to memory address signals [31:5], respectively. |
| 4–1 | Reserved, should be cleared. Bits reserved for future use and should be cleared. |
| 0 | Terminate (T). 1=pointer is invalid. 0=Pointer is valid (points to a valid Transfer Element Descriptor). This bit indicates to the Device Controller that there are no more valid entries in the queue. |

Table 17-80 describes the next dTD token fields.

Table 17-80. dTD Token

| Bits | Description |
|-------|---|
| 31 | Reserved, should be cleared. Bit reserved for future use and should be cleared. |
| 30–16 | <p>Total Bytes. This field specifies the total number of bytes to be moved with this transfer descriptor. This field is decremented by the number of bytes actually moved during the transaction and only on the successful completion of the transaction.</p> <p>The maximum value software may store in the field is 5*4K(5000H). This is the maximum number of bytes 5 page pointers can access. Although it is possible to create a transfer up to 20K this assumes the 1st offset into the first page is 0. When the offset cannot be predetermined, crossing past the 5th page can be guaranteed by limiting the total bytes to 16K**. Therefore, the maximum recommended transfer is 16K(4000H).</p> <p>If the value of the field is zero when the host controller fetches this transfer descriptor (and the active bit is set), the device controller executes a zero-length transaction and retires the transfer descriptor.</p> <p>It is not a requirement for IN transfers that Total Bytes To Transfer be an even multiple of Maximum Packet Length. If software builds such a transfer descriptor for an IN transfer, the last transaction will always be less than Maximum Packet Length.</p> |
| 15 | Interrupt On Complete (IOC). This bit is used to indicate if USBINT is to be set in response to device controller being finished with this dTD. |
| 14–12 | Reserved, should be cleared. Bits reserved for future use and should be cleared. |
| 11–10 | <p>Multiplier Override (MultiO). This field can be used for transmit ISO's (that is, ISO-IN) to override the multiplier in the QH. This field must be zero for all packet types that are not transmit-ISO.</p> <p>Example:</p> <p>if QH.multiplier = 3; Maximum packet size = 8; Total bytes = 15; MultiO = 0 [default] Three packets are sent: {Data2(8); Data1(7); Data0(0)}</p> <p>if QH.multiplier = 3; Maximum packet size = 8; Total bytes = 15; MultiO = 2 Two packets are sent: {Data1(8); Data0(7)}</p> <p>For maximal efficiency, software should compute MultiO = greatest integer of (Total Bytes/Max. Packet Size) except for the case when Total bytes = 0; then MultiO should be 1.</p> <p>Note: Non-ISO and non-TX endpoints must set MultiO = 00.</p> |
| 9–8 | Reserved, should be cleared. Bits reserved for future use and should be cleared. |
| 7–0 | <p>Status. This field is used by the Device Controller to communicate individual command execution states back to the Device Controller software. This field contains the status of the last transaction performed on this qTD. The bit encodings are:</p> <p>BitStatus Field Description</p> <p>7 Active 6 Halted 5 Data Buffer Error 3 Transaction Error 4,2,0Reserved, should be cleared</p> |

Table 17-81–Table 17-83 the buffer pointer page *n* fields.

Table 17-81. Buffer Pointer Page 0

| Bits | Description |
|-------|--|
| 31–12 | Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers. |
| 11–0 | Current Offset. Offset into the 4kb buffer where the packet is to begin. |

Table 17-82. Buffer Pointer Page 1

| Bits | Description |
|-------|---|
| 31–12 | Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers. |
| 11 | Reserved |
| 10–0 | Frame Number. Written by the device controller to indicate the frame number in which a packet finishes. This is typically be used to correlate relative completion times of packets on an ISO endpoint. |

Table 17-83. Buffer Pointer Pages 2–4

| Bits | Description |
|-------|--|
| 31–12 | Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers. |
| 11–0 | Reserved |

17.8 Device Operational Model

The function of the device operation is to transfer a request in the memory image to and from the Universal Serial Bus. Using a set of linked list transfer descriptors, pointed to by a queue head, the device controller will perform the data transfers. The following sections explain the use of the device controller from the device controller driver (DCD) point-of-view and further describe how specific USB bus events relate to status changes in the device controller programmer's interface.

17.8.1 Device Controller Initialization

After hardware reset, the USB DR module is disabled until the run/stop bit (USBCMD[RS]) is set to a '1'. In the disabled state, the pull-up on the USB D+ is not active which prevents an attach event from occurring. At a minimum, it is necessary to have the queue heads setup for endpoint zero before the device attach occurs. Shortly after the device is enabled, a USB reset will occur followed by setup packet arriving at endpoint 0. A queue head must be prepared so that the device controller can store the incoming setup packet.

The device can be configured for either the internal UTMI PHY or an external ULPI PHY enabled with various clocking options. The configuration of each PHY interface and the various clocking options are detailed below.

To configure the internal UTMI PHY the following initialization sequence is required:

1. After power-on reset the UTMI PHY will be in disabled state and the PLL will be held reset.
2. Set the CONTROL[REFSEL] bits, 24–25 to correspond to the UTMI PLL the appropriate reference clock frequency.
3. Set the CONTROL[CLKIN_SEL] bits to select the source and divider of the UTMI reference clock.
4. Set the CONTROL[PHY_CLK_SEL] bits to select the UTMI PHY as the source of USB controller PHY clock.

5. Set the CONTROL[UTMI_PHY_EN] to enable the UTMI PHY and release the PLL.
6. Wait for the PHY clock to become valid. This can be determined by polling the CONTROL[PHY_CLK_VALID] status bit.

Once the PHY clock is valid the user can proceed to the host controller initialization phase.

To configure the external ULPI PHY the following initialization sequence is required.

1. After power-on reset the UTMI PHY will be in disabled state and the PLL will be held reset. The UTMI PHY should remain disabled if the ULPI is being used.
2. Set the CONTROL[PHY_CLK_SEL] bits to select the ULPI PHY as the source of USB controller PHY clock.
3. Wait for PHY clock to become valid. This can be determined by polling the CONTROL[PHY_CLK_VALID] status bit. Note that this bit is not valid once the CONTROL[USB_EN] bit is set.

Once the PHY clock is valid the user can proceed to the device controller initialization phase.

In order to initialize a device, the software should perform the following steps:

1. Set the controller mode to device mode. Optionally set USBMODE[SDIS] (streaming disable).

NOTE

Transitioning from host mode to device mode requires a device controller reset before modifying USBMODE.

2. Optionally modify the BURSTSIZE register.
3. Program PORTSC[PTS] if using a non-ULPI PHY.
4. Set CONTROL[USB_EN]
5. Allocate and initialize device queue heads in system memory Minimum: Initialize device queue heads 0 Tx and 0 Rx.

NOTE

All device queue heads must be initialized for control endpoints before the endpoint is enabled. Device queue heads for non-control endpoints must be initialized before the endpoint can be used.

For information on device queue heads, refer to [Section 17.7, “Device Data Structures.”](#)

6. Configure the ENDPOINTLISTADDR pointer.
For additional information on ENDPOINTLISTADDR, refer to the register table.
7. Enable the microprocessor interrupt associated with the USB DR module and optionally change setting of USBCMD[ITC].
Recommended: enable all device interrupts including: USBINT, USBERRINT, Port Change Detect, USB Reset Received, DCSuspend.
For a list of available interrupts refer to the USBINTR and the USBSTS register tables.
8. Set USBCMD[RS] to run mode.

After the run bit is set, a device reset will occur. The DCD must monitor the reset event and set the DEVICEADDR register, set the ENDPTCTRLx registers, and adjust the software state as described in [Section 17.8.2.1, “Bus Reset.”](#)

NOTE

Endpoint 0 is designed as a control endpoint only and does not need to be configured using ENDPTCTRL0 register.

It is also not necessary to initially prime Endpoint 0 because the first packet received will always be a setup packet. The contents of the first setup packet will require a response in accordance with USB device framework command set.

17.8.2 Port State and Control

From a chip or system reset, the USB_DR enters the powered state. A transition from the powered state to the attach state occurs when the run/stop bit (USBCMD[RS]) is set to a ‘1’. After receiving a reset on the bus, the port will enter the defaultFS or defaultHS state in accordance with the protocol reset described in Appendix C.2 of the *USB Specification Rev. 2.0*. [Figure 17-63](#) depicts the state of a USB 2.0 device.

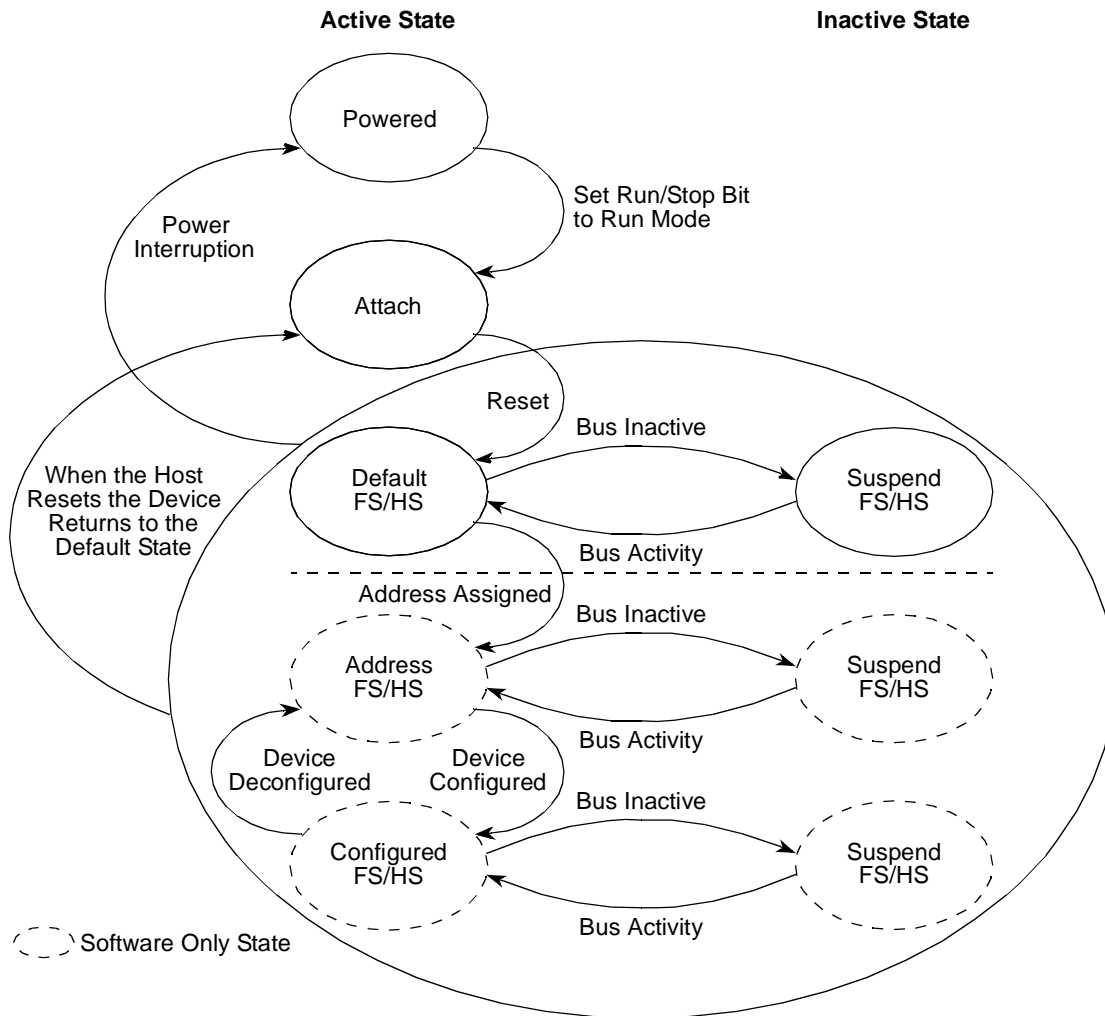


Figure 17-63. USB 2.0 Device States

States powered, attach, defaultFS/HS, suspendFS/HS are implemented in the USB_DR and are communicated to the DCD using status bits, as shown in [Table 17-84](#):

Table 17-84. Device Controller State Information Bits

| Bits | Register |
|--------------------------|----------|
| DCSuspend (SLI) | USBSTS |
| USB Reset Received (URI) | USBSTS |
| Port Change Detect (PCI) | USBSTS |
| High-Speed Port | PORTSC |

It is the responsibility of the DCD to maintain a state variable to differentiate between the defaultFS/HS state and the address/configured states. Change of state from default to address and the configured states is part of the enumeration process described in the device framework section of the USB 2.0 Specification.

As a result of entering the address state, the device address register (DEVICEADDR) must be programmed by the DCD.

Entry into the configured indicates that all endpoints to be used in the operation of the device have been properly initialized by programming the ENDPTCTRL n registers and initializing the associated queue heads.

17.8.2.1 Bus Reset

A bus reset is used by the host to initialize downstream devices. When a bus reset is detected, the USB_DR controller will renegotiate its attachment speed, reset the device address to 0, and notify the DCD by interrupt (assuming the USB reset interrupt enable bit, USBINTR[URE], is set). After a reset is received, all endpoints (except endpoint 0) are disabled and any primed transactions are canceled by the device controller. The concept of priming is clarified below, but the DCD must perform the following tasks when a reset is received:

1. Clear all setup token semaphores by reading the ENDPTSETUPSTAT register and writing the same value back to the ENDPTSETUPSTAT register.
2. Clear all the endpoint complete status bits by reading the ENDPTCOMPLETE register and writing the same value back to the ENDPTCOMPLETE register.
3. Cancel all primed status by waiting until all bits in the ENDPTPRIME are 0 and then writing 0xFFFF_FFFF to ENDPTFLUSH.
4. Read the reset bit in the PORTSC register (PORTSC[PR]) and make sure that it is still active.
 - A USB reset occurs for a minimum of 3 ms, and the DCD must reach this point in the reset cleanup before end of the reset occurs.
 - If it does not, a hardware reset of the device controller is recommended. A hardware reset can be performed by writing a one to the USB_DR reset bit in (USBCMD[RST]). Note that a hardware reset will cause the device to detach from the bus by clearing USBCMD[RS] bit. Thus, the DCD must completely re-initialize the USB_DR after a hardware reset.
5. Free all allocated dTDs because they will no longer be executed by the device controller. If this is the first time the DCD is processing a USB reset event, then it is likely that no dTDs have been allocated.

At this time, the DCD may release control back to the OS because no further changes to the device controller are permitted until a Port Change Detect is indicated.

After a Port Change Detect, the device has reached the default state and the DCD can read the PORTSC to determine if the device is operating in FS or HS mode. At this time, the device controller has reached normal operating mode and DCD can begin enumeration according to the USB Chapter 9, Device Framework.

NOTE

The device DCD may use the FS/HS mode information to determine the bandwidth mode of the device.

In some applications, it may not be possible to enable one or more pipes while in FS mode. Beyond the data rate issue, there is no difference in DCD operation between FS and HS modes.

17.8.2.2 Suspend/Resume

This section discusses the suspend and resume functions.

17.8.2.2.1 Suspend Description

In order to conserve power, USB_DR automatically enters the suspended state when no bus traffic has been observed for a specified period. When suspended, the USB_DR maintains any internal status, including its address and configuration. Attached devices must be prepared to suspend at any time they are powered, regardless of if they have been assigned a non-default address, are configured, or neither. Bus activity may cease due to the host entering a suspend mode of its own. In addition, a USB device shall also enter the suspended state when the hub port it is attached to is disabled.

The USB_DR exits suspend mode when there is bus activity. It may also request the host to exit suspend mode or selective suspend by using electrical signaling to indicate remote wake-up. The ability of a device to signal remote wake-up is optional. The USB_DR is capable of remote wake-up signaling. When the USB_DR is reset, remote wake-up signaling must be disabled.

17.8.2.2.2 Suspend Operational Model

The USB_DR moves into the suspend state when suspend signaling is detected or activity is missing on the upstream port for more than a specific period. After the device controller enters the suspend state, the DCD is notified by an interrupt (assuming DC Suspend Interrupt is enabled). When the USBSTS[SLI] (device controller suspend) is set, the device controller is suspended.

DCD response when the device controller is suspended is application specific and may involve switching to low-power operation. Information on the bus power limits in suspend state can be found in USB 2.0 specification.

17.8.2.2.3 Resume

If the USB_DR is suspended, its operation is resumed when any non-idle signaling is received on its upstream facing port. In addition, the USB_DR can signal the system to resume operation by forcing resume signaling to the upstream port. Resume signaling is sent upstream by writing a '1' to the PORTSC[FPR] (resume bit) while the device is in suspend state. Sending resume signal to an upstream port should cause the host to issue resume signaling and bring the suspended bus segment (one more devices) back to the active condition.

NOTE

Before resume signaling can be used, the host must enable it by using the Set Feature command defined in device framework (Chapter 9) of the USB 2.0 Specification.

17.8.3 Managing Endpoints

The USB 2.0 specification defines an endpoint, also called a device endpoint or an address endpoint as a uniquely addressable portion of a USB device that can source or sink data in a communications channel

between the host and the device. The endpoint address is specified by the combination of the endpoint number and the endpoint direction.

The channel between the host and an endpoint at a specific device represents a data pipe. Endpoint 0 for a device is always a control type data channel used for device discovery and enumeration. Other types of endpoints support by USB include bulk, interrupt, and isochronous. Each endpoint type has specific behavior related to packet response and error handling. More detail on endpoint operation can be found in the USB 2.0 specification.

The USB_DR supports up to three endpoint specified numbers. The DCD can enable, disable, and configure each endpoint.

Each endpoint direction is essentially independent and can be configured with differing behavior in each direction. For example, the DCD can configure endpoint 1-IN to be a bulk endpoint and endpoint 1-OUT to be an isochronous endpoint. This helps to conserve the total number of endpoints required for device operation. The only exception is that control endpoints must use both directions on a single endpoint number to function as a control endpoint. Endpoint 0 is, for example, is always a control endpoint and uses the pair of directions.

Each endpoint direction requires a queue head allocated in memory. If the maximum of 6 endpoint numbers, one for each endpoint direction are being used by the device controller, then 12 queue heads are required. The operation of an endpoint and use of queue heads are described later in this document.

17.8.3.1 Endpoint Initialization

After hardware reset, all endpoints except endpoint zero are uninitialized and disabled. The DCD must configure and enable each endpoint by writing to configuration bit in the `ENDPTCTRLn` register. Each 32-bit `ENDPTCTRLn` is split into an upper and lower half. The lower half of `ENDPTCTRLn` is used to configure the receive or OUT endpoint and the upper half is likewise used to configure the corresponding transmit or IN endpoint. Control endpoints must be configured the same in both the upper and lower half of the `ENDPTCTRLn` register otherwise the behavior is undefined. [Table 17-85](#) shows how to construct a configuration word for endpoint initialization.

Table 17-85. Device Controller Endpoint Initialization

| Field | Value |
|---------------------|---|
| Data Toggle Reset | 1 |
| Data Toggle Inhibit | 0 |
| Endpoint Type | 00 Control 01 Isochronous 10 Bulk 11 Interrupt |
| Endpoint Stall | 0 |

17.8.3.1.1 Stalling

There are two occasions where the USB_DR may need to return to the host a STALL.

The first occasion is the functional stall, which is a condition set by the DCD as described in the USB 2.0 device framework (Chapter 9). A functional stall is only used on non-control endpoints and can be enabled in the device controller by setting the endpoint stall bit in the `ENDPTCTRLn` register associated with the given endpoint and the given direction. In a functional stall condition, the device controller will continue to return STALL responses to all transactions occurring on the respective endpoint and direction until the endpoint stall bit is cleared by the DCD.

A protocol stall, unlike a function stall, is used on control endpoints is automatically cleared by the device controller at the start of a new control transaction (setup phase). When enabling a protocol stall, the DCD should enable the stall bits (both directions) as a pair. A single write to the `ENDPTCTRLn` register can ensure that both stall bits are set at the same instant.

NOTE

Any write to the `ENDPTCTRLn` register during operational mode must preserve the endpoint type field (that is, perform a read-modify-write).

Table 17-86 describes the device controller stall response matrix.

Table 17-86. Device Controller Stall Response Matrix

| USB Packet | Endpoint Stall Bit. | Effect on STALL Bit. | USB Response |
|---|---------------------|----------------------|--------------|
| SETUP packet received by a non-control endpoint | N/A | None | STALL |
| IN/OUT/PING packet received by a non-control endpoint | '1 | None | STALL |
| IN/OUT/PING packet received by a non-control endpoint | '0 | None | ACK/NAK/NYET |
| SETUP packet received by a control endpoint | N/A | Cleared | ACK |
| IN/OUT/PING packet received by a control endpoint | '1 | None | STALL |
| IN/OUT/PING packet received by a control endpoint | '0 | None | ACK/NAK/NYET |

17.8.3.2 Data Toggle

Data toggle is a mechanism to maintain data coherency between host and device for any given data pipe. For more information on data toggle, refer to the *Universal Serial Bus Revision 2.0 Specification*.

17.8.3.2.1 Data Toggle Reset

The DCD may reset the data toggle state bit and cause the data toggle sequence to reset in the device controller by writing a '1' to the data toggle reset bit in the `ENDPTCTRLn` register. This should only be necessary when configuring/initializing an endpoint or returning from a STALL condition.

17.8.3.2.2 Data Toggle Inhibit

This feature is for test purposes only and should never be used during normal device controller operation.

Setting the data toggle Inhibit bit active ('1') causes the USB_DR to ignore the data toggle pattern that is normally sent and accept all incoming data packets regardless of the data toggle state.

In normal operation, the USB_DR checks the DATA0/DATA1 bit against the data toggle to determine if the packet is valid. If Data PID does not match the data toggle state bit maintained by the device controller for that endpoint, the Data toggle is considered not valid. If the data toggle is not valid, the device controller assumes the packet was already received and discards the packet (not reporting it to the DCD). To prevent the USB_DR from re-sending the same packet, the device controller will respond to the error packet by acknowledging it with either an ACK or NYET response.

17.8.3.3 Device For Packet Transfers

All transactions on the USB bus are initiated by the host and in turn, the device must respond to any request from the host within the turnaround time stated in the *Universal Serial Bus Revision 2.0 Specification*.

A USB host will send requests to the USB_DR in an order that can not be precisely predicted as a single pipeline, so it is not possible to prepare a single packet for the device controller to execute. However, the order of packet requests is predictable when the endpoint number and direction is considered. For example, if endpoint 2 (transmit direction) is configured as a bulk pipe, then it is expected that the host will send IN requests to that endpoint. This USB_DR prepares packets for each endpoint/direction in anticipation of the host request. The process of preparing the device controller to send or receive data in response to host initiated transaction on the bus is referred to as 'priming' the endpoint. This term is used throughout the following documentation to describe the USB_DR operation so the DCD can be architected properly use priming. Further, note that the term 'flushing' is used to describe the action of clearing a packet that was queued for execution.

17.8.3.3.1 Priming Transmit Endpoints

Priming a transmit endpoint will cause the device controller to fetch the device transfer descriptor (dTD) for the transaction pointed to by the device queue head (dQH). After the dTD is fetched, it is stored in the dQH until the device controller completes the transfer described by the dTD. Storing the dTD in the dQH allows the device controller to fetch the operating context needed to handle a request from the host without the need to follow the linked list, starting at the dQH when the host request is received.

After the device has loaded the dTD, the leading data in the packet is stored in a FIFO in the device controller. This FIFO is split into virtual channels so that the leading data can be stored for any endpoint up to the maximum number of endpoints configured at device synthesis time.

After a priming request is complete, an endpoint state of primed is indicated in the ENDPTSTATUS register. For a primed transmit endpoint, the device controller can respond to an IN request from the host and meet the stringent bus turnaround time of High Speed USB.

Since only the leading data is stored in the device controller FIFO, it is necessary for the device controller to begin filling in behind leading data after the transaction starts. The FIFO must be sized to account for the maximum latency that can be incurred by the system memory bus.

17.8.3.3.2 Priming Receive Endpoints

Priming receive endpoints is identical to priming of transmit endpoints from the point of view of the DCD. At the device controller the major difference in the operational model is that there is no data movement of the leading packet data simply because the data is to be received from the host.

Note as part of the architecture, the FIFO for the receive endpoints is not partitioned into multiple channels like the transmit FIFO. Thus, the size of the RX FIFO does not scale with the number of endpoints.

17.8.3.4 Interrupt/Bulk Endpoint Operational Model

The behaviors of the device controller for interrupt and bulk endpoints are identical. All valid IN and OUT transactions to bulk pipes will handshake with a NAK unless the endpoint had been primed. Once the endpoint has been primed, data delivery will commence.

A dTD is retired by the device controller when the packets described in the transfer descriptor have been completed. Each dTD describes N packets to be transferred according to the USB Variable Length transfer protocol. The formula, [Table 17-87](#), and [Table 17-88](#) describe how the device controller computes the number and length of the packets to be sent/received by the USB vary according to the total number of bytes and maximum packet length.

With Zero Length Termination (ZLT) = 0

$$N = \text{INT}(\text{number of bytes}/\text{max. packet length}) + 1$$

With Zero Length Termination (ZLT) = 1

$$N = \text{MAXINT}(\text{number of bytes}/\text{max. packet length})$$

Table 17-87. Variable Length Transfer Protocol Example (ZLT = 0)

| Bytes (dTD) | Max. Packet Length (dQH) | N | P1 | P2 | P3 |
|-------------|--------------------------|---|-----|-----|----|
| 511 | 256 | 2 | 256 | 255 | — |
| 512 | 256 | 3 | 256 | 256 | 0 |
| 512 | 512 | 2 | 512 | 0 | — |

Table 17-88. Variable Length Transfer Protocol Example (ZLT = 1)

| Bytes (dTD) | Max. Packet Length (dQH) | N | P1 | P2 | P3 |
|-------------|--------------------------|---|-----|-----|----|
| 511 | 256 | 2 | 256 | 255 | — |
| 512 | 256 | 2 | 256 | 256 | — |
| 512 | 512 | 1 | 512 | — | — |

NOTE

The MULT field in the dQH must be set to '00' for bulk, interrupt, and control endpoints.

TX-dTD is complete when:

- All packets described dTD were successfully transmitted. *** Total bytes in dTD will equal zero when this occurs.

RX-dTD is complete when:

- All packets described in dTD were successfully received. *** Total bytes in dTD will equal zero when this occurs.
- A short packet (number of bytes < maximum packet length) was received. *** This is a successful transfer completion; DCD must check Total Bytes in dTD to determine the number of bytes that are remaining. From the total bytes remaining in the dTD, the DCD can compute the actual bytes received.
- A long packet was received (number of bytes > maximum packet size) OR (total bytes received > total bytes specified). *** This is an error condition. The device controller will discard the remaining packet, and set the Buffer Error bit in the dTD. In addition, the endpoint is flushed and the USBERR interrupt will become active.

On the successful completion of the packet(s) described by the dTD, the active bit in the dTD is cleared and the next pointer will be followed when the Terminate bit is clear. When the Terminate bit is set, the USB_DR will flush the endpoint/direction and cease operations for that endpoint/direction.

On the unsuccessful completion of a packet (see long packet above), the dQH is left pointing to the dTD that was in error. In order to recover from this error condition, the DCD must properly re-initialize the dQH by clearing the active bit and update the nextTD pointer before attempting to re-prime the endpoint.

NOTE

All packet level errors such as a missing handshake or CRC error will be retried automatically by the device controller.

There is no required interaction with the DCD for handling such errors.

17.8.3.4.1 Interrupt/Bulk Endpoint Bus Response Matrix

Table 17-89 shows the interrupt/bulk endpoint bus response matrix.

Table 17-89. Interrupt/Bulk Endpoint Bus Response Matrix

| | Stall | Not Primed | Primed | Underflow | Overflow |
|----------------|--------|------------|---------------------------------|-----------------------|----------|
| Setup | Ignore | Ignore | Ignore | N/A | N/A |
| In | STALL | NAK | Transmit | BS Error ¹ | N/A |
| Out | STALL | NAK | Receive + NYET/ACK ² | N/A | NAK |
| Ping | STALL | NAK | ACK | N/A | N/A |
| Invalid | Ignore | Ignore | Ignore | Ignore | Ignore |

¹ Force Bit Stuff Error.

² NYET/ACK—NYET unless the Transfer Descriptor has packets remaining according to the USB variable length protocol then ACK.
SYSERR—System error should never occur when the latency FIFOs are correctly sized and the DCD is responsive.

17.8.3.5 Control Endpoint Operation Model

This section discusses the control endpoint operation model.

17.8.3.5.1 Setup Phase

All requests to a control endpoint begin with a setup phase followed by an optional data phase and a required status phase. The USB_DR will always accept the setup phase unless the setup lockout is engaged.

The setup lockout will engage so that future setup packets are ignored. Lockout of setup packets ensures that while software is reading the setup packet stored in the queue head, that data is not written as it is being read potentially causing an invalid setup packet.

The setup lockout mechanism can be disabled and a tripwire type semaphore will ensure that the setup packet payload is extracted from the queue head without being corrupted by an incoming setup packet. This is the preferred behavior because ignoring repeated setup packets due to long software interrupt latency would be a compliance issue.

Setup Packet Handling

- Disable Setup Lockout by writing '1' to Setup Lockout Mode (SLOM) in USBMODE (once at initialization). Setup lockout is not necessary when using the tripwire as described below.

NOTE

Leaving the Setup Lockout Mode as '0' will result in a potential compliance issue.

- After receiving an interrupt and inspecting ENDPTSETUPSTAT to determine that a setup packet was received on a particular pipe:
 - Write '1' to clear corresponding bit ENDPTSETUPSTAT.

- Write ‘1’ to Setup Tripwire (SUTW) in USBCMD register.
- Duplicate contents of dQH.SetupBuffer into local software byte array.
- Read Setup TripWire (SUTW) in USBCMD register. (if set—continue; if cleared—goto 2)
- Write ‘0’ to clear Setup Tripwire (SUTW) in USBCMD register.
- Process setup packet using local software byte array copy and execute status/handshake phases.

NOTE

After receiving a new setup packet the status and/or handshake phases may still be pending from a previous control sequence. These should be flushed and de-allocated before linking a new status and/or handshake dTD for the most recent setup packet.

17.8.3.5.2 Data Phase

Following the setup phase, the DCD must create a device transfer descriptor for the data phase and prime the transfer.

After priming the packet, the DCD must verify a new setup packet has not been received by reading the ENDPTSETUPSTAT register immediately verifying that the prime had completed. A prime will complete when the associated bit in the ENDPTPRIME register is zero and the associated bit in the ENDPTSTATUS register is a one. If a prime fails, that is, The ENDPTPRIME bit goes to zero and the ENDPTSTATUS bit is not set, then the prime has failed. This can only be due to improper setup of the dQH, dTD or a setup arriving during the prime operation. If a new setup packet is indicated after the ENDPTPRIME bit is cleared, then the transfer descriptor can be freed and the DCD must reinterpret the setup packet.

Should a setup arrive after the data stage is primed, the device controller will automatically clear the prime status (ENDPTSTATUS) to enforce data coherency with the setup packet.

NOTE

The MULT field in the dQH must be set to ‘00’ for bulk, interrupt, and control endpoints.

NOTE

Error handling of data phase packets is the same as bulk packets described previously.

17.8.3.5.3 Status Phase

Similar to the data phase, the DCD must create a transfer descriptor (with byte length equal zero) and prime the endpoint for the status phase. The DCD must also perform the same checks of the ENDPTSETUPSTAT as described above in the data phase.

NOTE

The MULT field in the dQH must be set to ‘00’ for bulk, interrupt, and control endpoints.

NOTE

Error handling of data phase packets is the same as bulk packets described previously.

17.8.3.5.4 Control Endpoint Bus Response Matrix

Table 17-90 shows the device controller response to packets on a control endpoint, according to the device controller state.

Table 17-90. Control Endpoint Bus Response Matrix

| Token Type | Endpoint State | | | | | Setup Lockout |
|------------|----------------|------------|---------------------------------|-----------------------|---------------------|---------------|
| | Stall | Not Primed | Primed | Underflow | Overflow | |
| Setup | ACK | ACK | ACK | N/A | SYSERR ¹ | |
| In | STALL | NAK | Transmit | BS Error ² | N/A | N/A |
| Out | STALL | NAK | Receive + NYET/ACK ³ | N/A | NAK | N/A |

Table 17-90. Control Endpoint Bus Response Matrix (continued)

| Token Type | Endpoint State | | | | | Setup Lockout |
|------------|----------------|------------|--------|-----------|----------|---------------|
| | Stall | Not Primed | Primed | Underflow | Overflow | |
| Ping | STALL | NAK | ACK | N/A | N/A | N/A |
| Invalid | Ignore | Ignore | Ignore | Ignore | Ignore | Ignore |

¹ SYSERR—System error should never occur when the latency FIFOs are correctly sized and the DCD is responsive.

² Force Bit Stuff Error.

³ NYET/ACK—NYET unless the Transfer Descriptor has packets remaining according to the USB variable length protocol then ACK.

17.8.3.6 Isochronous Endpoint Operational Model

Isochronous endpoints are used for real-time scheduled delivery of data and their operational model is significantly different than the host throttled Bulk, Interrupt, and Control data pipes. Real time delivery by the USB_DR will be accomplished by the following:

- Exactly MULT Packets per (micro)Frame are transmitted/received. Note that MULT is a two-bit field in the device Queue Head. The variable length packet protocol is not used on isochronous endpoints.
- NAK responses are not used. Instead, zero length packets are sent in response to an IN request to unprimed endpoints. For unprimed RX endpoints, the response to an OUT transaction is to ignore the packet within the device controller.
- Prime requests always schedule the transfer described in the dTD for the next (micro)frame. If the ISO-dTD is still active after that frame, then the ISO-dTD is held ready until executed or canceled by the DCD.

The USB_DR in host mode uses the periodic frame list to schedule data exchanges to Isochronous endpoints. The operational model for device mode does not use such a data structure. Instead, the same dTD used for Control/Bulk/Interrupt endpoints is also used for isochronous endpoints. The difference is in the handling of the dTD.

The first difference between bulk and ISO-endpoints is that priming an ISO-endpoint is a delayed operation such that an endpoint will become primed only after a SOF is received. After the DCD writes the prime bit, the prime bit is cleared as usual to indicate to software that the device controller completed priming the dTD for transfer. Internal to the design, the device controller hardware masks that prime start until the next frame boundary. This behavior is hidden from the DCD but occurs so that the device controller can match the dTD to a specific (micro)frame.

Another difference with isochronous endpoints is that the transaction must wholly complete in a (micro)frame. Once an ISO transaction is started in a (micro)frame it will retire the corresponding dTD when MULT transactions occur or the device controller finds a fulfillment condition.

The transaction error bit set in the status field indicates a fulfillment error condition. When a fulfillment error occurs, the frame after the transfer failed to complete wholly, the device controller will force retire the ISO-dTD and move to the next ISO-dTD.

It is important to note that fulfillment errors are only caused due to partially completed packets. If no activity occurs to a primed ISO-dTD, the transaction will stay primed indefinitely. This means it is up to software discard transmit ISO-dTDs that pile up from a failure of the host to move the data.

Finally, the last difference with ISO packets is in the data level error handling. When a CRC error occurs on a received packet, the packet is not retried similar to bulk and control endpoints. Instead, the CRC is noted by setting the Transaction Error bit and the data is stored as usual for the application software to sort out.

- TX Packet Retired
 - MULT counter reaches zero.
 - Fulfillment Error [Transaction Error bit is set]
 - #Packets Occurred > 0 AND # Packets Occurred < MULT

NOTE

For TX-ISO, MULT Counter can be loaded with a lesser value in the dTD Multiplier Override field. If the Multiplier Override is zero, the MULT Counter is initialized to the Multiplier in the QH.

- RX Packet Retired:
 - MULT counter reaches zero.
 - Non-MDATA Data PID is received
 - Overflow Error:
 - Packet received is > maximum packet length. [Buffer Error bit is set]
 - Packet received exceeds total bytes allocated in dTD. [Buffer Error bit is set]
 - Fulfillment Error [Transaction Error bit is set]
 - # Packets Occurred > 0 AND # Packets Occurred < MULT
 - CRC Error [Transaction Error bit is set]

NOTE

For ISO, when a dTD is retired, the next dTD is primed for the next frame. For continuous (micro)frame to (micro)frame operation the DCD should ensure that the dTD linked-list is out ahead of the device controller by at least two (micro)frames.

17.8.3.6.1 Isochronous Pipe Synchronization

When it is necessary to synchronize an isochronous data pipe to the host, the (micro)frame number (FRINDEX register) can be used as a marker. To cause a packet transfer to occur at a specific (micro)frame number [N], the DCD should interrupt on SOF during frame N – 1. When the FRINDEX = N – 1, the DCD

must write the prime bit. The USB_DR will prime the isochronous endpoint in (micro)frame N – 1 so that the device controller will execute delivery during (micro)frame N.

CAUTION

Priming an endpoint towards the end of (micro)frame N – 1 will not guarantee delivery in (micro)frame N. The delivery may actually occur in (micro)frame N + 1 if device controller does not have enough time to complete the prime before the SOF for packet N is received.

17.8.3.6.2 Isochronous Endpoint Bus Response Matrix

Table 17-91 shows the isochronous endpoint bus response matrix.

Table 17-91. Isochronous Endpoint Bus Response Matrix

| | Stall | Not Primed | Primed | Underflow | Overflow |
|---------|--------------------------|-------------|----------|-----------------------|-------------|
| Setup | STALL | STALL | STALL | N/A | N/A |
| In | NULL ¹ Packet | NULL Packet | Transmit | BS Error ² | N/A |
| Out | Ignore | Ignore | Receive | N/A | Drop Packet |
| Ping | Ignore | Ignore | Ignore | Ignore | Ignore |
| Invalid | Ignore | Ignore | Ignore | Ignore | Ignore |

¹ Zero Length Packet.

² Force Bit Stuff Error.

17.8.4 Managing Queue Heads

Figure 17-64 shows the endpoint queue head diagram.

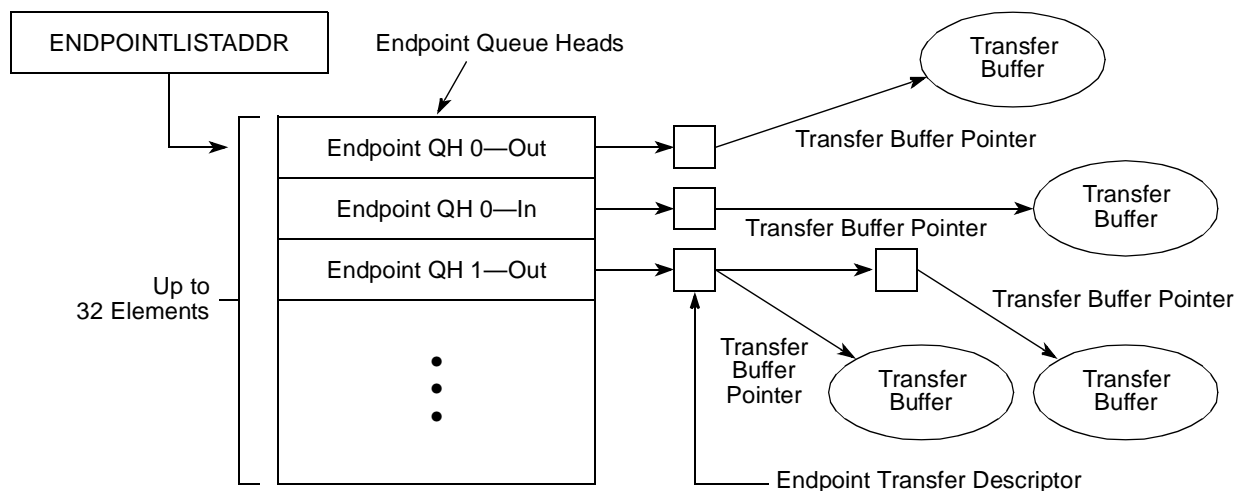


Figure 17-64. Endpoint Queue Head Diagram

The device queue head (dQH) points to the linked list of transfer tasks, each depicted by the device Transfer Descriptor (dTDD). An area of memory pointed to by ENDPOINTLISTADDR contains a group of all dQHs in a sequential list as shown in [Figure 17-64](#). The even elements in the list of dQH's are used for receive endpoints (OUT/SETUP) and the odd elements are used for transmit endpoints (IN/INTERRUPT). Device transfer descriptors are linked head to tail starting at the queue head and ending at a terminate bit. Once the dTD has been retired, it will no longer be part of the linked list from the queue head. Therefore, software is required to track all transfer descriptors since pointers will no longer exist within the queue head once the dTD is retired (see [Section 17.8.5.1, "Software Link Pointers"](#)).

In addition to the current and next pointers and the dTD overlay, the dQH also contains the following parameters for the associated endpoint: Multiplier, Maximum Packet Length, Interrupt On Setup. The complete initialization of the dQH including these fields is demonstrated in the next section.

17.8.4.1 Queue Head Initialization

One pair of device queue heads must be initialized for each active endpoint. To initialize a device queue head:

- Write the MaxPacketSize field as required by the USB Chapter 9 or application specific protocol.
- Write the multiplier field to 0 for control, bulk, and interrupt endpoints. For ISO endpoints, set the multiplier to 1, 2, or 3 as required bandwidth and in conjunction with the USB Chapter 9 protocol. Note that in FS mode, the multiplier field can only be 1 for ISO endpoints.
- Write the next dTD Terminate bit field to '1.'
- Write the Active bit in the status field to '0.'
- Write the Halt bit in the status field to '0.'

NOTE

The DCD must only modify dQH if the associated endpoint is not primed and there are no outstanding dTDs.

17.8.4.2 Operational Model for Setup Transfers

As discussed in [Section 17.8.3.5, "Control Endpoint Operation Model,"](#) setup transfer requires special treatment by the DCD. A setup transfer does not use a dTD but instead stores the incoming data from a setup packet in an 8-byte buffer within the dQH.

Upon receiving notification of the setup packet, the DCD should handle the setup transfer as demonstrated here:

1. Copy setup buffer contents from dQH - RX to software buffer.
2. Acknowledge setup backup by writing a '1' to the corresponding bit in ENDPTSETUPSTAT.

NOTE

The acknowledge must occur before continuing to process the setup packet.

NOTE

After the acknowledge has occurred, the DCD must not attempt to access the setup buffer in the dQH - RX. Only the local software copy should be examined.

3. Check for pending data or status dTD's from previous control transfers and flush if any exist as discussed in [Section 17.8.5.5, "Flushing/Depriming an Endpoint."](#)

NOTE

It is possible for the device controller to receive setup packets before previous control transfers complete. Existing control packets in progress must be flushed and the new control packet completed.

4. Decode setup packet and prepare data phase [optional] and status phase transfer as required by the USB Chapter 9 or application specific protocol.

17.8.5 Managing Transfers with Transfer Descriptors

17.8.5.1 Software Link Pointers

It is necessary for the DCD software to maintain head and tail pointers to the for the linked list of dTDs for each respective queue head. This is necessary because the dQH only maintains pointers to the current working dTD and the next dTD to be executed. The operations described in next section for managing dTD will assume the DCD can use reference the head and tail of the dTD linked list.

NOTE

To conserve memory, the reserved fields at the end of the dQH can be used to store the Head and Tail pointers but it still remains the responsibility of the DCD to maintain the pointers.

Figure 17-65 shows the software link pointers.

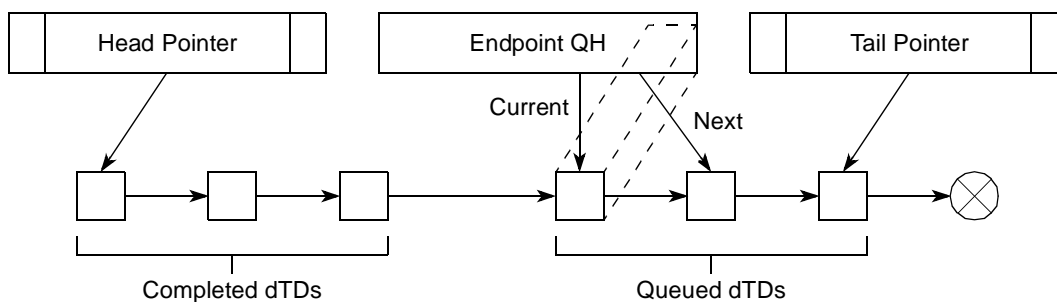


Figure 17-65. Software Link Pointers

17.8.5.2 Building a Transfer Descriptor

Before a transfer can be executed from the linked list, a dTD must be built to describe the transfer. Use the following procedure for building dTDs.

Allocate 8-DWord dTD block of memory aligned to 8-DWord boundaries. Example: bit address 4–0 would be equal to ‘00000’.

Write the following fields:

1. Initialize first seven DWords to ‘0’.
2. Set the terminate bit to ‘1’.
3. Fill in total bytes with transfer size.
4. Set the interrupt on complete if desired.
5. Initialize the status field with the active bit set to ‘1’ and all remaining status bits set to ‘0’.
6. Fill in buffer pointer page 0 and the current offset to point to the start of the data buffer.
7. Initialize buffer pointer page 1 through page 4 to be one greater than each of the previous buffer pointer.

17.8.5.3 Executing a Transfer Descriptor

To safely add a dTD, the DCD must account for the event in which the device controller reaches the end of the dTD list at the same time a new dTD is being added to the end of the list.

First, determine whether the link list is empty by checking the DCD driver to see if the pipe is empty (internal representation of linked-list should indicate if any packets are outstanding). Then follow the sequence of actions in the following list as appropriate, depending on whether the link list is empty or not empty.

- Link list is empty
 - a) Write dQH next pointer AND dQH terminate bit to ‘0’ as a single DWord operation.
 - b) Clear active and halt bit in dQH (in case set from a previous error).
 - c) Prime endpoint by writing ‘1’ to correct bit position in ENDPTPRIME.
- Link list is not empty
 - a) Add dTD to end of linked list.
 - b) Read correct prime bit in ENDPTPRIME—if ‘1’ DONE.
 - c) Set ATDTW bit in USBCMD register to ‘1’.
 - d) Read correct status bit in ENDPTSTATUS. (store in tmp. variable for later).
 - e) Read ATDTW bit in USBCMD register.
 - If ‘0’ goto 3.
 - If ‘1’ continue to 6.
 - f) Write ATDTW bit in USBCMD register to ‘0’.
 - g) If status bit read in (3) is ‘1’ DONE.
 - h) If status bit read in (3) is ‘0’ then Goto Case 1: Step 1.

17.8.5.4 Transfer Completion

After a dTD has been initialized and the associated endpoint primed, the device controller will execute the transfer upon the host-initiated request. The DCD is notified with a USB interrupt if the Interrupt On

Complete bit was set or alternately, the DCD can poll the endpoint complete register to find when the dTD had been executed. After a dTD has been executed, DCD can check the status bits to determine success or failure.

NOTE

Multiple dTD can be completed in a single endpoint complete notification. After clearing the notification, DCD must search the dTD linked list and retire all dTDs that have finished (Active bit cleared).

By reading the status fields of the completed dTDs, the DCD can determine if the transfers completed successfully. Success is determined with the following combination of status bits:

- Active = 0
- Halted = 0
- Transaction Error = 0
- Data Buffer Error = 0

Should any combination other than the one shown above exist, the DCD must take proper action. Transfer failure mechanisms are indicated in the Device Error Matrix.

In addition to checking the status bit the DCD must read the Transfer Bytes field to determine the actual bytes transferred. When a transfer is complete, the Total Bytes transferred is by decremented by the actual bytes transferred. For Transmit packets, a packet is only complete after the actual bytes reaches zero, but for receive packets, the host may send fewer bytes in the transfer according the USB variable length packet protocol.

17.8.5.5 Flushing/Depriming an Endpoint

It is necessary for the DCD to flush to deprime one more endpoints on a USB device reset or during a broken control transfer. There may also be application specific requirements to stop transfers in progress. The following procedure can be used by the DCD to stop a transfer in progress:

1. Write a '1' to the corresponding bit(s) in ENDPTFLUSH.
2. Wait until all bits in ENDPTFLUSH are '0'.
3. Software note: this operation may take a large amount of time depending on the USB bus activity. It is not desirable to have this wait loop within an interrupt service routine.
4. Read ENDPTSTATUS to ensure that for all endpoints commanded to be flushed, that the corresponding bits are now '0' If the corresponding bits are '1' after step #2 has finished, then the flush failed as described in the following:

Explanation: In very rare cases, a packet is in progress to the particular endpoint when commanded flush using ENDPTFLUSH. A safeguard is in place to refuse the flush to ensure that the packet in progress completes successfully. The DCD may need to repeatedly flush any endpoints that fail to flush by repeating steps 1–3 until each endpoint is successfully flushed.

17.8.5.6 Device Error Matrix

Table 17-92 summarizes packet errors that are not automatically handled by the USB controller.

Table 17-92. Device Error Matrix

| Error | Direction | Packet Type | Data Buffer Error Bit | Transaction Error Bit |
|-----------------------|-----------|-------------|-----------------------|-----------------------|
| Overflow ** | RX | Any | 1 | 0 |
| ISO Packet Error | RX | ISO | 0 | 1 |
| ISO Fulfillment Error | Both | ISO | 0 | 1 |

Notice that the device controller handles all errors on Bulk/Control/Interrupt Endpoints except for a data buffer overflow. However, for ISO endpoints, errors packets are not retried and errors are tagged as indicated. Table 17-93 provides the error descriptions.

Table 17-93. Error Descriptions

| | |
|------------------------------|---|
| Overflow | Number of bytes received exceeded max. packet size or total buffer length. Note: This error also sets the Halt bit in the dQH. If there are dTDs remaining in the linked list for the endpoint, they will not be executed. |
| ISO Packet Error | CRC Error on received ISO packet. Contents not guaranteed to be correct. |
| ISO Fulfillment Error | Host failed to complete the number of packets defined in the dQH mult field within the given (micro)frame. For scheduled data delivery the DCD may need to readjust the data queue because a fulfillment error will cause Device Controller to cease data transfers on the pipe for one (micro)frame. During the dead (micro)frame, the Device Controller reports error on the pipe and primes for the following frame. |

17.8.6 Servicing Interrupts

The interrupt service routine must consider that there are high-frequency, low-frequency operations, and error operations and order accordingly.

17.8.6.1 High-Frequency Interrupts

High frequency interrupts in particular should be handed in the order shown in Table 17-94. The most important of these is listed first because the DCD must acknowledge a setup buffer in the timeliest manner possible.

Table 17-94. Interrupt Handling Order

| Execution Order | Interrupt | Action |
|-----------------|--|--|
| 1a | USB Interrupt ¹ ENDPTSETUPSTATUS | Copy contents of setup buffer and acknowledge setup packet (as indicated in Section 17.8.4, "Managing Queue Heads"). Process setup packet according to USB 2.0 Chapter 9 or application specific protocol. |

Table 17-94. Interrupt Handling Order (continued)

| Execution Order | Interrupt | Action |
|-----------------|-----------------------------|---|
| 1b | USB Interrupt ENDPTCOMPLETE | Handle completion of dTD as indicated in Section 17.8.4, “Managing Queue Heads” . |
| 2 | SOF Interrupt | Action as deemed necessary by application. This interrupt may not have a use in all applications. |

¹ It is likely that multiple interrupts to stack up on any call to the Interrupt Service Routine AND during the Interrupt Service Routine.

17.8.6.2 Low-Frequency Interrupts

The low frequency events include the interrupts shown in [Table 17-95](#). These interrupts can be handled in any order because they do not occur often in comparison to the high-frequency interrupts.

Table 17-95. Low Frequency Interrupt Events

| Interrupt | Action |
|------------------------|---|
| Port Change | Change software state information. |
| Sleep Enable (Suspend) | Change software state information. Low power handling as necessary. |
| Reset Received | Change software state information. Abort pending transfers. |

17.8.6.3 Error Interrupts

Error interrupts are the least frequently occurring events. They should be placed last in the interrupt service routine. [Table 17-96](#) shows the error interrupt events.

Table 17-96. Error Interrupt Events

| Interrupt | Action |
|---------------------|---|
| USB Error Interrupt | This error is redundant because it combines USB Interrupt and an error status in the dTD. The DCD will more aptly handle packet-level errors by checking dTD status field upon receipt of USB Interrupt (w/ ENDPTCOMPLETE). |
| System Error | Unrecoverable error. Immediate Reset of core; free transfers buffers in progress and restart the DCD. |

17.9 Deviations from the EHCI Specifications

The host mode operation of the USB DR module is nearly EHCI-compatible with few minor differences. For the most part, the module conforms to the data structures and operations described in Section 3, “Data Structures,” and Section 4, “Operational Model,” in the EHCI specification. The particulars of the deviations occur in the following areas:

- Embedded transaction translator—Allows direct attachment of FS and LS devices in host mode without the need for a companion controller.

- Device operation—In host mode, the device operational registers are generally disabled and thus device mode is mostly transparent when in host mode. However, there are a couple exceptions documented in the following sections.
- Embedded design interface—The module does not have a PCI interface and therefore the PCI configuration registers described in the EHCI specification are not applicable.
- For the purposes of the DR implementing dual-role host/device controller with support for OTG applications, it is necessary to deviate from the EHCI specification. Device operation and OTG operation are not specified in the EHCI and thus the implementation supported in the DR module is proprietary.

17.9.1 Embedded Transaction Translator Function

The DR module supports directly connected full and low speed devices without requiring a companion controller by including the capabilities of a USB 2.0 high speed hub transaction translator. Although there is no separate transaction translator block in the system, the transaction translator function normally associated with a high speed hub has been implemented within the DMA and Protocol engine blocks. The embedded transaction translator function is an extension to EHCI interface, but makes use of the standard data structures and operational models that exist in the EHCI specification to support full and low speed devices.

17.9.1.1 Capability Registers

The following additions have been added to the capability registers to support the embedded transaction translator Function:

- N_TT added to HSCPARAMS—Host Controller Structural Parameters
- N_PTT added to HSCPARAMS—Host Controller Structural Parameters

See [Section 17.3.1.3, “Host Controller Structural Parameters \(HSCPARAMS\),”](#) for usage information.

17.9.1.2 Operational Registers

The following additions have been added to the operational registers to support the embedded TT:

- ASYNCTTSTS is a new register.
- Addition of two-bit Port Speed (PSPD) to the PORTSC register.

17.9.1.3 Discovery

In a standard EHCI controller design, the EHCI host controller driver detects a Full speed (FS) or Low speed (LS) device by noting if the port enable bit is set after the port reset operation. The port enable will only be set in a standard EHCI controller implementation after the port reset operation and when the host and device negotiate a High-Speed connection (that is, Chirp completes successfully).

The module always sets the port enable after the port reset operation regardless of the result of the host device chirp result. The resulting port speed is indicated by the PSPD field in PORTSC. Therefore, the standard EHCI host controller driver requires an alteration to handle directly connected full- and

low-speed devices or hubs. [Table 17-97](#) summarizes the functional differences between EHCI and EHCI with embedded TT.

Table 17-97. Functional Differences Between EHCI and EHCI with Embedded TT

| Standard EHCI | EHCI with Embedded Transaction Translator |
|---|---|
| After port enable bit is set following a connection and reset sequence, the device/hub is assumed to be HS. | After port enable bit is set following a connection and reset sequence, the device/hub speed is noted from PORTSC. |
| FS and LS devices are assumed to be downstream from a HS hub thus, all port-level control is performed through the Hub Class to the nearest Hub. | FS and LS device can be either downstream from a HS hub or directly attached. When the FS/LS device is downstream from a HS hub, then port-level control is done using the Hub Class through the nearest Hub. When a FS/LS device is directly attached, then port-level control is accomplished using PORTSC. |
| FS and LS devices are assumed to be downstream from a HS hub with HubAddr=X. [where HubAddr > 0 and HubAddr is the address of the Hub where the bus transitions from HS to FS/LS (that is, Split target hub)] | FS and LS device can be either downstream from a HS hub with HubAddr = X [HubAddr > 0] or directly attached [where HubAddr = 0 and HubAddr is the address of the Root Hub where the bus transitions from HS to FS/LS (that is, Split target hub is the root hub)] |

17.9.1.4 Data Structures

The same data structures used for FS/LS transactions through a HS hub are also used for transactions through the Root Hub. The following list demonstrates how the Hub Address and Endpoint Speed fields should be set for directly attached FS/LS devices and hubs:

1. QH (for direct attach FS/LS)—Async. (Bulk/Control Endpoints) Periodic (Interrupt)
 - Hub Address = 0
 - Transactions to direct attached device/hub.
 - QH.EPS = Port Speed
 - Transactions to a device downstream from direct attached FS hub.
 - QH.EPS = Downstream Device Speed

NOTE

When QH.EPS = 01 (LS) and PORTSC[PSPD] = 00 (FS), a LS-pre-pid is sent before the transmitting LS traffic.

Maximum Packet Size must be less than or equal 64 or undefined behavior may result.

2. siTD (for direct attach FS)—Periodic (ISO Endpoint)
 - All FS ISO transactions:
 - Hub Address = 0
 - siTD.EPS = 00 (full speed)

Maximum Packet Size must less than or equal to 1023 or undefined behavior may result.

17.9.1.5 Operational Model

The operational models are well defined for the behavior of the transaction translator (see *Universal Serial Bus Revision 2.0 Specification*) and for the EHCI controller moving packets between system memory and a USB-HS hub. Since the embedded transaction translator exists within the DR module there is no physical

bus between EHCI host controller driver and the USB FS/LS bus. These sections will briefly discuss the operational model for how the EHCI and transaction translator operational models are combined without the physical bus between. The following sections assume the reader is familiar with both the EHCI and USB 2.0 transaction translator operational models.

17.9.1.5.1 Microframe Pipeline

The EHCI operational model uses the concept of H-frames and B-frames to describe the pipeline between the Host (H) and the Bus (B). The embedded transaction translator shall use the same pipeline algorithms specified in the *Universal Serial Bus Revision 2.0 Specification* for a Hub-based transaction translator.

All periodic transfers always begin at B-frame 0 (after SOF) and continue until the stored periodic transfers are complete. As an example of the microframe pipeline implemented in the embedded transaction translator, all periodic transfers that are tagged in EHCI to execute in H-frame 0 will be ready to execute on the bus in B-frame 0.

It is important to note that when programming the S-mask and C-masks in the EHCI data structures to schedule periodic transfers for the embedded transaction translator, the EHCI host controller driver must follow the same rules specified in EHCI for programming the S-mask and C-mask for downstream Hub-based transaction translators.

Once periodic transfers are exhausted, any stored asynchronous transfer are moved. Asynchronous transfers are opportunistic in that they shall execute whenever possible and their operation is not tied to H-frame and B-frame boundaries with the exception that an asynchronous transfer can not babble through the SOF (start of B-frame 0).

17.9.1.5.2 Split State Machines

The start and complete split operational model differs from EHCI slightly because there is no bus medium between the EHCI controller and the embedded transaction translator. Where a start or complete-split operation would occur by requesting the split to the HS hub, the start/complete split operation is simple an internal operation to the embedded transaction translator. [Table 17-98](#) summarizes the conditions where handshakes are emulated from internal state instead of actual handshakes to HS split bus traffic.

Table 17-98. Emulated Handshakes

| Condition | Emulate TT Response |
|---|--------------------------------------|
| Start-Split: All asynchronous buffers full | NAK |
| Start-Split: All periodic buffers full | ERR |
| Start-Split: Success for start of Async. Transaction | ACK |
| Start-Split: Start Periodic Transaction | No Handshake (Ok) |
| Complete-Split: Failed to find transaction in queue | Bus Time Out |
| Complete-Split: Transaction in Queue is Busy | NYET |
| Complete-Split: Transaction in Queue is Complete | [Actual Handshake from FS/LS device] |

17.9.1.5.3 Asynchronous Transaction Scheduling and Buffer Management

The following *Universal Serial Bus Revision 2.0 Specification* items are implemented in the embedded transaction translator:

- USB 2.0–11.17.3
 - Sequencing is provided and a packet length estimator ensures no full-speed/low-speed packet babbles into SOF time.
- USB 2.0–11.17.4
 - Transaction tracking for 2 data pipes.
- USB 2.0–11.17.5
 - Clear_TT_Buffer capability provided

17.9.1.5.4 Periodic Transaction Scheduling and Buffer Management

The following *Universal Serial Bus Revision 2.0 Specification* items are implemented in the embedded transaction translator:

- USB 2.0–11.18.6.[1-2]
 - Abort of pending start-splits
 - EOF (and not started in microframes 6)
 - Idle for more than 4 microframes
 - Abort of pending complete-splits
 - EOF
 - Idle for more than 4 microframes

NOTE

There is no data schedule mechanism for these transactions other than the microframe pipeline. The embedded TT assumes the number of packets scheduled in a frame does not exceed the frame duration (1 msec) or else undefined behavior may result.

17.9.1.5.5 Multiple Transaction Translators

The maximum number of embedded transaction translators that is currently supported is one as indicated by the N_TT field in the HCSPARAMS register. See [Section 17.3.1.3, “Host Controller Structural Parameters \(HCSPARAMS\),”](#) for more information.

17.9.2 Device Operation

The co-existence of a device operational controller within the DR module has little effect on EHCI compatibility for host operation except as noted in this section.

17.9.3 Non-Zero Fields the Register File

Some of the reserved fields and reserved addresses in the capability registers and operational registers have use in device mode, the following must be adhered to:

- Write operations to all EHCI reserved fields (some of which are device fields in the DR module) in the operation registers should always be written to zero. This is an EHCI requirement of the device controller driver that must be adhered to.
- Read operations by the module must properly mask EHCI reserved fields (some of which are device fields in the DR module registers).

17.9.4 SOF Interrupt

The SOF interrupt is a free running 125 μ sec interrupt for host mode. EHCI does not specify this interrupt, but it has been added for convenience and as a potential software time base. Note that the free running interrupt is shared with the device-mode start-of-frame interrupt. See [Section 17.3.2.2, “USB Status Register \(USBSTS\),”](#) and [Section 17.3.2.3, “USB Interrupt Enable Register \(USBINTR\),”](#) for more information.

17.9.5 Embedded Design

This is an Embedded USB Host Controller as defined by the EHCI specification and thus does not implement the PCI configuration registers.

17.9.5.1 Frame Adjust Register

Given that the optional PCI configuration registers are not included in this implementation, there is no corresponding bit level timing adjustments like those provided by the Frame Adjust register in the PCI configuration registers. Starts of microframes are timed precisely to 125 μ sec using the transceiver clock as a reference clock. That is, 60 MHz transceiver clock for 8-bit physical interfaces and full-speed serial interfaces or 30 MHz transceiver clock for 16-bit physical interfaces.

17.9.6 Miscellaneous Variations from EHCI

The modules support multiple physical interfaces which can operate in different modes when the module is configured with the software programmable Physical Interface Modes. The control bits for selecting the PHY operating mode have been added to the PORTSC register providing a capability that is not defined by the EHCI specification.

17.9.6.1 Discovery

This section discusses port reset and port speed detection.

17.9.6.1.1 Port Reset

The port connect methods specified by EHCI require setting the port reset bit in the register for a duration of 10 msec. Due to the complexity required to support the attachment of devices that are not high speed

there are counter already present in the design that can count the 10 msec reset pulse to alleviate the requirement of the software to measure this duration. Therefore, the basic connection is then summarized as the following:

- [Port Change Interrupt] Port connect change occurs to notify the host controller driver that a device has attached.
- Software shall write a '1' to the reset the device.
- Software shall write a '0' to the reset the device after 10 msec.
 - This step, which is necessary in a standard EHCI design, may be omitted with this implementation. Should the EHCI host controller driver attempt to write a '0' to the reset bit while a reset is in progress the write will simple be ignored and the reset will continue until completion.
- [Port Change Interrupt] Port enable change occurs to notify the host controller that the device in now operational and at this point the port speed has been determined.

17.9.6.1.2 Port Speed Detection

After the port change interrupt indicates that a port is enabled, the EHCI stack should determine the port speed. Unlike the EHCI implementation which will re-assign the port owner for any device that does not connect at High-Speed, this host controller supports direct attach of non-HS devices. Therefore, the following differences are important regarding port speed detection:

- Port owner is read-only and always reads 0.
- A 2-bit port speed indicator has been added to PORTSC to provide the current operating speed of the port to the host controller driver.
- A 1-bit high-speed indicator has been added to PORTSC to signify that the port is in HS vs. FS/LS

17.10 Timing Diagrams

This section contains diagrams showing the basic operation of the ULPI interface. For a more detailed description refer to the ULPI Specifications.

Figure 17-66 shows ULPI timing.




Figure 17-66. ULPI Timing

Table 17-99 summarizes the ULPI timing parameters.

Table 17-99. ULPI Timing

| Parameter | Symbol | Min | Max | Units |
|---------------------------|-----------------|-----|-----|-------|
| Control signal setup time | T _{SC} | — | 4 | ns |
| Data setup time | T _{SD} | — | 4 | ns |
| Control signal hold time | T _{HC} | 0 | — | ns |
| Data hold time | T _{HD} | 0 | — | ns |
| Control output delay | T _{DC} | 2 | 7 | ns |
| Data output delay | T _{DD} | 2 | 7 | ns |

Figure 17-67 shows the diagram for the sending of RX CMD.

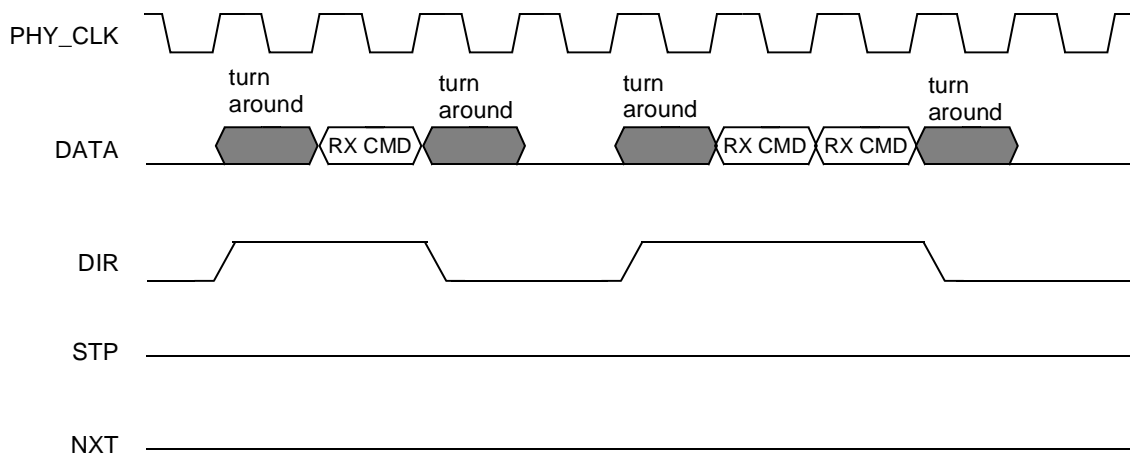


Figure 17-67. Sending of RX CMD

Figure 17-68 shows ULPI data transmit—NOPID.

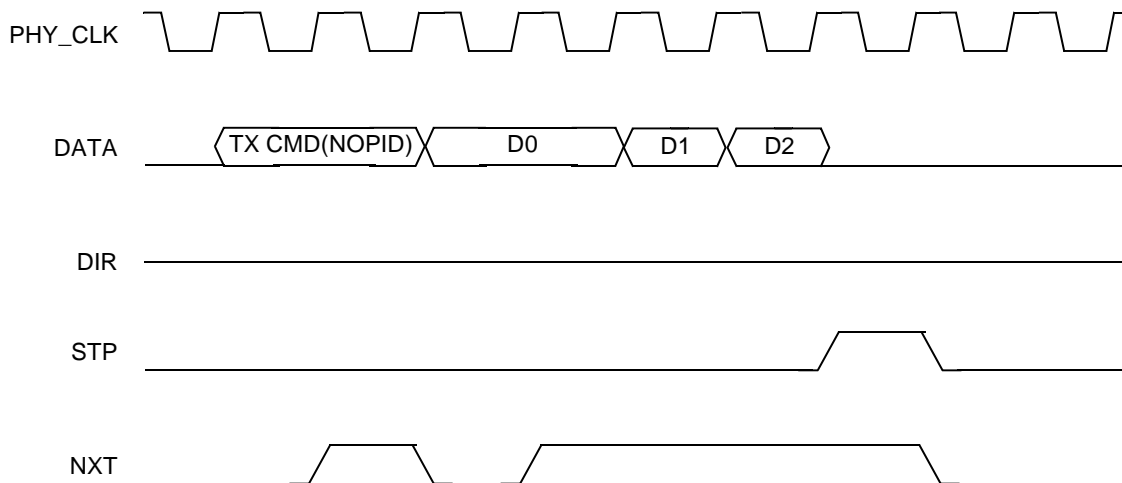


Figure 17-68. ULPI Data Transmit (NOPID)

Figure 17-69 shows ULPI data transmit—PID.

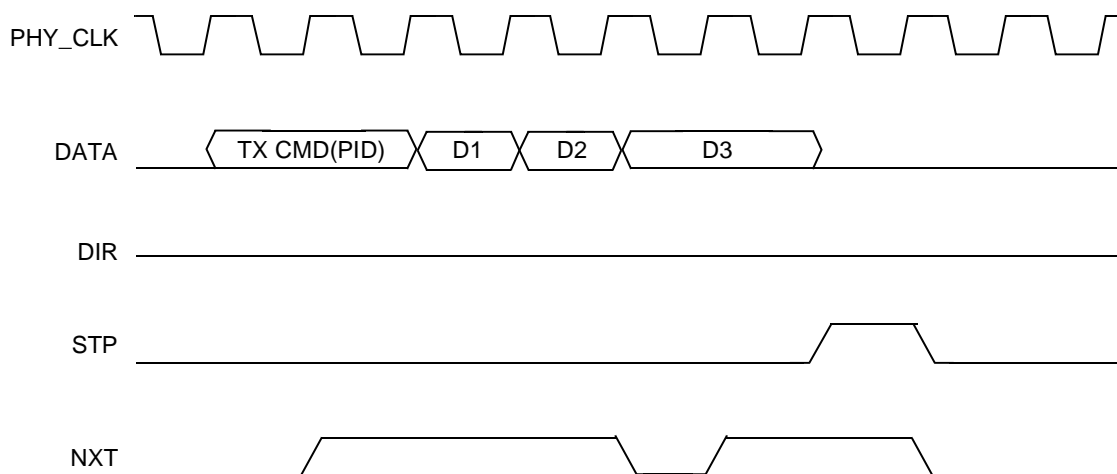


Figure 17-69. ULPI Data Transmit (PID)

Figure 17-70 shows ULPI data receive.

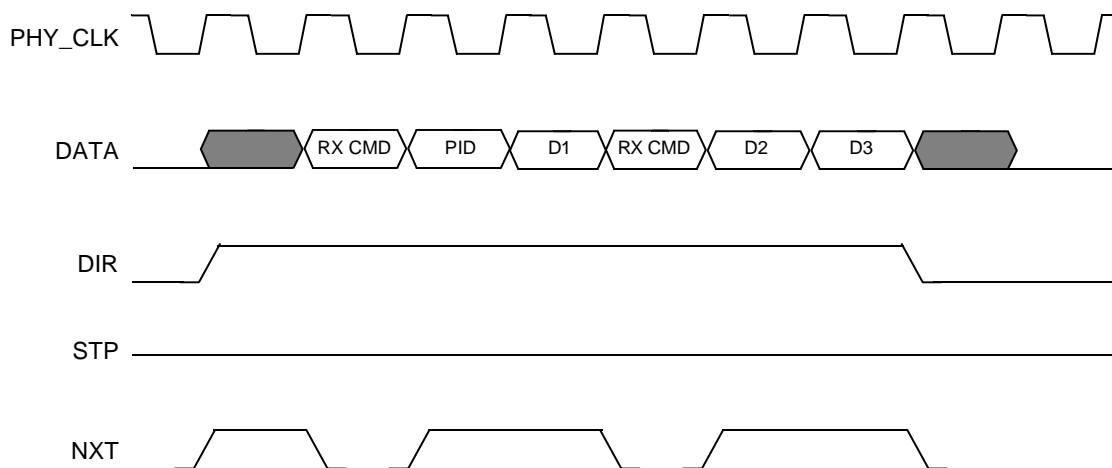


Figure 17-70. ULPI Data Receive

Figure 17-71 shows ULPI register write.




Figure 17-71. ULPI Register Write

Figure 17-72 shows ULPI register read.




Figure 17-72. ULPI Register Read



Chapter 18

Security Engine (SEC) 3.3

This chapter describes the functionality of the device's integrated security engine (SEC 3.3). It addresses the following topics:

- [Section 18.1, “SEC 3.3 Architecture Overview”](#)
- [Section 18.2, “Configuration of Internal Memory Space”](#)
- [Section 18.3, “Descriptors”](#)
- [Section 18.5, “Polychannel”](#)
- [Section 18.6, “Controller”](#)
- [Section 18.6.1, “Bus Transfers”](#)
- [Section 18.8, “Execution Units”](#)

The SEC 3.3 is designed to off-load computationally intensive security functions, such as key generation and exchange, authentication, and bulk encryption from the processor core of the device. It is optimized to process cryptographic algorithms associated with IPsec, IKE, SSL/TLS, iSCSI, SRTP, 802.11i, 802.16 (WiMAX), and 802.1AE (MACSec). The SEC 3.3 is derived from integrated security cores found in other members of the PowerQUICC II Pro family.

The security engine includes six different execution units (EUs). Where data flows in and out of an EU, each has buffer FIFOs of at least 256 bytes. EU types and features include the following:

- AESU—Advanced Encryption Standard Unit
 - Implements the Rijndael symmetric key cipher
 - Modes providing data confidentiality: ECB, CBC, CCM, Counter, GCM, CBC-RBP, OFB-128, and CFB-128.
 - Modes providing data authentication: CCM, GCM, CMAC, and XCBC-MAC.
 - 128, 192, 256 bit key lengths (only 128 bit keys in XCBC-MAC)
 - ICV checking in CCM, GCM, CMAC, and XCBC-MAC mode
 - XOR operations on 2 - 6 sources for RAID 5
- CRCU—Cyclical Redundancy Check Unit
 - Implements CRC32C as required for iSCSI header and payload checksums, CRC32 as required for IEEE Std. 802 packets, and programmable CRC modes
 - ICV checking
- DEU—Data Encryption Standard Execution Unit
 - DES, 3DES
 - Two key (K1, K2, K1) or Three Key (K1, K2, K3)
 - ECB, CBC, CFB-64 and OFB-64 modes for both DES and 3DES

- MDEU—Message Digest Execution Unit
 - Implements SHA-1 with 160-bit, 224-bit, 256-bit, 384-bit, and 512-bit message digest (as specified by the FIPS 180-2 standard)
 - Implements MD5 with 128-bit message digest (as specified by RFC 1321)
 - Implements HMAC computation with either message digest algorithm (as specified in RFC 2104 and FIPS-198)
 - Implements SSL MAC computation
 - ICV checking
- PKEU—Public Key Execution Unit
 - RSA and Diffie-Hellman
 - Programmable field size up to 4096 bits (increased from 2048)
 - Elliptic curve cryptography
 - F_{2^m} and F_p modes
 - Programmable field size up to 1023 bits (increased from 511)
 - Run time equalization to protect against timing and power attacks
- RNGU—Random Number Generator
 - Combines a True Random Number Generator (TRNG) and NIST-approved Pseudo-Random Number Generator (PRNG) (as described in Annex C of FIPS140-2 and ANSI X9.62).

In addition to the execution units, SEC 3.3 also includes:

- Four-channel operation, where each channel:
 - Supports a queue of commands (descriptor pointers) to be executed
 - Provides dynamic arbitration for needed crypto-execution units
 - Manages up to two execution units (one ciphering and one hashing), and configures for any required data transfers from one to another
 - Performs flow-control management of buffer FIFOs on the inputs and outputs of execution units
 - Supports scatter/gather of input and output data, enabling concatenation of multiple segments of memory when reading or writing data
 - Fetches data bursts on 32-byte boundaries to optimize bus throughput
- Master and slave interfaces, with DMA capability
 - 32 or 36 bit address/64 -bit data
 - Operates at DDR clock rate
 - Master interface allows multiple pipelined requests
 - DMA data blocks can start and end on any byte boundary

18.1 SEC 3.3 Architecture Overview

The SEC can act as a master on the internal system bus, allowing it to off-load the data movement bottleneck normally associated with slave-only cores. The host processor accesses the SEC through its

device drivers using system memory for data storage. The SEC resides in the peripheral memory map of the processor. When an application requires cryptographic functions, it simply creates descriptors for the SEC, which define the cryptographic function to be performed and the locations of the data. With a single short write, the host processor can enqueue a descriptor pointer in the SEC. The SEC's bus-mastering capability then enables it to execute the entire cryptographic task, performing reads and writes on system memory as needed.

A block diagram of the SEC internal architecture is shown in [Figure 18-1](#).

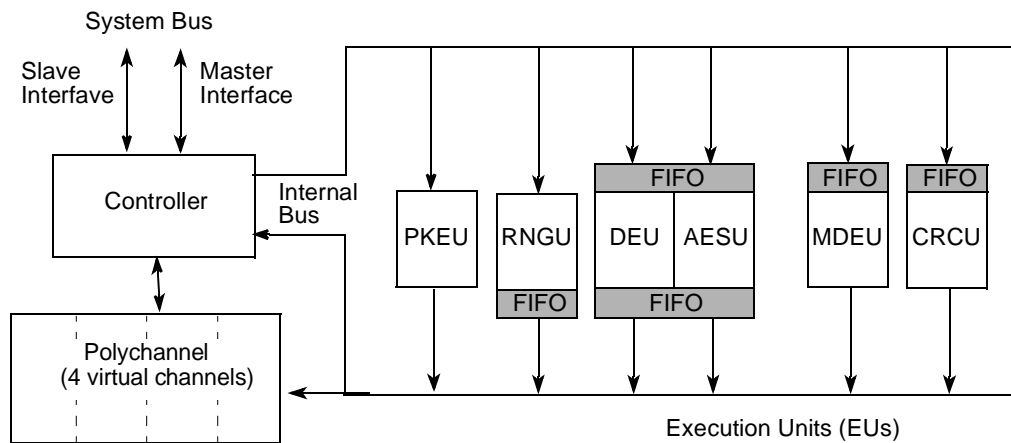


Figure 18-1. SEC Functional Modules

The SEC interfaces with the system buses through the controller. Using the slave interface, the host processor can do 64-bit reads or writes on any register or FIFO inside the SEC. Using the master interface, the controller can transfer blocks of 64-bit words between system memory and SEC FIFOs or registers.

A typical SEC operation begins when the host writes a descriptor pointer to the fetch FIFO in one of the four SEC virtual channels. This write operation uses the slave interface (where the host is master and SEC is the slave). From this point on, the channel directs the sequence of operations, using the master interface (where SEC is master). The channel uses the descriptor pointer to read the descriptor, then decodes the first word of the descriptor to determine the operation to be performed and the crypto-execution units needed to perform it. If necessary, the channel waits for the needed crypto-execution units to be free. Next, the channel requests that the controller transfer keys, context and data from memory locations specified in the descriptor to the appropriate execution units. The controller satisfies the requests through its master interface. Data is fed into the execution units through their registers and input FIFOs. The execution units read from their input FIFOs and write processed data to their output FIFOs and registers. The channel requests the controller to write data from the output FIFOs and registers back to system memory.

The channel can signal to the host that it is done with a descriptor via interrupt or by a writeback of the descriptor header into host memory. For more about this signaling, see [Section 18.1.2, “Polychannel Overview.”](#)

Upon completion of a descriptor, the channel checks the next entry in its fetch FIFO, and, if non-zero, the channel requests a read of the next descriptor.

For most packets, the entire payload is too long to fit in an execution unit’s input or output FIFO, so the channel uses a flow control scheme for reading and writing data. The channel directs the controller to read bursts of input as necessary to keep refilling the input FIFO, until the entire payload has been fetched. Similarly, the channel directs the controller to write bursts of output whenever enough accumulates in the execution unit’s output FIFO.

The Polychannel can process up to four descriptors concurrently by implementing four virtual channels. Channels arbitrate for use of execution units, and wait if the needed execution unit is currently reserved by another channel. Each channel has its own FIFO of descriptor pointers to fetch and execute, and its own internal storage. The four channels, however, time-share a single control and datapath unit, and hence they are referred to as virtual channels. A programmable priority scheme allows for round-robin or weighted priorities among these channels.

18.1.1 Descriptor Overview

The SEC controller has been designed for easy use and integration with existing systems and software. All cryptographic functions are accessible through descriptors. A descriptor specifies a cryptographic function to be performed, and contains reference address pointers to all necessary input data and to the places where output data is to be written. Some descriptor types perform multiple functions to facilitate particular protocols. A sample descriptor is diagrammed in [Table 18-1](#).

Each descriptor contains eight dwords (64 bits each), consisting of the following:

- One dword of header—The header describes the required services and encodes information that indicates which EUs to use and which modes to set. It also indicates whether notification should be sent to the host when the descriptor operation is complete.
- Seven dwords containing pointers and lengths used to locate input or output data. Each pointer can either point directly to the data, or can point to a link table that lists a set of data segments to be concatenated.

Table 18-1. Example Descriptor

| Field Name | Value | Description |
|--------------------------------|-----------------------------------|--|
| Header | 0x2053_1E08_0000_0000 | Example header for IPsec ESP outbound using DES and MD-5 |
| Length0 Extent0 Pointer0 | 16 0 (32 or 36-bit pointer) | Number of bytes in authenticate key Unused Pointer to authentication key |
| Length1 Extent1 Pointer1 | 16 0 (32 or 36-bit pointer) | Number of bytes in authentication-only data Unused Pointer to authentication-only data |
| Length2 Extent2 Pointer2 | 8 0 (32 or 36-bit pointer) | Length of input context (IV) Unused Pointer to input context |
| Length3 Extent3 Pointer3 | 8 0 (32 or 36-bit pointer) | Number of bytes in cipher key Unused Pointer to cipher key |

Table 18-1. Example Descriptor (continued)

| Field Name | Value | Description |
|--------------------------------|--------------------------------------|--|
| Length4 Extent4 Pointer4 | 1500 0 (32 or 36-bit pointer) | Number of bytes of data to be ciphered Unused Pointer to input data to perform ciphering upon |
| Length5 Extent5 Pointer5 | 1500 12 (32 or 36-bit pointer) | Number of bytes of data after ciphering Number of bytes in authentication result (ICV) Pointer to location where cipher output is to be written, followed by ICV |
| Length6 Extent6 Pointer6 | 8 0 (32 or 36-bit pointer) | Length of output context (IV) Unused Pointer to location where altered Context is to be written |

For more information, refer to [Section 18.3, “Descriptors.”](#)

18.1.2 Polychannel Overview

The polychannel block implements four channels for processing descriptors. Each channel contains the following addressable structures:

- A fetch FIFO, which holds a queue of pointers to descriptors waiting to be processed
- A configuration register, which allows the user a number of options for SEC event signalling
- Monitor registers containing information about the transaction in process
- A status register containing an indication of the last unfulfilled bus request
- A descriptor buffer memory used to store the active descriptor and other temporary

Whenever a channel is idle and its fetch FIFO is non-empty, the channel reads the next descriptor pointer from the fetch FIFO. Using this pointer, the channel fetches the descriptor and places it in its descriptor buffer. The channel’s processing of descriptors is described in more detail in #####”

A channel can signal to the host that it is done with a descriptor via interrupt and/or by a writeback of the descriptor header into host memory. In the case of writeback, the value written back is identical to the header that was read, with the exception that a DONE byte is set to 0xFF. The channels’ done signaling is described in more detail in #####”.

An EU operation can include generating an ICV and then comparing it against a received ICV. The result of the ICV checking can be signalled to the host either via interrupt or by a writeback of the descriptor header (but not by both methods). In the case of writeback, the user can opt to do it at end of every descriptor, or only at the end of descriptors that call for ICV checking.

In case of an error condition in a channel or its reserved EUs, the channel issues an interrupt to the host. The channel can be configured to either abort the current descriptor and proceed to the next one, or halt and wait for host intervention.

For more about configuring signalling see [Section 18.5.4.1, “Channel Configuration Register \(CCR\),”](#) and for detail on the writeback fields see [Section 18.3.4, “Link Table Format.”](#)

Many security protocols involve both encryption and hashing of packet payloads. To accomplish this without requiring two passes through the data, channels can configure data flows through two EUs. In such

cases, one EU is designated the “primary EU”, and the other as the “secondary EU”. The primary EU receives its data from memory via the controller, and the secondary EU receives its data by “snooping” the SEC buses.

There are two types of snooping:

- Input data can be fed to the primary EU and the same input data snooped by the secondary EU. This is called “in-snooping”.
- Output data from the primary EU can be snooped by the secondary EU. This is called “out-snooping”.

In the SEC, only MDEU and CRCU are used as secondary EUs.

For more information, refer to [Section 18.5, “Polychannel.”](#)

18.1.3 Controller Overview

The controller manages the master and slave interfaces to the system bus and the internal buses that connect all the various modules. It receives service requests from the host (via the slave interface) and from the channels, and schedules the required data transfers. The system bus interface and access to system memory are critical factors in performance, and the 64-bit master and slave interfaces of the SEC controller enable it to achieve performance unattainable on secondary buses.

The controller provides for two ways of operating the execution units, channel-controlled access and host-controlled access.

- In channel-controlled access (the SEC’s normal operating mode), all interactions with EUs are directed by a channel executing a descriptor. The host is involved only in initially supplying the descriptor pointer and in handling results once descriptor processing is complete.
- In host-controlled access (intended primarily for debug purposes), the host moves data in and out of execution units directly through memory-mapped EU registers. No descriptor is involved.

18.1.4 Execution Units (EUs) Overview

“Execution unit” (EU) is the generic term for a functional block that performs the cryptographic computations. The EUs are compatible with many protocols, and can work together to perform high-level cryptographic tasks. The SEC’s execution units are as follows:

- PKEU for computing asymmetric key operations, including modular exponentiation (and other modular arithmetic functions) or ECC point arithmetic
- DEU for performing block cipher, symmetric key cryptography using DES and 3DES
- AESU for performing the Advanced Encryption Standard algorithm in various modes
- MDEU for performing security hashing using MD-5, SHA-1, SHA-224, SHA-256, SHA-384 or SHA-512
- CRCU for generating cyclical redundancy check values
- RNGU for random number generation

The following sections give an overview of the EUs. Operational details of each EU are given in [Section 18.8, “Execution Units.”](#)

18.1.4.1 Public Key Execution Unit (PKEU)

The PKEU is capable of performing many advanced mathematical functions to support both RSA and ECC public key cryptographic algorithms. ECC is supported in both F_{2^m} (polynomial field) and F_p (prime field) modes.

To assist the host in performing its desired cryptographic functions, the PKEU supports functions with various levels of complexity. For example, at the highest level, the accelerator performs modular exponentiations to support RSA and performs point multiplies to support ECC. At a lower level, the PKEU can perform simple operations such as modular adds and multiplies. For more information, refer to [Section 18.8.5, “Public Key Execution Units \(PKEU\).”](#)

18.1.4.1.1 Elliptic Curve Operations

The PKEU has its own data and control units, including a general-purpose register file in the programmable-size arithmetic unit. The field or modulus size can be programmed to any value between 33 bits and 1024 bits in programmable increments of 8, with each programmable value i supporting all actual field sizes from $8i - 7$ to $8i$. The result is hardware supporting a wide range of cryptographic security. Larger field / modulus sizes result in greater security but lower performance.

Compared to RSA, elliptic curve cryptography provides greater security with smaller field sizes. For example, an elliptic curve field size of 160 is roughly equivalent to the security provided by 1024 bit RSA. A field size set to 224 roughly equates to 2048 bits of RSA security.

The PKEU contains routines implementing the atomic functions for elliptic curve processing—point arithmetic and finite field arithmetic. The point operations (multiplication, addition and doubling) involve one or more finite field operations which are addition, multiplication, inverse, and squaring. Point add and double each use of all four finite field operations. Similarly, point multiplication uses all elliptic curve point operations as well as the finite field operations. All these functions are supported both in prime fields and polynomial fields.

18.1.4.1.2 Modular Exponentiation Operations

The PKEU is also capable of performing integer modulo arithmetic. This arithmetic is an integral part of the RSA public key algorithm; however, it can also play a role in the generation of ECC digital signatures (including ECDSA) and Diffie-Hellman key exchanges.

Modular arithmetic functions supported by the SEC’s PKEU include the following:

- $R^2 \bmod N$
- $(A \times B) R^{-1} \bmod N$
- $(A \times B) R^{-2} \bmod N$
- $(A + B) \bmod N$
- $(A - B) \bmod N$

(where N is the modulus vector, A and B are input vectors, E is the exponent vector, and R is $2^{Sz'(N)}$, where $Sz'(N)$ is the bit length of the N vector rounded up to the nearest multiple of 32.

The PKEU can perform modular arithmetic on operands up to 4096 bits in length. The modulus must be larger than or equal to 33 bits (5 bytes). This is not seen as a limitation since for sizes smaller than this, no useful cryptographic application exists. Furthermore, for small data sizes the overhead of using a hardware accelerator would not be justified. The PKEU uses the Montgomery modular multiplication algorithm to perform core functions. The addition and subtraction functions help support known methods of the Chinese Remainder Theorem (CRT) for efficient implementation of the RSA algorithm.

18.1.4.2 Data Encryption Standard Execution Unit (DEU)

The DES Execution Unit (DEU) performs bulk data encryption/decryption, in compliance with the Data Encryption Standard algorithm (NIST FIPS 46-3). The DEU can also compute 3DES, an extension of the DES algorithm in which each 64-bit input block is processed three times. The SEC supports 2 key ($K1 = K3$) or 3 key 3DES.

The DEU operates by permuting 64-bit data blocks with a shared 56-bit key and an initialization vector (IV). The SEC supports four modes of operation:

- Electronic Code Book (ECB)
- Cipher Block Chaining (CBC)
- 64-bit Cipher Feedback Mode (CFB-64)
- 64-bit Output Feedback Mode (OFB-64).

For more information, refer to [Section 18.8.3, “Data Encryption Standard Execution Unit \(DEU\).”](#)

18.1.4.3 Advanced Encryption Standard Execution Unit (AESU)

The AESU is used to accelerate bulk data encryption/decryption in compliance with the Advanced Encryption Standard algorithm Rijndael specified by NIST standard FIPS-197. The AESU executes on 128 bit blocks with a choice of key sizes: 128, 192, or 256 bits.

AESA is a symmetric key algorithm, meaning the sender and receiver use the same key for encryption and decryption. The session key and IV are supplied to the AESU module prior to encryption. The processor supplies data to the module that is processed as 128 bit input.

AESU implements the following confidentiality Modes from NIST Recommendation 800-38A:

- Electronic Codebook mode (ECB)
- Cipher Block Chaining mode (CBC)
- Output Feedback mode (OFB)
- 128-bit Cipher Feedback mode (CFB-128)
- Counter mode (CTR)

AESU also implements other NIST recommended modes providing Authentication (two of which also provide confidentiality):

- Counter with CBC-MAC (CCM) per NIST recommendation 800-38C

- Galois Counter Mode (GCM) per NIST draft recommendation 800-38D
- Cipher-based MAC (CMAC) per NIST recommendation 800-38B

AESU modes also implement the following modes not sanctioned by NIST:

- CBC-RBP (Residual Block Processing from 802.16)
- XCBC-MAC as specified by IETF RFC-3566

In all modes supporting Authentication, the AESU hashes data to produce an integrity check vector (ICV). If a reference ICV is supplied to the AESU, it can do a bitwise check of this supplied ICV against the one computed by the AESU (ICV checking).

For more information, refer to [Section 18.8.1, “Advanced Encryption Standard Execution Unit \(AESU\).”](#)

18.1.4.4 Message Digest Execution Unit (MDEU)

The MDEU computes a single message digest (or hash or integrity check) value of all the data presented on the input bus, using either the MD5, SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 algorithms for bulk data hashing.

With any hash algorithm, the larger message is mapped onto a smaller output “digest”, so it is possible for two messages to produce the same digest. The hash digest lengths (128 bits or more) are sufficiently large, however, that such collisions are extremely rare. The security of the hash function is based on the difficulty of locating collisions. That is, it is computationally infeasible to construct two distinct but similar messages that produce the same hash output.

- The MD5 generates a 128-bit hash, and the algorithm is specified in RFC 1321.
- SHA-1 is a 160-bit hash function, specified by the NIST FIPS 180-1 standard.
- SHA-224 is a 224-bit hash function, specified by the NIST FIPS 180-2 standard.
- SHA-256 is a 256-bit hash function that provides 256 bits of security against collision attacks, specified by the NIST FIPS 180-2 standard.
- SHA-384 is a 384-bit hash function, specified by the NIST FIPS 180-2 standard.
- SHA-512 is a 512-bit hash function, specified by the NIST FIPS 180-2 standard.
- The MDEU also supports HMAC computations, as specified by the NIST FIPS-198 standard.

If a digest is supplied to the MDEU, it can do a bitwise check of this supplied digest against the one computed by the MDEU (ICV checking).

For more information, refer to [Section 18.8.4, “Message Digest Execution Unit \(MDEU\).”](#)

18.1.4.5 Cyclical Redundancy Check Unit (CRCU)

The CRCU computes a single 32-bit cyclic redundancy code (checksum) from all data presented on the input bus.

The CRC algorithm treats a message stream of bits as coefficients of a massive polynomial and computes the remainder of the modulo two division by an order 32 divisor polynomial. The divisor polynomial is specific to the protocol and chosen to conform to certain mathematical properties to ensure that single bit errors can be detected. Cyclic redundancy codes are used to ensure data integrity over potentially unreliable

channels. There are two major CRC protocol algorithms: CRC32 and CRC32C. IEEE 802 defines the CRC32 algorithm, while iSCSI defines the CRC32C algorithm. Both protocols bit swap, byte swap, and then complement the calculated remainder to generate the checksum.

- The CRC32 algorithm is specified in IEEE 802.1.
- The CRC32C algorithm is specified in RFC3385.
- The CRCU also supports a programmable polynomial mode with no remainder bit mangling. This can be used to implement proprietary protocols.

If a CRC is supplied to the CRCU, it can do a bitwise check of this supplied checksum against the one computed by the CRCU (ICV checking).

For more information, refer to [Section 18.8.2, “Cyclical Redundancy Check Unit \(CRCU\).”](#)

18.1.4.6 Random Number Generator (RNGU)

The RNGU is a functional block capable of generating 64-bit random numbers. It is designed to comply with FIPS 140-1 standards for randomness and non-determinism.

Because many cryptographic algorithms use random numbers as a source for generating a secret value (a nonce), it is desirable to have a private RNG for use by the SEC. The anonymity of each random number must be maintained, as well as the unpredictability of the next random number. The FIPS-140 ‘common criteria’ compliant private RNG allows the system to develop random challenges or random secret keys. The secret key can thus remain hidden from even the high-level application code, providing an added measure of physical security.

For more information, refer to [Section 18.8.6, “Random Number Generator \(RNGU\).”](#)

18.2 Configuration of Internal Memory Space

[Table 18-2](#) is the base address map, showing the blocks of addresses assigned to each SEC sub-block. [Table 18-3](#) is the detailed address map, including all registers in each sub-block. The 18-bit SEC address bus value is shown. These address values are offsets from the SoC’s base address register (consult the SoC documentation for specific register name).

Table 18-2. SEC Address Map

| Byte Address Offset (AD 17–0) | Module | Description | Type | Section/Page |
|-------------------------------|------------|---|------------|--------------------------------|
| 0x3_1000–0x3_10FF | Controller | Arbiter/controller control register space | Controller | 18.6/18-41 |
| 0x3_1100–0x3_11FF | Channel_1 | Channel 1 | Channels | 18.5/18-29 |
| 0x3_1200–0x3_12FF | Channel_2 | Channel 2 | | |
| 0x3_1300–0x3_13FF | Channel_3 | Channel 3 | | |
| 0x3_1400–0x3_14FF | Channel_4 | Channel 4 | | |
| 0x3_1500–0x3_16FF | PolyChn | PolyChannel Memory | | |
| 0x3_1BF8 | Controller | IP Block Revision Register | Read only | 18.6.4.6/18-50 |
| 0x3_2000–0x3_2FFF | DEU | DES/3DES Execution Unit | Crypto EU | 18.8.3/18-89 |
| 0x3_4000–0x3_4FFF | AESU | AES Execution Unit | | 18.8.1/18-53 |
| 0x3_6000–0x3_6FFF | MDEU | Message Digest Execution Unit | | 18.8.4/18-97 |
| 0x3_8000–0x3_8FFF | — | Reserved | | — |
| 0x3_A000–0x3_AFFF | RNGU | Random Number Generator | | 18.8.6/18-119 |
| 0x3_C000–0x3_CFFF | | Public Key Execution Unit | | 18.8.5/18-111 |
| 0x3_E000–0x3_EFFF | — | Reserved | | — |
| 0x3_F000–0x3_FFFF | CRCU | Cyclical Redundancy Check Accelerator | | 18.8.6/18-119 |

All regions not listed are reserved for future use.

[Table 18-3](#) shows the system address map showing all functional registers.

All SEC registers are allocated 8 bytes (one dword), and the addresses listed in the table are all at 8-byte boundaries (that is, ending in 0 or 8). It is possible, however, to access the registers by 4-byte word, or in some cases by byte. The “Write by” column distinguishes these cases. Possible entries in this column are:

- **Byte:** This register can be written by byte (using any address), by word (using an address ending in 0, 4, 8, or C), or by dword (using an address ending in 0 or 8).
- **Word:** This register can be written by dword (using an address ending in 0 or 8), or by word (using an address ending in 0, 4, 8, or C), but not by byte.

Reads can always be done by byte, word, or dword.

Table 18-3. SEC Address Map

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|----------------------------|----------------------------|------------------------|----------|--------------------------------|
| Byte Address Offset (AD 17–0) | Module | Register | Access | Write By | Section/Page |
| 0x3_1008 | Controller | Interrupt Enable | R/W | Byte | 18.6.4.2/18-45 |
| 0x3_1010 | | Interrupt Status | R | — | 18.6.4.3/18-47 |
| 0x3_1018 | | Interrupt Clear | R/W | Byte | 18.6.4.4/18-48 |
| 0x3_1020 | | Identification | R | — | 18.6.4.5/18-49 |
| 0x3_1028 | | EU Assignment Status | R | — | 18.6.4.1/18-44 |
| 0x3_1030 | | Master Control | R/W | Byte | 18.6.4.7/18-51 |
| 0x3_1108 | | Channel_1 | Configuration Register | R/W | Word |
| 0x3_1110 | Pointer Status | | R/W | — | 18.5.4.2/18-36 |
| 0x3_1140 | Current Descriptor Pointer | | R | — | 18.5.4.3/18-38 |
| 0x3_1148 | Fetch FIFO | | W | Word | 18.5.4.4/18-39 |
| 0x3_1180–0x3_11BF | Descriptor Buffer | | R | — | 18.5.5.1/18-40 |
| 0x3_11C0–0x3_11DF | Gather Link Tables | | W | Byte | 18.5.5.2/18-40 |
| 0x3_11E0–0x3_11FF | Scatter Link Tables | | W | Byte | 18.5.5.2/18-40 |
| 0x3_1208 | Channel_2 | Config Register | R/W | Word | 18.5.4.1/18-34 |
| 0x3_1210 | | Pointer Status | R/W | — | 18.5.4.2/18-36 |
| 0x3_1240 | | Current Descriptor Pointer | R | — | 18.5.4.3/18-38 |
| 0x3_1248 | | Fetch FIFO | W | Word | 18.5.4.4/18-39 |
| 0x3_1280–0x3_12BF | | Descriptor Buffer | R | — | 18.5.5.1/18-40 |
| 0x3_12C0–0x3_12DF | | Gather Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x3_12E0 - 0x3_12FF | | Scatter Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x3_1308 | Channel_3 | Config Register | R/W | Word | 18.5.4.1/18-34 |
| 0x3_1310 | | Pointer Status | R/W | — | 18.5.4.2/18-36 |
| 0x3_1340 | | Current Descriptor Pointer | R | — | 18.5.4.3/18-38 |
| 0x3_1348 | | Fetch FIFO | W | Word | 18.5.4.4/18-39 |
| 0x3_1380–0x3_13BF | | Descriptor Buffer | R | — | 18.5.5.1/18-40 |
| 0x3_13C0–0x3_13DF | | Gather Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x3_13E0–0x3_13FF | | Scatter Link Tables | W | Byte | 18.5.5.2/18-40 |

Table 18-3. SEC Address Map (continued)

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|------------|----------------------------|------------------|----------|---------------------------------|
| Byte Address Offset (AD 17–0) | Module | Register | Access | Write By | Section/Page |
| 0x3_1408 | Channel_4 | Config Register | R/W | Word | 18.5.4.1/18-34 |
| 0x3_1410 | | Pointer Status | R/W | — | 18.5.4.2/18-36 |
| 0x3_1440 | | Current Descriptor Pointer | R | — | 18.5.4.3/18-38 |
| 0x3_1448 | | Fetch FIFO | W | Word | 18.5.4.4/18-39 |
| 0x3_1480–0x3_14BF | | Descriptor Buffer | R | — | 18.5.5.1/18-40 |
| 0x3_14C0–0x3_14DF | | Gather Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x3_14E0–0x3_14FF | | Scatter Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x3_1BF8 | Controller | IP block revision | R | — | 18.6.4.6/18-50 |
| 0x3_2000 | DEU | Mode Register | R/W | Word | 18.8.3.1/18-89 |
| 0x3_2008 | | Key Size Register | R/W | Word | 18.8.3.2/18-90 |
| 0x3_2010 | | Data Size Register | R/W | Word | 18.8.3.3/18-90 |
| 0x3_2018 | | Reset Control Register | R/W | Word | 18.8.3.4/18-91 |
| 0x3_2028 | | Status Register | R | — | 18.8.3.5/18-92 |
| 0x3_2030 | | Interrupt Status Register | R/W | Word | 18.8.3.6/18-92 |
| 0x3_2038 | | Interrupt Mask Register | R/W | Word | 18.8.3.7/18-94 |
| 0x3_2050 | | EU-Go | W | Word | 18.8.3.8/18-96 |
| 0x3_2100 | | IV Register | R/W | Word | 18.8.3.9/18-96 |
| 0x3_2400 | | Key 1 Register | W | Byte | 18.8.3.10/18-96 |
| 0x3_2408 | | Key 2 Register | W | Byte | 18.8.3.10/18-96 |
| 0x3_2410 | | Key 3 Register | W | Byte | 18.8.3.10/18-96 |
| 0x3_2800–0x3_2FFF | | Input FIFO / Output FIFO | R/W ¹ | Byte | 18.8.3.11/18-97 |

Table 18-3. SEC Address Map (continued)

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|--------|---------------------------|------------------|----------|-----------------------------------|
| Byte Address Offset (AD 17–0) | Module | Register | Access | Write By | Section/Page |
| 0x3_4000 | AESU | Mode Register | R/W | Word | 18.8.1.2/18-54 |
| 0x3_4008 | | Key Size Register | R/W | Word | 18.8.1.3/18-57 |
| 0x3_4010 | | Data Size Register | R/W | Word | 18.8.1.4/18-57 |
| 0x3_4018 | | Reset Control Register | R/W | Word | 18.8.1.5/18-58 |
| 0x3_4028 | | Status Register | R | — | 18.8.1.6/18-58 |
| 0x3_4030 | | Interrupt Status Register | R/W | Word | 18.8.1.7/18-59 |
| 0x3_4038 | | Interrupt Mask Register | R/W | Word | 18.8.1.8/18-61 |
| 0x3_4050 | | End Of Message Register | W | Word | 18.8.1.10/18-63 |
| 0x3_4100–0x3_415F | | Context | R/W | Byte | 18.8.1.11/18-64 |
| 0x3_4400–0x3_441F | | Key Memory | R/W | Byte | 18.8.1.11.8/18-79 |
| 0x3_4800–0x3_4FFF | | Input FIFO / Output FIFO | R/W ¹ | Byte | 18.8.1.11.9/18-79 |
| 0x3_6000 | MDEU | Mode Register | R/W | Word | 18.8.4.2/18-98 |
| 0x3_6008 | | Key Size Register | R/W | Word | 18.8.4.4/18-102 |
| 0x3_6010 | | Data Size Register | R/W | Word | 18.8.4.5/18-102 |
| 0x3_6018 | | Reset Control Register | R/W | Word | 18.8.4.6/18-102 |
| 0x3_6028 | | Status Register | R | — | 18.8.4.7/18-103 |
| 0x3_6030 | | Interrupt Status Register | R/W | Word | 18.8.4.8/18-104 |
| 0x3_6038 | | Interrupt Mask Register | R/W | Word | 18.8.4.9/18-106 |
| 0x3_6040 | | ICV Size Register | W | Word | 18.8.4.10/18-107 |
| 0x3_6050 | | End_of_message | W | Word | 18.8.4.11/18-107 |
| 0x3_6100–0x3_6140 | | Context Registers | R/W | Byte | 18.8.4.12/18-108 |
| 0x3_6400–0x3_647F | | Key Registers | W | Byte | 18.8.4.13/18-111 |
| 0x3_6800–0x3_6FFF | | Input FIFO | W ¹ | Byte | 18.8.4.14/18-111 |

Table 18-3. SEC Address Map (continued)

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|--------|---------------------------|----------------|----------|------------------|
| Byte Address Offset (AD 17–0) | Module | Register | Access | Write By | Section/Page |
| 0x3_A000 | RNGU | Mode Register | R/W | Word | 18.8.6.1/18-120 |
| 0x3_A010 | | Data Size Register | R/W | Word | 18.8.6.2/18-120 |
| 0x3_A018 | | Reset Control Register | R/W | Word | 18.8.6.3/18-121 |
| 0x3_A028 | | Status Register | R | — | 18.8.6.4/18-121 |
| 0x3_A030 | | Interrupt Status Register | R/W | Word | 18.8.6.5/18-122 |
| 0x3_A038 | | Interrupt Mask Register | R/W | Word | 18.8.6.6/18-123 |
| 0x3_A050 | | End_of_message | W | Word | 18.8.6.7/18-124 |
| 0x3_A400–0x3_A43F | | Entropy Registers | W | Word | 18.8.6.8/18-125 |
| 0x3_A800–0x3_AFFF | | Output FIFO | R ¹ | — | 18.8.6.8/18-125 |
| 0x3_C000 | PKEU | Mode Register | R/W | Word | 18.8.5.1/18-111 |
| 0x3_C008 | | Key Size Register | R/W | Word | 18.8.5.2/18-112 |
| 0x3_C010 | | Data Size Register | R/W | Word | 18.8.5.4/18-114 |
| 0x3_C018 | | Reset Control Register | R/W | Word | 18.8.5.5/18-114 |
| 0x3_C028 | | Status Register | R | — | 18.8.5.6/18-115 |
| 0x3_C030 | | Interrupt Status Register | R/W | Word | 18.8.5.7/18-116 |
| 0x3_C038 | | Interrupt Mask Register | R/W | Word | 18.8.5.8/18-117 |
| 0x3_C040 | | ABSize | R/W | Word | 18.8.5.3/18-113 |
| 0x3_C050 | | End_of_message | W | Word | 18.8.5.9/18-118 |
| 0x3_C200–0x3_C27F | | Parameter Memory A0 | R/W | Byte | 18.8.5.10/18-119 |
| 0x3_C280–0x3_C2FF | | Parameter Memory A1 | R/W | Byte | |
| 0x3_C300–0x3_C37F | | Parameter Memory A2 | R/W | Byte | |
| 0x3_C380–0x3_C3FF | | Parameter Memory A3 | R/W | Byte | |
| 0x3_C400–0x3_C47F | | Parameter Memory B0 | R/W | Byte | |
| 0x3_C480–0x3_C4FF | | Parameter Memory B1 | R/W | Byte | |
| 0x3_C500–0x3_C57F | | Parameter memory B2 | R/W | Byte | |
| 0x3_C580–0x3_C5FF | | Parameter Memory B3 | R/W | Byte | |
| 0x3_C800–0x3_C9FF | | Parameter Memory N | R/W | Byte | |
| 0x3_CA00–0x3_CBFF | | Parameter Memory E | W | Byte | |

Table 18-3. SEC Address Map (continued)

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|--------|---------------------------|----------------|----------|---------------------------------|
| Byte Address Offset (AD 17–0) | Module | Register | Access | Write By | Section/Page |
| 0x3_F000 | CRCU | Mode Register | R/W | Word | 18.8.2.2/18-80 |
| 0x3_F008 | | Key Size Register | R/W | Word | 18.8.2.3/18-81 |
| 0x3_F010 | | Data Size Register | R/W | Word | 18.8.2.4/18-81 |
| 0x3_F018 | | Reset Control Register | R/W | Word | 18.8.2.5/18-82 |
| 0x3_F020 | | Control | R/W | Word | 18.8.2.6/18-82 |
| 0x3_F028 | | Status Register | R | — | 18.8.2.7/18-83 |
| 0x3_F030 | | Interrupt Status Register | R/W | Word | 18.8.2.8/18-84 |
| 0x3_F038 | | Interrupt Mask Register | R/W | Word | 18.8.2.9/18-86 |
| 0x3_F048 | | Data Out Register | R | Byte | 18.8.2.10/18-87 |
| 0x3_F050 | | End Of Message Register | W | Word | 18.8.2.11/18-87 |
| 0x3_F100 | | Received ICV Register | R/W | Byte | 18.8.2.12/18-87 |
| 0x3_F108 | | Context Register | R/W | Byte | 18.8.2.13/18-87 |
| 0x3_F400 | | Key Register | R/W | Byte | 18.8.2.14/18-88 |
| 0x3_F800–0x3_FFFF | | Input FIFO | W ¹ | Byte | 18.8.2.15/18-89 |

¹ For the EU FIFOs, write operations anywhere in the address range enqueue to the input FIFO, and read operations anywhere in the address range dequeue from the output FIFO. See the referenced section for more detailed information.

18.3 Descriptors

The host processor maintains a record of current secure sessions and the corresponding keys and contexts of those sessions. Once the host has determined that a security operation is required, it creates a “descriptor” containing all the information the SEC needs to perform the security operation. The host creates the descriptor in main memory, then writes a pointer to the descriptor into the Fetch FIFO of one of the SEC channels. The channel uses this pointer to read the descriptor into its descriptor buffer. Once it obtains the descriptor, the SEC uses its bus mastering capability to obtain inputs and write results, thus off-loading data movement and encryption operations from the host processor.

Descriptors are only used in channel-controlled accesses to SEC, and not in host-controlled accesses. For more information about host-controlled access, see #####’.

18.3.1 Descriptor Structure

SEC descriptors are conceptually similar to descriptors used by most devices with DMA capability. SEC descriptors are designed to support the cryptographic computation of a single packet using a single descriptor. SEC descriptors have a fixed length of 64 bytes, that is, eight 64-bit words (referred to as dwords). A descriptor consists of one “header dword” and seven “pointer dwords”, as seen in [Figure 18-2](#).

| | 0 | 15 | 16 | 17 | 23 | 24 | 27 | 28 | 31 | 32 | 63 |
|---------------------|---------------|----|---------|----|-------|----------|----|----|----|----|-----------------|
| Header Dword | Header | | | | | | | | | | Reserved |
| Pointer Dword 0 | Length0 | J0 | Extent0 | — | Eptr0 | Pointer0 | | | | | |
| Pointer Dword 1 | Length1 | J1 | Extent1 | — | Eptr1 | Pointer1 | | | | | |
| Pointer Dword 2 | Length2 | J2 | Extent2 | — | Eptr2 | Pointer2 | | | | | |
| Pointer Dword 3 | Length3 | J3 | Extent3 | — | Eptr3 | Pointer3 | | | | | |
| Pointer Dword 4 | Length4 | J4 | Extent4 | — | Eptr4 | Pointer4 | | | | | |
| Pointer Dword 5 | Length5 | J5 | Extent5 | — | Eptr5 | Pointer5 | | | | | |
| Pointer Dword 6 | Length6 | J6 | Extent6 | — | Eptr6 | Pointer6 | | | | | |

Figure 18-2. Descriptor Format

As shown in [Figure 18-2](#), the first and second halves of the header dword are denoted as descriptor control and descriptor feedback fields, respectively. The descriptor control field of the header dword specifies the security operation to be performed, the execution unit(s) needed, and the modes for each execution unit. The descriptor feedback field is written to by the security engine upon completion of descriptor processing, when the “channel done writeback” feature is enabled. Further details about the header dword may be found in [Section 18.3.2, “Descriptor Format: Header Dword”](#).

The pointer dwords, all of which have the same format, contain pointer and length information for locating input or output parcels (such as keys, context, or text data). The large number of pointers provided in the descriptor allows for multi-algorithm operations that require fetching of multiple keys, as well as fetch and return of contexts. Any pointer dword that is not needed may be assigned a length of zero. Further details about the pointer dwords may be found in [Section 18.3.3, “Descriptor Format: Pointer Dwords”](#).

SEC descriptors include scatter/gather capability, which means that each pointer in a descriptor can be either a direct pointer to a contiguous parcel of data, or a pointer to a “link table” which is a list of pointers and lengths used to assemble the parcel. When a link table is used to read input data, this is referred to as a “gather” operation; when used to write output data, it is referred to as a “scatter” operation. Further details about scatter/gather capability may be found in [Section 18.3.4, “Link Table Format”](#).

18.3.2 Descriptor Format: Header Dword

Descriptors are created by the host to guide the SEC through required cryptographic operations. The header dword defines the operations to be performed, the mode for each operation, and internal addressing used by the controller and channel for internal data movement. The fields that must be supplied to SEC are shown in the “Field” rows of [Figure 18-3](#), and described in [Table 18-4](#). The SEC device drivers allow the host to create proper headers for each cryptographic operation.

SEC processing of a descriptor sometimes includes performing a header writeback, that is, writing the original header dword back to system memory with certain fields modified. The modified fields are shown in the “Writeback” rows of [Figure 18-3](#), and described in [Table 18-5](#).




Figure 18-3. Header Dword

Header dword bit definitions are described in [Table 18-4](#).

Table 18-4. Header Dword Bit Definitions

| Bits | Name | Description |
|----------------|-----------|--|
| 0–3 4–11 | OP_0: | When done writeback is used, then at the completion of descriptor processing this byte is written with the value 0xFF. To determine when done writeback is used, see the CDWE, NT, and CDIE fields in the channel configuration register (see Table 18-11). |
| | EU_SEL0 | Reserved |
| | MODE0 | Integrity check comparison result from primary: These bits are supplied by the primary EU when descriptor processing is complete. 00 No integrity check comparison was performed 01 The integrity check comparison passed 10 The integrity check comparison failed 11 Reserved |
| 12–15 16–23 | OP_1: | Reserved |
| | EU_SEL1 | Integrity check comparison result from secondary: These bits are supplied by the secondary EU (if any) when descriptor processing is complete. 00 No integrity check comparison was performed 01 The integrity check comparison passed 10 The integrity check comparison failed 11 Reserved |
| | MODE1 | Reserved |
| 24–28 | DESC_TYPE | Descriptor Type: This, along with the DIR field, determines the sequence of actions to be performed by the channel and selected EUs using the blocks of data listed in the rest of the descriptor. The attributes determined include the direction of data flow for each data block, which EU (primary or secondary) is accessed, what snooping options are used, and which internal EU addresses are accessed. See Section 18.3.2.2, “Selecting Descriptor Type—DESC_TYPE,” for possible values. |
| 29 | — | Reserved. |

Table 18-4. Header Dword Bit Definitions (continued)

| Bits | Name | Description |
|------|------|--|
| 30 | DIR | Direction: direction of overall data flow: 0 Outbound 1 Inbound This, along with the DESC_TYPE field, helps determine the sequence of actions to be performed by the channel and selected EUs. |
| 31 | DN | Done notification: 0 No done notification. 1 Signal “done” to the host on completion of this descriptor. This enables done notification if the NT field is 1 in the channel configuration register (see Table 18-11). The done notification can take the form of an interrupt, a writeback in the DONE field of this header dword (see Table 18-5), or both, depending upon the states of the CDIE (channel done interrupt enable) and CDWE (channel done writeback enable) bits in the channel configuration register. |

A detailed description of the header dword fields used in writeback notification is provided in [Table 18-5](#).

Table 18-5. Header Dword Writeback Bit Definitions

| Bits | Name | Description |
|-------|-------|---|
| 0–7 | DONE | When done writeback is used, then at the completion of descriptor processing this byte is written with the value 0xFF. To determine when done writeback is used, see the CDWE, NT, and CDIE fields in the channel configuration register (see Table 18-11). |
| 8–34 | — | Reserved |
| 35–36 | ICCR0 | Integrity check comparison result from primary. These bits are supplied by the primary EU when descriptor processing is complete. 00 No integrity check comparison was performed 01 The integrity check comparison passed 10 The integrity check comparison failed 11 Reserved |
| 37–42 | — | Reserved |
| 43–44 | ICCR1 | Integrity check comparison result from secondary. These bits are supplied by the secondary EU (if any) when descriptor processing is complete. 00 No integrity check comparison was performed 01 The integrity check comparison passed 10 The integrity check comparison failed 11 Reserved |
| 45–63 | — | Reserved |

18.3.2.1 Selecting Execution Units—EU_SEL0 and EU_SEL1

[Table 18-6](#) shows the values for EU_SEL0 and EU_SEL1 in the descriptor header. The following rules govern the choices for these fields:

1. EU_SEL0 values of “No EU selected” or “Reserved” will result in an “Unrecognized Header Error” condition during processing of the descriptor header.
2. The only valid choices for EU_SEL1 are “No EU selected”, CRCU, or MDEU. Any other choice will result in an “Unrecognized Header” error condition.

- If EU_SEL1 is CRCU or MDEU, then EU_SEL0 must be DEU, or AESU. All other values of EU_SEL0 will result in an “Unrecognized header” error condition.

Table 18-6. EU_SEL0 and EU_SEL1 Values

| Value (Binary) | Selected EU |
|----------------|-------------------------------|
| 0000 | No EU selected |
| 0001 | Reserved |
| 0010 | DEU |
| 0011 | MDEU-A |
| 1011 | MDEU-B |
| 0100 | RNGU |
| 0101 | PKEU |
| 0110 | AESU |
| 0111 | Reserved |
| 1000 | CRCU |
| Others | Reserved |
| 1111 | Reserved for header writeback |

MDEU uses two different designators to refer to the same Message Digest Execution Unit. If MDEU-B is selected, then the Channel configures MDEU to perform SHA-224, SHA-256, SHA-384, and SHA-512. If MDEU-A is selected, then the Channel configures MDHA to perform SHA-160, SHA-224, SHA-256, or MD5. This configuration is achieved automatically; the channel will set bit 51 of the MDEU mode register as it inserts the MODE0 (or MODE1) value into the MDEU mode register. For more information, see [Section 18.8.4.2, “MDEU Mode Register.”](#)

18.3.2.2 Selecting Descriptor Type—DESC_TYPE

[Table 18-7](#) shows the permissible values for the DESC_TYPE field in the descriptor header. Descriptor types from the SEC1.0, which have “0” in the last bit, are listed first, followed by new SEC 2.x types, which have “1” in the last bit.

Table 18-7. Descriptor Types

| Value (Binary) | Descriptor Type | Notes |
|----------------|--------------------|--|
| 0000_0 | aesu_ctr_nonsnoop | AESU CTR nonsnooping ¹ |
| 0001_0 | common_nonsnoop | Common, non-snooping, non-PKEU, non-AFEU |
| 0010_0 | hmac_snoop_no_afeu | Snooping, HMAC, non-AFEU |
| 0011_0 | — | Reserved |
| 0100_0 | — | Reserved |

Table 18-7. Descriptor Types

| Value (Binary) | Descriptor Type | Notes |
|----------------|---------------------|---|
| 0101_0 | — | Reserved |
| 0110_0 | — | Reserved |
| 0111_0 | — | Reserved |
| 1000_0 | pkeu_mm | PKEU-Montgomery Multiplication |
| 1001_0 | — | Reserved |
| 1010_0 | — | Reserved |
| 1011_0 | — | Reserved |
| 1100_0 | hmac_snoop_aesu_ctr | AESU CTR hmac snooping ² |
| 1101_0 | — | Reserved |
| 1110_0 | — | Reserved |
| 1111_0 | — | Reserved |
| 0000_1 | ipsec_esp | IPsec ESP mode encryption and hashing |
| 0001_1 | 802.11i_aes_ccmp | CCMP encryption and hashing, suitable for 802.11i |
| 0010_1 | srtp | SRTP encryption and hashing |
| 0011_1 | pkeu_build | pkeu_build Elliptic Curve Cryptography |
| 0100_1 | pkeu_ptmul | pkeu_ptmul Elliptic Curve Cryptography |
| 0101_1 | pkeu_ptadd_dbl | pkeu_ptadd_dbl Elliptic Curve Cryptography |
| 0110_1 | — | Reserved |
| 0111_1 | — | Reserved |
| 1000_1 | tls_ssl_block | TLS/SSL generic block cipher |
| 1001_1 | — | Reserved |
| 1010_1 | raid_xor | XOR 2-6 sources together |
| 1011_1 | ipsec_aes_gcm | IPsec ESP mode using AES GCM encryption and hashing |
| 1100_1 | dbl_digest | Do two digest operations (CRC or hash) |
| Others | — | Reserved |

¹ Type 0000_0 is for AES-CTR operations. Type 0001_0 also supports AES-CTR, however to use AES-CTR with 0001_0, the user must pre-pend zeros to the AES-Context before loading the AES Context Registers.

² Type 1100_0 is for AES-CTR operations with HMAC. Type 0010_0 also supports AES-CTR with HMAC, however to use AES-CTR with 0010_0, the user must pre-pend zeros to the AES-Context before loading the AES Context Registers.

For more about descriptor types and the data used for each type, see [Section 18.4, “Descriptor Types.”](#)

18.3.3 Descriptor Format: Pointer Dwords

The descriptor contains seven “pointer dwords” which define where in memory the SEC should access its input and output parcels. The pointer dwords are numbered 0 to 6 as shown in [Figure 18-2](#). The channel determines how it uses each of the pointer dwords based on the descriptor type and direction fields in the header (see [Table 18-4](#)).

The pointer dword bit fields as shown in [Figure 18-4](#), and are described in [Table 18-8](#).




Figure 18-4. Pointer Dword

Table 18-8. Pointer Dword Field Definitions

| Bits | Name | Description |
|-------|---------|---|
| 0–15 | LENGTH | Length. A number of bytes in the range 0 to 65535. The use of this field depends on the “Descriptor Type” and “Direction” in the header dword. A value of zero may cause the channel to skip this dword. |
| 16 | J | Jump. Determines whether to “jump” to a link table whenever the POINTER field in this same word is used. 0 The POINTER field points to data. 1 The POINTER field points to a link table, and scatter/gather is enabled. |
| 17–23 | EXTENT | Extent. A number of bytes in the range 0 to 127. The use of this field depends on the “Descriptor Type” and “Direction” in the header dword. |
| 24–27 | — | Reserved |
| 28–31 | EPTR | Extended Pointer. Concatenated as the top 4 bits of the pointer when EAE is high (see the EAE bit in Table 18-11). |
| 32–63 | POINTER | Pointer. A memory address. |

On occasion, a descriptor field may not be applicable to the requested service. With seven pointer dwords, it is possible that not all these dwords are required to specify the input and output parameters (for instance, some operations do not require context.) Wherever a particular field is not used, it should be set to zero.

The channel proceeds linearly through the descriptor, fetching LENGTH data beginning at location POINTER. If the EAE (extend address enable) bit is set in the channel configuration register (see [Table 18-11](#)), then the four EPTR bits are concatenated with the POINTER field to form a 36-bit pointer address.

If all the data of LENGTH is found contiguously beginning at POINTER, then the Jump bit is not set. Otherwise, POINTER indicates the location of a link table (scatter-gather list). For more details, see [Section 18.3.4, “Link Table Format”](#).

LENGTH and EXTENT fields normally specify the sizes of parcels: often (but not always) the size of the parcel located at the address contained in the matching POINTER field¹. However, in some cases the POINTER field is zero, and the LENGTH and/or EXTENT fields simply specify values to be written to an EU. The specific use of these fields in each channel depends on the descriptor type and direction fields in the descriptor’s header dword (see [Table 18-4](#)).

The Raid-XOR descriptor type does not support scatter/gather capability. However, scatter/gather is available for all pointer dwords for all other descriptor types (provided the J bit is set).

18.3.4 Link Table Format

Link tables implement scatter/gather capability. For “gather” operations, a link table specifies a list of “memory segments” that are to be concatenated in the process of assembling parcels. For “scatter” operations, a link table specifies a list of memory segments into which the output data should be written. Scatter or gather of a parcel may be specified by a single link table or by a chain of link tables that are linked together with pointers, as shown in [Figure 18-6](#).

A link table may contain any number of dword entries. Link table entry format is shown in [Figure 18-5](#), and the field definitions are given in [Table 18-9](#).




Figure 18-5. Link Table Entry

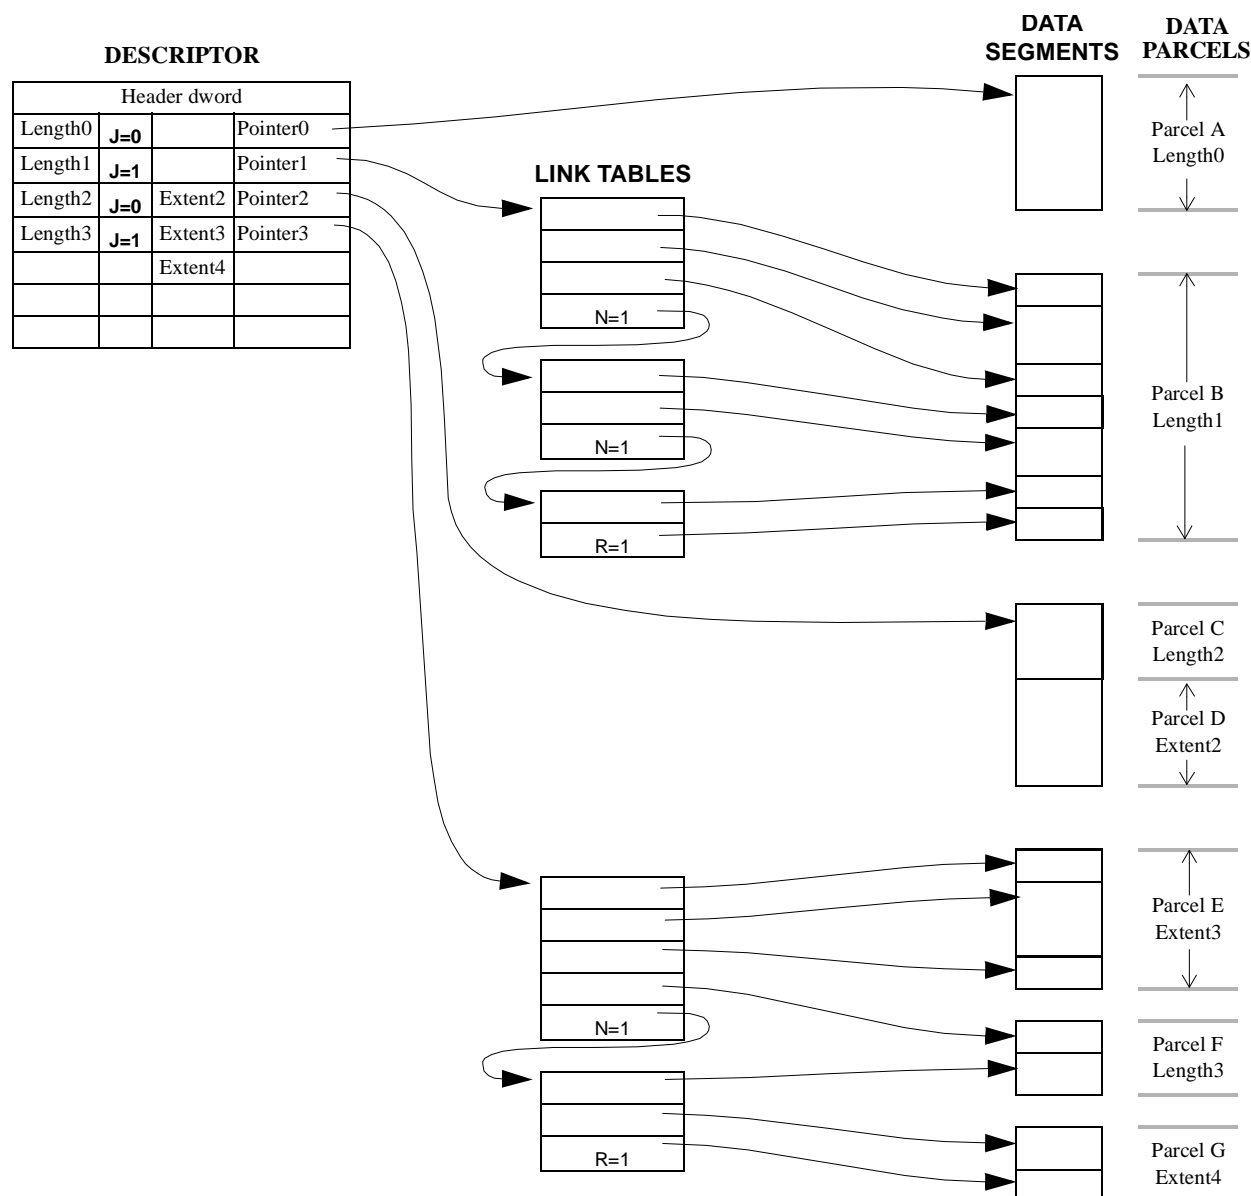
1. Sometimes an EXTENT field refers to data in a pointer which is not in the same dword. For example, with the CCMP descriptor type (see [Table 18-10](#)), the length of the CRC check field appears in Extent0, but the field that is Extent0 bytes in length is referred to either by Pointer4 or Pointer5, depending on the direction bit in the descriptor’s header dword.

Table 18-9. Link Table Field Definitions

| Bits | Name | Description |
|-------|--------|--|
| 0–15 | SEGLEN | Length. 0-15: When N = 0, a number in the range 1 to 65535, specifying the number of bytes in the memory segment pointed to by SEGPTR. A value of 0 will cause the SGZL error bit to be set in the Channel Status (see Section 18.5.4.2, “Channel Status Register (CSR)”). When N = 1, must be 0. |
| 16–21 | — | Reserved |
| 22 | R | Return. When N = 0: 0 No special action. 1 This is the last entry in the chain of link tables. If this entry does not specify the right number of bytes to complete the last parcel, a G-STATE or S-STATE error will be set in the Channel Status Register (see Section 18.5.4.2, “Channel Status Register (CSR)”). When N = 1, ignored. |
| 23 | N | Next. 0 No special action. 1 This is the last dword in the current link table. The SEGPTR field is the address of the next link table in the chain. |
| 24–27 | — | Reserved |
| 28–31 | EPTR | Extended pointer. Concatenated as the top 4 bits of the segment pointer when EAE is high (see the EAE bit in Table 18-11). |
| 32–63 | SEGPTR | Segment pointer. A memory address. |

There are two kinds of link table entries: “regular” entries or “next” entries (which have the N bit cleared or set, respectively). Each “regular” entry specifies a memory segment by means of a 36-bit starting address (SEGPTR) and a 16-bit length (SEGLEN). A “next” entry is used at the end of a link table to specify that the list of memory segments is continued in another link table. In a “next” entry, the N bit is set, the SEGPTR field gives the address of the next link table, and the SEGLEN field must be cleared. A chain of link tables may contain any number of link tables. Whether the list of memory segments is in a single link table or split into several link tables, the last entry in the last link table is a “regular” entry with the R (return) bit set. The R bit signifies the end of link table operations so that the channel returns to the descriptor for its next pointer (if any).

The construction and use of link tables is illustrated in [Figure 18-6](#).



This figure illustrates various ways that a descriptor may specify parcels:

The first pointer dword in the descriptor specifies Parcel A using the simplest method—the parcel is specified directly through Pointer0 and Length0.

The next pointer dword uses a chain of link tables to specify Parcel B. Since J=1, Pointer1 is used as the address of a link table. The link table specifies several “regular” entries specifying data segments to be concatenated. The last word of the link table is a “next” entry indicating that the list continues in the next link table. The last entry in the last link table of the chain has the R bit set.

The last cases illustrate how one pointer in a descriptor can be used to specify multiple parcels. Pointer2 and Length2 specify Parcel C, then Parcel D follows immediately afterwards, with length specified by Extent2. Pointer3 is used for three parcels (E, F, and G), this time using link tables.

Figure 18-6. Descriptors, Link Tables, and Parcels

As shown in [Figure 18-6](#), in some cases, a single parcel is accessed through a given POINTER, and the chain of link tables specifies only that parcel (this is the most common situation). In other cases, the

descriptor POINTER is used multiple times to access a sequence of parcels, and the chain of link tables must supply data for the entire sequence.

18.3.4.1 Example of Link Table Operation

To further clarify the link table's operation, it is explained in detail the case where the fourth pointer dword in the descriptor in [Figure 18-6](#) is used to access parcels. In this example, the descriptor type is such that Pointer3 is used to access successive parcels of size Extent3, Length3, and Extent4 respectively (refer to [Table 18-10](#) for the significance of POINTER, EXTENT, and LENGTH fields in various descriptor types).

Since the J3 bit is set, Pointer3 is used as the address of a link table and not a data address. The channel begins by reading the first four dwords of the link table starting at Pointer3 into an internal "gather table buffer".

Using the first entry of the gather table buffer, the channel starts accessing the parcel by reading SEGLEN bytes beginning at SEGPTR. If the required parcel size (specified by 'Extent3' in the pointer dword) is greater than this first segment length, the channel moves on to the next entry of the gather table buffer, and reads SEGLEN bytes starting at SEGPTR. This process continues as long as there are more bytes to be read in the parcel. If all the link table entries in the channel's gather table buffer have been exhausted, then the channel reads the next four dwords of the link table into its gather table buffer. If a gather table buffer entry is encountered in which the N bit is set, the channel uses the SEGPTR field in that word to find the next link table in the chain.

Now assume that the channel accesses its next parcel using Pointer3 again, this time with length given by Length3. In this case the channel continues to the next line of the link table, and begins reading the memory segment specified there. As before, the channel concatenates memory segments from as many link table entries as necessary to obtain the required number of bytes (Length3).

Similarly, the next parcel is obtained by using Pointer3 yet again, this time with length given by Extent4.

Assume that for the current descriptor type, the Extent4 parcel is the last one to be accessed through Pointer3. Then the link table entry that supplies the last memory segment for Extent4 has the R bit set, signifying that this is the last entry in the chain of link tables.

NOTE

The link table or chain of link tables accessed through a descriptor pointer must specify enough memory segments to hold precisely *all the data that will be accessed through that pointer*. This means that the combined lengths of the parcels associated with that pointer (where each parcel length is specified by a particular LENGTH or EXTENT field in the descriptor) must equal the combined lengths of the link table memory segments (SEGLEN fields). Otherwise the channel sets the error state in the SGLM bit of the channel status register (see [Section 18.5.4.2, "Channel Status Register \(CSR\)"](#)).

18.4 Descriptor Types

Table 18-10 shows in summary form how the pointer dwords are used with the various descriptor types. Detailed information about each descriptor type is given in the remainder of this chapter. Additional explanation of the use of certain descriptor types, can be found in the SEC 3.0 Descriptor Programmer's Guide.

In Table 18-10, older descriptor types which end in 0 are listed first, followed by newer types which end in 1.

Table 18-10. Descriptor Format Summary

| Descriptor Type | Field Type | Pointer Dword0 | Pointer Dword1 | Pointer Dword2 | Pointer Dword3 | Pointer Dword4 | Pointer Dword5 | Pointer Dword6 |
|--|------------|----------------|-------------------|----------------|-------------------|----------------|-----------------------------|----------------|
| 0000_0 aesu_ctr_ nosnoop | Length | reserved | Cipher Context In | Cipher Key | Main Data In | Data Out | Cipher Context Out | reserved |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 0001_0 common_ nosnoop for DES, RNGU, AES-CCM | Length | reserved | Context In | Key | Main Data In | Data Out | Context Out (incl. ICV out) | reserved |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 0001_0 common_ nosnoop for MDEU | Length | reserved | Context In | Key | Main Data In | ICV In | Context Out | reserved |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 0001_0 common_ nosnoop for AES-XCBC, AES-CMAC, CRCU | Length | reserved | Context In | Key | Main Data In | reserved | Context Out | ICV Out |
| | Extent | reserved | reserved | reserved | reserved | ICV In | reserved | reserved |
| 0010_0 hmac_snoop_ _no_afeu | Length | Hash Key | Hash-only Header | Cipher Key | Cipher Context In | Main Data In | Data Out | ICV Out |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 1000_0 pkeu_mm | Length | "N" In | "B" In | "A" In | "E" In | "B" Out | reserved | reserved |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 1100_0 hmac_snoop_ aesu_ctr | Length | Hash Key | Hash-only Header | AES Key | AES Context In | Main Data In | Data Out | ICV Out |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 0000_1 ipsec_esp | Length | HMAC Key | Hash-only Header | Cipher IV In | Cipher Key | Main Data In | Data Out | Cipher IV Out |
| | Extent | reserved | reserved | reserved | reserved | ICV In | ICV Out | reserved |

Table 18-10. Descriptor Format Summary (continued)

| Descriptor Type | Field Type | Pointer Dword0 | Pointer Dword1 | Pointer Dword2 | Pointer Dword3 | Pointer Dword4 | Pointer Dword5 | Pointer Dword6 |
|---|------------|---------------------|---------------------|---------------------|---------------------|------------------------|----------------------|--------------------|
| 0001_1 802.11i AES ccmp | Length | CRC-only Header | AES Context In | AES Key | Hash-only Header | Main Data In | Data Out | AES Context Out |
| | Extent | | reserved | reserved | reserved | MIC In | MIC Out | reserved |
| 0010_1 srtp inbound | Length | HMAC Key | AES Context In | AES Key | Main Data In | Data Out | HMAC Out | AES Context Out |
| | Extent | | reserved | reserved | reserved | Hash-only Header | Hash-only Trailer | reserved |
| 0010_1 srtp outbound | Length | HMAC Key | AES Context In | AES Key | Main Data In | HMAC In | Data Out | AES Context Out |
| | Extent | | reserved | reserved | reserved | Hash-only Header | Hash-only Trailer | HMAC Out |
| 0011_1 pkeu_build | Length | "A0" In | "A1" In | "A2" In | "A3" In | "B0" In | "B1" In | "Build" Out |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 0100_1 pkeu_ptmul | Length | "N" In | "E" In | "Build" In | "B1" Out | "B2" Out | "B3" Out | reserved |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 0101_1 pkeu_ptadd_ dbl | Length | "N" In | "Build" In | "B2" In | "B3" In | "B1" Out | "B2" Out | "B3" Out |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |
| 1000_1 outbound tls_ssl_ block | Length | MAC Key | Cipher IV In | Cipher Key | Main Data In | Cipher-only Trailer | Data Out | Cipher IV Out |
| | Extent | reserved | reserved | reserved | Hash-only Header | ICV Out | reserved | reserved |
| 1000_1 inbound tls_ssl_ block | Length | MAC Key | Cipher IV In | Cipher Key | reserved | Main Data In | Data Out | Cipher IV Out |
| | Extent | reserved | reserved | reserved | Hash-only Header | ICV In | ICV Out | reserved |
| 1010_1 raid_xor | Length | Source F Data In | Source E Data In | Source D Data In | Source C Data In | Source B Data In | Source A Data In | Data Out |
| | Extent | reserved | reserved | reserved | reserved | reserved | reserved | reserved |

Table 18-10. Descriptor Format Summary (continued)

| Descriptor Type | Field Type | Pointer Dword0 | Pointer Dword1 | Pointer Dword2 | Pointer Dword3 | Pointer Dword4 | Pointer Dword5 | Pointer Dword6 |
|-------------------------|------------|----------------|----------------|-----------------|-----------------|----------------|----------------|----------------|
| 1011_1 ipsec_aes_gcm | Length | AES Ctx In | AAD In | Nonce Part 2 In | AES Key | Main Data In | Data Out | Cipher Ctx Out |
| | Extent | reserved | reserved | reserved | Nonce Part 1 In | AES ICV In | ICV Out | CRC ICV In/Out |
| 1100_1 dbl_digest | Length | Header In | Payload In | reserved | reserved | reserved | reserved | reserved |
| | Extent | Header ICV | Payload ICV | Header ICV Out | Payload ICV Out | reserved | reserved | reserved |
| others | | Reserved | | | | | | |

18.5 Polychannel

The polychannel is the main control unit in the SEC. It implements four independent channels.

Each cryptographic task performed by the SEC is managed by a channel and makes use of one or more of the SEC's execution units (EUs). Control information and data pointers for a given task are stored in the form of a descriptor (see [Section 18.3.1, "Descriptor Structure"](#)) placed in system memory. A descriptor determines what EUs are used, how they are configured, where to fetch needed data, and where to store the results.

The following subsections describe the operation (including descriptor processing, arbitration, and host notification), registers, and interrupts of the polychannel.

18.5.1 Channel Operation

18.5.1.1 Channel Descriptor Processing

To invoke a cryptographic task, the host constructs a descriptor, selects a channel, and writes a pointer to the descriptor into the selected channel's fetch FIFO. Each fetch FIFO can store up to 24 pointers.

Typical operations performed by a channel to process a descriptor are:

1. Analyze the descriptor header to determine the cryptographic services required, and arbitrate for the appropriate EUs. If required EUs are already reserved by another channel, wait for the EUs to be available. When available, reserve them.
2. Set the mode register in each reserved EU(s) for the required EU function.
3. Fetch "parcels" (up to 64K-1 bytes long) from system memory using pointers from the descriptor buffer, and place them in either an EU input FIFO or EU registers, as appropriate. The term "parcel" refers here to any input or output of an EU algorithm, such as a key, hash result, input context, output context, or text data. "Context" refers to either an IV (initialization vector) or other

internal EU state that can be read out or loaded in. “Text data” refers to plaintext or ciphertext to be operated on. Each parcel transfer may involve using link tables to gather input data that has been split into multiple segments in system memory.

4. Take data accumulated in the EU output FIFO and write it to system memory using pointers from the descriptor buffer. This may again involve using link tables to scatter output data into multiple segments in system memory.
5. If the data size is greater than EU FIFO size, continue fetching input data and writing output data to memory as needed.
6. After writing the last input data to each EU’s input FIFO, write to the end of message register in the EU.
7. Wait for EU(s) to complete processing of text data.
8. Unload final results from output FIFOs and context registers and write them to external memory using pointers from the descriptor buffer. This may again involve using link tables to scatter output data into multiple segments in system memory.
9. Reset and release the EUs.
10. If enabled, then notify the host of descriptor completion (see [Section 18.5.1.3, “Channel Host Notification”](#)).

18.5.1.2 Channel Arbitration

All channels share a set of common resources, including the EUs and the SEC’s bus master interface (managed by the SEC controller). When multiple channels are used in parallel, arbitration may be required to determine which channel is serviced. The different arbitration schemes are described in [Section 18.6.2, “Arbitration Algorithms.”](#)

Generally speaking, no arbitration for use of the controller/bus master interface is required. The channels within the polychannel execute one at a time, so individual channels do not experience contention when requests to the controller. In effect, when a channel wins arbitration for use of the polychannel, it wins use of the controller as well.

The same is not true of EUs. Once the controller has assigned an EU to a channel, that channel owns the EU for the duration of descriptor processing. The maximum amount of data that can be processed by a single descriptor is 64 Kbytes, which prevents a channel from owning an EU for an unbounded length of time.

While one channel owns a particular EU, it is possible for two or more other channels to request access to the same EU; in this case, an arbitration scheme determines which channel is granted next access. EU arbitration schemes are similar to channel arbitration mentioned above, and are described in [Section 18.6.2, “Arbitration Algorithms”](#).

If a channel needs two EUs, a primary and a secondary, it requests them one at a time. Sometimes a channel reserves one EU and then has to wait for some other channel(s) to finish before obtaining the second requested EU. Though such waiting may occur, the requests are always eventually satisfied. Deadlock is avoided through the following design rules:

The channel always requests the secondary EU first.

In cases where both a primary and secondary are used, the choices for primaries and secondaries are distinct sets. Primaries are AESU and DES, and the secondaries are MDEU and CRCU.

18.5.1.3 Channel Host Notification

When a channel completes operation on a descriptor, it can notify the host that it is done through interrupt and/or through a writeback of the descriptor header dword. In case the descriptor operation is not completed or completed with a known error, the host may be notified by an error interrupt. The error interrupts, done interrupts and header writeback are described as follows:

- Error interrupts are always enabled at the channel level, but can be masked at the controller level. For more details concerning these interrupts, see [Section 18.5.2.2, “Channel Error Interrupt”](#).
- Done interrupts are enabled on a per-channel basis by programming the channel’s configuration register. For programming details, refer to [Section 18.5.2.1, “Channel Done Interrupt”](#) and [Section 18.5.4.1, “Channel Configuration Register \(CCR\)”](#).
- Independently of the done interrupt, channels can inform software of their completion status via header writebacks. Like done interrupts, writebacks are enabled on a per-channel basis by programming the CCR. If enabled, then upon completion the channel writes 0xFF to the DONE byte in the original descriptor header (see [Table 18-5](#)), allowing software to poll for completion of a specific descriptor. The CCR can also be programmed so that the channel writes back a status code indicating whether an integrity checking EU has encountered a mismatch between the received ICV and the recalculated ICV. [Table 18-5](#) shows the specific bytes in the descriptor header that are updated in this case.

For more details on programming the CCR for writeback, see [Section 18.5.4.1, “Channel Configuration Register \(CCR\)”](#)

NOTE

The done and status writebacks are not performed should the channel signal any error during processing. For example, there are no writebacks in case of a failing, unmasked ICV check in an EU.

18.5.2 Channel Interrupts

Active channels can assert done and error interrupts to the controller. As with all SEC interrupt events, channel done and error interrupts are reflected in the controller’s interrupt status register. Channel do not have internal interrupt masks, but the controller can be programmed to disable channel interrupts through its interrupt enable register. For more details on interrupt types and disablement, see [Section 18.6.4.2, “Interrupt Enable Register \(IER\)”](#).

18.5.2.1 Channel Done Interrupt

Channel done interrupt generation depend on the setting of the CDIE (channel done interrupt enable) and NT (notification type) bits in the channel configuration register (see [Section 18.5.4.1, “Channel Configuration Register \(CCR\)”](#)). If both CDIE and NT are set, the channel generates an interrupt event after every successfully completed descriptor; If CDIE is set and NT is cleared, an interrupt is generated

after each successfully completed descriptor with the DN (done notification) bit set in the descriptor’s header word. If the EU(s) signal any error during processing, the channel done interrupt is not generated.

Even if multiple channel done interrupt events are generated by a channel before the first can be cleared by the host, the interrupt events are not lost. The controller keeps count of the backlog of channel done interrupts from each channel (see [Section 18.6.3, “Controller Interrupts”](#)).

18.5.2.2 Channel Error Interrupt

The channel error interrupt is generated when an error condition occurs during descriptor processing. The error could be a bus error for a transaction requested by the channel; or it could be in one of the EUs reserved by the channel, or in the channel itself. The channel error interrupt is asserted as soon as the error condition is detected. The type of error condition is reflected in the ERROR field of the channel status register (CSR).

For most error types, the error causes the corresponding channel to halt. Any EUs reserved by the halted channel continue to be reserved until the channel reset occurs. Other channels continue normal processing, though they may be held up if they need an EU that is reserved by a halted channel.

Handling of errors depends on the error type. Details of each error type are given in [Table 18-13](#). For some types, the host must clear the source of the error before restarting the channel. If the channel is halted, the host restarts it by setting the no-pop-reset, continue or reset bits of the CCR (see [Section 18.5.4.1, “Channel Configuration Register \(CCR\)”](#)).

18.5.3 Polychannel Registers

The polychannel has several aggregate performance counters, which are common to all channels; plus a set of channel-specific registers, descriptor buffers, and link tables which are duplicated for each channel. The following subsections describes the format and function of all of these objects in the SEC’s memory.

18.5.3.1 Traffic Counters

The SEC maintains several counters, which are described in the following subsections.

18.5.3.1.1 Fetch FIFO Enqueue Counter

The Fetch FIFO Enqueue Counter, shown in [Figure 18-7](#), indicates the total number of descriptor addresses that have been enqueued to the channel fetch FIFOs. If this counter reaches all ones, the next count causes it to roll over to all zeros and set the FF_ENQ_CNT bit of the interrupt enable register (see [Section 18.6.4.2, “Interrupt Enable Register \(IER\)_x”](#)).




Figure 18-7. Fetch FIFO Enqueue Counter

18.5.3.1.2 Descriptor Finished Counter

The descriptor finished counter, shown in [Figure 18-8](#), indicates the total number of descriptors that have successfully completed processing. It does not count descriptors that halt due to error. If this counter reaches all ones, the next count causes it to roll over to all zeros and set the DF_CNT bit of the interrupt enable register (see [Section 18.6.4.2, “Interrupt Enable Register \(IER\)”](#)).




Figure 18-8. Descriptor Finished Counter

18.5.3.1.3 Data Bytes In Counter

The data bytes in counter, shown in [Figure 18-9](#), indicates the total number of bytes written into a primary EU input FIFO. If other parcels such as context or ICV are placed in the input FIFO, they are not counted. For cases where a secondary EU is used, data going only to the secondary EU (such as a hash-only region) is counted. If this counter reaches all ones, the next count causes it to roll over to all zeros and set the DI_CNT bit of the interrupt enable register (see [Section 18.6.4.2, “Interrupt Enable Register \(IER\)”](#)).




Figure 18-9. Data Bytes In Counter

18.5.3.1.4 Data Bytes Out Counter

The data bytes out counter, shown in [Figure 18-10](#), indicates the total number of payload bytes read from an EU output FIFO. If other parcels such as context or ICV are read from the output FIFO, they are not counted. If this counter reaches all ones, the next count causes it to roll over to all zeros and set the DO_CNT bit of the interrupt enable register (see [Section 18.6.4.2, “Interrupt Enable Register \(IER\)”](#)).




Figure 18-10. Data Bytes Out Counter

18.5.4 Channel Registers

These registers are replicated for each of the 4 channels in the polychannel.

18.5.4.1 Channel Configuration Register (CCR)

CCR contains bits that allow the user to configure and reset the channel, as shown in [Figure 18-11](#). [Table 18-11](#) describes each field of the CCR.




Figure 18-11. Channel Configuration Register (CCR)

Table 18-11. CCR Bit Definitions

| Bits | Name | Description |
|-------|------|--|
| 0–28 | — | Reserved, set to zero. |
| 29 | NPR | No-pop-reset. 0 No special action. 1 Causes the same channel reset actions as bit CON, except that the Fetch FIFO is left unchanged, such that the channel will pickup by re-fetching the previous descriptor. This permits debug of a descriptor in-place without having to rewrite the descriptor pointer into the Fetch FIFO. |
| 30 | CON | Continue bit. 0 No special action. 1 Causes the same channel reset actions as bit R, except that the Fetch FIFO and the lower half of the CCR register (bits 32–63) are not cleared. After the reset sequence is complete, this bit automatically returns to 0 and the channel resumes normal operation, servicing the next descriptor pointer in the Fetch FIFO, if any. |
| 31 | R | Reset channel. 0 No special action. 1 Causes a software reset of the channel, clearing all its registers. Some actions depend on what the channel is doing when the bit is set: <ul style="list-style-type: none"> • If the R bit is set while the channel is requesting an EU assignment from the controller, the channel cancels its request. • If the R bit is set after the channel has been assigned one or more EUs, the channel requests a write from the controller to set the software reset bit of each reserved EU. The channel then releases the EU(s). After the reset sequence is complete, the channel returns to the idle state and the R bit automatically returns to 0 and resumes normal operation. |
| 32–50 | — | Reserved, set to zero. |

Table 18-11. CCR Bit Definitions (continued)

| Bits | Name | Description |
|-------|------|---|
| 51 | PBS | Permit byte summing. 0 Bytes written to EU input FIFOs and read from EU output FIFOs are not counted in the data bytes counters 1 Bytes written to EU input FIFOs and read from EU output FIFOs are counted in the data bytes counters |
| 52–54 | — | Reserved, set to zero. |
| 55 | BS | Burst size. The SEC accesses long text-parcels in main memory through bursts of programmable size. 0 Burst size is 64 bytes 1 Burst size is 128 bytes |
| 56 | IWSE | ICV writeback status enable. 0 No special action. 1 If the descriptor calls for ICV checking, then at the completion of descriptor processing, the channel writes back to the descriptor header all of the writeback information shown in Table 18-5 , that is, the DONE, ICCR0, and ICCR1 fields. ¹ |
| 57 | — | Reserved, set to zero. |
| 58 | EAE | Extend address enable. This bit determines whether the channel uses a 36-bit address bus or a 32-bit address bus. 0 Channel's address bus is 32 bits. 1 Channel's address bus is 36 bits. |
| 59 | CDWE | Channel done writeback enable. 0 Channel done writeback disabled. 1 Channel done writeback enabled. Upon successful completion of descriptor processing, if the NT bit is reset (global), or if the DN (done notification) bit is set in the header word of the descriptor, then the channel notifies the host by writing back the descriptor header with the DONE field shown in Table 18-5 . This enables the host to poll the memory location of the original descriptor header to determine if that descriptor has been completed. ¹ |
| 60 | AWSE | Always writeback status enable. 0 No special action. 1 At the completion of processing each descriptor, the channel writes back to the descriptor header all of the writeback information shown in Table 18-5 , that is, the DONE, ICCR0, and ICCR1 fields. In this case, IWSE has no effect. ¹ |
| 61 | NT | Notification type. This bit controls when the channel will generate channel done notification. Channel done notification can take the form of an interrupt or modified header writeback or both, depending on the state of the CDIE and CDWE control bits. 0 Global notification—the channel will generate channel done notification (if enabled) at the end of each descriptor. 1 Selected notification—the channel will generate channel done notification (if enabled) at the end of every descriptor with the DN bit set in the descriptor header. |
| 62 | CDIE | Channel done interrupt enable 0 Channel done interrupt disabled 1 Channel done interrupt enabled. Upon successful completion of descriptor processing, if the NT bit is set for Global, or if the DN (Done Notification) bit is set in the header word of the descriptor, then notify the host by asserting an interrupt. ¹ Refer to Section 18.5.2, "Channel Interrupts," for complete description of channel interrupt operation. |
| 63 | — | Reserved, set to zero. |

¹ WARNING: The done interrupt, done writeback, and status writeback will not occur if an EU produces an error interrupt to the channel. In particular, if the ICV check error interrupt is enabled in the EU (see the ICE bit in the EU's Interrupt Mask Register), and the ICV check finds a mismatch, then the channel will produce an error interrupt, but no channel done interrupt and no writebacks.

18.5.4.2 Channel Status Register (CSR)

CSR contains status fields and counters that provide status information on the channel processing of the current descriptor. This register is for debug use. For a channel error, the error field indicates the type of error.




Figure 18-12. Channel Status Register (CSR)

Table 18-12 describes the channel status register fields.

Table 18-12. CSR Bit Descriptions

| Bits | Name | Description |
|-------|------------|--|
| 0 | — | Reserved |
| 1–7 | GET_STATE | Get State Machine State. This field reflects the state of the Get State Machine when it last went to sleep, or the state captured when an error occurred. For debug purposes only. |
| 8 | — | Reserved |
| 9–15 | PUT_STATE | Put State Machine State. This field reflects the state of the Put State Machine when it last went to sleep, or the state captured when an error occurred. For debug purposes only. |
| 16–22 | — | Reserved |
| 23–31 | MAIN_STATE | Main State Machine State. This field reflects the state of the Main State Machine when it last went to sleep, or the state captured when an error occurred. For debug purposes only. |
| 32–34 | — | Reserved, set to zero |
| 35–39 | FF_LEVEL | Fetch FIFO Level. This 5 bit counter indicates how many pointers are currently stored in the Fetch FIFO. |
| 40–43 | — | Reserved, set to zero |

Table 18-12. CSR Bit Descriptions (continued)

| Bits | Name | Description |
|-------|-------|---|
| 44 | PRD | Primary EU reset done. The PRI_RST_DONE bit reflects the state of the reset done signal from the assigned primary EU. 0 The assigned primary EU reset done signal is inactive. 1 The assigned primary EU reset done signal is active indicating its reset sequence has completed and it is ready to accept data. |
| 45 | SRD | Secondary EU reset done. The SEC_RST_DONE bit reflects the state of the reset done signal from the assigned secondary EU. 0 The assigned secondary EU reset done signal is inactive. 1 The assigned secondary EU reset done signal is active indicating its reset sequence has completed and it is ready to accept data. |
| 46 | PD | Primary EU done. The PRI_DONE bit reflects the state of the done interrupt from the assigned primary EU. 0 The assigned primary EU done interrupt is inactive. 1 The assigned primary EU done interrupt is active indicating the EU has completed processing. This means that final values are available from EU registers. For EUs with output FIFOs, it means that all textdata output has been placed in the output FIFO. For EUs that provide context out through the output FIFO, the EU places the context in the output FIFO after asserting PRI_DONE. |
| 47 | SD | Secondary EU done. The SEC_DONE bit reflects the state of the done interrupt from the assigned secondary EU. 0 The assigned secondary EU done interrupt is inactive. 1 The assigned secondary EU done interrupt is active indicating the EU has completed processing. This means that final values are available from EU registers. |
| 48–59 | Error | Error bits for the channel. See below. |
| 60–63 | — | Reserved. |

Table 18-13 lists the types of errors in the error field of the channel status register. Multiple errors are possible. If any of these bits are set, a channel error interrupt is generated (see [Section 18.5.2.2, “Channel Error Interrupt”](#)). Most errors also halt the channel. For some error types, the host must take actions to clear the error bit before restarting the channel, as described in [Table 18-13](#). See the R and CON bits of the CCR ([Section 18.5.4.1, “Channel Configuration Register \(CCR\)”](#)) for information on restarting a channel.

Table 18-13. CSR Error Field Definitions

| Value | Error | Description |
|-------|-------|--|
| 48 | DOF | Double fetch FIFO write overflow error. This bit is set when the channel fetch FIFO is full, SOF is set, and another write has been made to the fetch FIFO. This error halts the channel. To clear this error, the host must write a 1 to this bit. |
| 49 | SOF | Single fetch FIFO write overflow error. This bit is set when the channel fetch FIFO is full and another write has been made to the fetch FIFO. The channel continues processing, but the descriptor pointer is lost. To clear this error, the host must write a 1 to this bit. |
| 50 | MDTE | Master data transfer error. When the SEC, while acting as a bus master, detects an error, the controller passes this error to the channel. This error halts the channel. Restarting the channel clears this bit. |
| 51–52 | — | Reserved |

Table 18-13. CSR Error Field Definitions (continued)

| Value | Error | Description |
|-------|-------|--|
| 53 | IDH | Illegal descriptor header. Possible causes of an illegal descriptor header are: <ul style="list-style-type: none"> • Invalid primary EU indicated by op0 field in descriptor header. • Invalid secondary EU indicated by op1 field in descriptor header. This error halts the channel. Restarting the channel clears this bit. |
| 54 | — | Reserved |
| 55 | EUE | EU error. An EU assigned to this channel has generated an error interrupt. This error may also be reflected in the controller's interrupt status register. This error halts the channel. To clear this error, the host must clear the error source in the EU that produced the error. |
| 56 | WDT | Watchdog timeout. The main state machine stayed asleep too long. This timer runs only after EUs have been reserved, and does not run if the primary EU is the RNGU or PKEU. The timeout interval is controlled by the FCC field of the Channel Configuration Register. This error halts the channel. Restarting the channel clears this bit. |
| 57 | SGLM | Scatter/gather length mismatch. Indicates the total data size covered by a gather link table did not match the total data size from the main descriptor. This error halts the channel. Restarting the channel clears this bit. |
| 58 | RSI | RAID Size Incorrect. The channel was provided with a descriptor of type RAID_XOR with data sizes not permitted. To clear this error, the host must write a '1' to this bit. |
| 59 | RSG | RAID scatter gather error. The channel was provided with a descriptor of type RAID_XOR with a j bit set. Use of scatter/gather is not permitted with RAID_XOR type descriptors. To clear this error, the host must write a '1' to this bit. |

18.5.4.3 Current Descriptor Pointer Register (CDPR)

The CDPR, shown in [Figure 18-13](#), reflects the value of the head-end of the fetch FIFO, which contains the address of the descriptor which the channel is currently processing.




Figure 18-13. Current Descriptor Pointer Register

The bits in the current descriptor pointer register perform the functions described in [Table 18-14](#).

Table 18-14. Current Descriptor Pointer Register Signals

| Bits | Name | Description |
|-------|--------|--|
| 0–15 | ID_TAG | Descriptor identification tag. Used by software. |
| 16–27 | — | Reserved, set to zero. |

Table 18-14. Current Descriptor Pointer Register Signals (continued)

| Bits | Name | Description |
|-------|------------------|--|
| 28–31 | EPTR | Extended pointer. Concatenated as the top 4 bits of the CUR_DES_PTR_ADRS when EAE is high (see the EAE bit in Table 18-11). |
| 32–63 | CUR_DES_PTR_ADRS | Current descriptor pointer address. Pointer to system memory location of the current descriptor. This field reflects the starting location in system memory of the descriptor currently loaded into the DB. This value is updated whenever the channel requests a fetch of a descriptor from the controller. The value from the fetch FIFO is transferred to the current descriptor pointer register immediately after the fetch is completed. This address will be used as the destination for writeback of the modified header dword, if header writeback notification is enabled. |


18.5.4.4 Fetch FIFO Enqueue Register (FFER)

Each channel contains a fetch FIFO to store a queue of pointers to descriptors which the channel will process. A pointer is added to the queue by writing to the FFER.

The register is shown in [Figure 18-14](#), and the fields are described in [Table 18-15](#).

Offset Channel 1: 0x3_01148, Channel 2: 0x3_01248,
Channel 3: 0x3_01348, Channel 4: 0x3_01448

Access: Write only


Figure 18-14. Fetch FIFO Enqueue Register (FFER)

[Table 18-15](#) describes the fetch FIFO fields.

Table 18-15. Fetch FIFO Field Descriptions

| Bits | Name | Description |
|-------|-----------|--|
| 0–15 | ID_TAG | Identification tag. A 16-bit value assigned by the host to a descriptor that will be written back to the ID_TAG field of the descriptor header (see Figure 18-3) whenever a status writeback occurs (per IWSE and AWSE bits of channel configuration register). |
| 16–27 | — | Reserved, set to zero. |
| 28–31 | EPTR | Extended pointer. Concatenated as the top 4 bits of the FETCH_ADRS when EAE is high (see the EAE bit in Table 18-11). |
| 32–63 | FETCH_ADR | Fetch address. Pointer to system memory location of a descriptor the host wants the SEC to fetch. |

In channel-driven access, the host CPU creates a descriptor in memory containing all relevant mode and location information for the SEC, then launches the descriptor by writing its address to the fetch FIFO enqueue register.

The fetch FIFO can hold up to 24 descriptor pointers at a time. When the current descriptor's processing is finished, the next fetch FIFO entry is read and the descriptor located at FETCH_ADR is launched.

NOTE

When extended addresses are enabled (by setting the EAE bit in channel configuration register), then the FFER’s EPTR field must be written before or concurrently with the FETCH_ADR field. This is necessary because writing the least significant byte (bits 56–63) is the “trigger” which causes the FFER contents to be added to the FIFO.

18.5.5 Channel Buffers and Tables

Besides the registers described in Section 18.5.4, “Channel Registers,” each channel has memory allocated for descriptors and scatter/gather link table entries (described in Section 18.3, “Descriptors”). The following subsections describe these features.

18.5.5.1 Descriptor Buffer (DB)

The descriptor buffer (DB) provides read-only access to the descriptor currently being processed by the channel. All descriptors are 8 dwords long. For descriptor format, see Figure 18-2. The address ranges of each channel’s DB are shown in Table 18-3.

Note that the DB is working storage and the channel may modify the contents of the DB during processing. In debug scenarios, it may be useful to read the contents of the DB to determine if a well formed descriptor is being fetched by the channel. Potential causes of malformed descriptors in the DB include:

- The descriptor is built incorrectly
- The descriptor is fully or partially overwritten by some other system bus master before the SEC can fetch the descriptor
- The descriptor is not built at the address written to the fetch FIFO

18.5.5.2 Scatter and Gather Link Tables (SLT, GLT)

A pointer dword in the descriptor buffer (DB) refers to a Gather Link Table (GLT) or a Scatter Link Table (SLT) if the J bit in the dword is set. As a channel works on a DB pointer entry, the GLT/SLT is loaded into channel memory. Reads from the GLT/SLT are enabled for debug purposes.

Figure 18-15 summarizes the entry format and address ranges for gather and scatter link table entries.

| | | | | | | | | | | | | |
|-------|--|----|----|----|----|----|------|----|--------|----|----|----|
| | 0 | 15 | 16 | 21 | 22 | 23 | 24 | 27 | 28 | 31 | 32 | 63 |
| Field | SEGLEN | | -- | R | N | — | EPTR | | SEGPTR | | | |
| Reset | (not reset) | | | | | | | | | | | |
| R/W | R/W | | | | | | | | | | | |
| Addr | Channel 1: 0x3_11c0-0x3_11df (Gather); 0x3_11e0-0x3_11ff (Scatter) Channel 2: 0x3_12c0-0x3_12df (Gather); 0x3_12e0-0x3_12ff (Scatter) Channel 3: 0x3_13c0-0x3_13df (Gather); 0x3_13e0-0x3_13ff (Scatter) Channel 4: 0x3_14c0-0x3_14df (Gather); 0x3_14e0-0x3_14ff (Scatter) | | | | | | | | | | | |

Figure 18-15. Gather/Scatter Link Table Entry Format and Memory Ranges

18.6 Controller

All transfers between the host and the EUs are moderated by the controller. Some of the main functions of the controller are as follows:

- Accept and execute commands from the slave system bus to read or write memory-mapped locations (up to 64 bits) anywhere in the SEC.
- Accept and execute requests from the polychannel to transfer blocks of bytes among system memory, EUs, and the channels.
- Arbitrate between channels when they contend for EUs and bus access
- Realign read and write data to the proper byte alignment
- Monitor interrupts from channels and pass them to the host

The remainder of this section discusses the controller's bus management, arbitration, interrupts, and registers.

18.6.1 Bus Transfers

As shown in [Figure 18-1](#), the SEC has an internal bus and connects to the SoC's system bus. The internal bus is a private 64-bit slave bus, with the controller block as the sole master. The SoC's system bus actually refers to two buses: a slave bus and a master bus, for which SEC's controller operates as slave and master, respectively. All accesses to SEC over the system bus go through the controller.

As mentioned in [Section 18.1.3, "Controller Overview"](#), there are two modes of access to the SEC, depending on whether the SEC's controller is slave or master. These two modes of access (host-controlled and channel-controlled) are discussed in the following subsections.

18.6.1.1 Host-Controlled Access

For host-controlled access, the host uses the SoC's slave bus to access the controller as a slave, and the controller relays the read or write accesses over the internal bus to the appropriate registers and FIFOs of the EUs. When a write command is received from the system bus, the controller takes the data and sends it to whichever internal location is indicated by the address. For a read, the controller goes to the internal location, fetches the requested data from the specified address (if allowed), and returns it over the system bus.

Host-controlled access is much more CPU-intensive than channel-controlled access, and requires a great deal of familiarity with the EU and controller registers and procedures. If host-controlled access is used, it is recommended that only a single EU be operated at a time. Snooping is not available through this interface.

NOTE

Host-controlled access of execution units is provided primarily for system debug purposes. The SEC contains no mechanism to arbitrate between host and channel accesses to EUs. Simultaneous use of an execution unit by a channel and a host is liable to force the execution unit into an error condition.

18.6.1.2 Channel-Controlled Access

Channel-controlled access is the SEC's normal operating mode. The controller performs data transfers based on information from the channels' descriptors. The controller can queue up to four requests. The controller dequeues requests and performs the required transfer. Most transfers involve not only the internal bus, but also the SoC's master bus with the controller as bus master.

When the SEC performs a read or write transaction as master, in some cases the intended target (for instance, system memory) may terminate the transaction due to an error. Once the transaction is posted to the SoC's target queue, it is the SoC's responsibility to either complete the transaction or signal an error. An error in an SEC-initiated transaction is also reported by the SEC through the channel interrupt status register (ISR). The host is able to determine which channel generated the interrupt by checking the ISR for the channel ERROR bit.

18.6.1.2.1 Channel Controlled Read—Detailed Description

A detailed description for a system bus read with controller as master is as follows:

1. Channel asserts bus read request to the controller
2. Channel furnishes external read address, internal write address, and transfer length
3. Controller asserts request to the system bus through the master interface
4. Controller waits for system bus read to begin
5. When bus read begins, controller receives data from the master interface and performs a write to the appropriate internal address supplied by the channel. Data may be realigned byte-wise by the controller if either:
 6. the external read address was not on an 8-byte boundary, or
 7. the internal write address was not on an 8-byte boundary.
8. Transfer continues until the bus read is completed and the controller has written all data to the appropriate internal address. The master interface continues making bus requests until the full data length has been read.

18.6.1.2.2 System Bus Master Write—Detailed Description

A detailed description for a system bus write with controller as master is as follows:

1. Channel asserts its bus write request to the controller.
2. Channel furnishes internal read address, external write address, and transfer length.
3. Controller performs a read from the appropriate internal address supplied by the channel, loads the write data into its FIFO, asserts a request to the system bus through the master interface, and waits for the system bus to become available.
4. When the system bus becomes available, controller writes data from its FIFO to the master interface.

18.6.2 Arbitration Algorithms

This section applies to both arbitration for use of the polychannel, and arbitration for use of execution units. Control fields for both are in the master control register (Section 18.6.4.7, “Master Control Register (MCR)”), as follows:

- CHN3_BUS_PR_CNT and CHN4_BUS_PR_CNT control polychannel arbitration
- CHN3_EU_PR_CNT and CHN4_EU_PR_CNT control EU arbitration

This section refers to generic control fields CHN3_XX_PR_CNT and CHN4_XX_PR_CNT, where “XX” refers to either “BUS” or “EU”.

If both CHN3_XX_PR_CNT and CHN4_XX_PR_CNT are zero (the default), the arbitration is round-robin (see Section 18.6.2.1); otherwise a weighted priority scheme is used (see Section 18.6.2.2).

18.6.2.1 Round-Robin Arbitration

In round-robin arbitration, requesting channels are granted access in rotating numerical order: 1, 2, 3, 4, 1, 2, ... and so on.

18.6.2.2 Weighted Priority Arbitration

In the weighted priority scheme, the priority is as follows:

- Channel 1—Highest priority
- Channel 2—Second highest priority, unless CHN3_XX_PR_CNT or CHN4_XX_PR_CNT has expired
- Channel 3—Third priority, unless CHN4_XX_PR_CNT expired
- Channel 4—Lowest priority, until CHN4_XX_PR_CNT expired

Initially, the priority is fixed from highest to lowest as channel 1, channel 2, channel 3, and channel 4, in that order. When channel 3 has lost arbitration the number of times specified in CHN3_XX_PR_CNT, channel 3 replaces channel 2 as the second-highest priority in the next round of arbitration. Likewise, when channel 4 has lost arbitration the number of times specified in CHN4_XX_PR_CNT, channel 4 replaces channel 2 as the second-highest priority in the next round of arbitration. These rules prevent channels 3 and 4 from being locked out.

Channel 1 always has the highest priority, but cannot make back-to-back requests. It follows that the second highest priority channel wins arbitration either immediately, or after one win for channel 1.

Note that the SEC does not dynamically adjust its own transaction priorities. System software, however, can adjust SEC transaction priority in real time, with the change in priority taking effect immediately.

18.6.3 Controller Interrupts

18.6.3.1 Controller Interrupt Conditions and Interrupt Generation

All interrupt outputs from other SEC blocks are fed to the controller as interrupt conditions. In addition, the controller itself detects some interrupt conditions. The controller maintains an interrupt status register

(ISR) with bits corresponding to all of these possible interrupt conditions. If an interrupt condition occurs and the corresponding bit of the interrupt enable register (IER) is set, then the associated ISR bit is set, indicating the presence of a pending interrupt.

A channel can generate frequent interrupts, especially if it is configured to interrupt at the completion of each descriptor. To make sure that the host receives the right number of interrupts, each channel done interrupt has a special “queuing” feature. If multiple channel done interrupts are generated before the first is cleared, then the additional interrupts are counted by the controller. Each time the host clears a channel interrupt, the count is decremented. If the host clears the channel interrupt and the count reaches zero, the channel done interrupt is negated. If the count does not reach zero, the controller negates the interrupt for one cycle and then re-asserts it.

Up to 15 interrupts can be queued for each channel. If the count of queued interrupts for any channel exceeds 15, then that channel’s done overflow bit is set in the channel’s ISR (if the corresponding IER bit is set), and the channel done interrupt is asserted.

18.6.3.2 Blocking of Interrupts

Interrupt conditions from the channels and controller can only be blocked through the controller’s IER, as described in Section 18.6.3.1. However, the EU interrupt conditions may be blocked at two different levels. There is an interrupt mask register in each EU which can block particular interrupt conditions before they reach the EU’s interrupt status register. In addition, interrupts from EUs can be individually blocked by bits of the controller’s IER before they reach the controller’s ISR. For normal operation, interrupts from EUs are typically disabled in the controller’s IER, but they still reach the channel, and the channel produces done or error interrupts to the host as needed.

18.6.3.3 Interrupt Handling

To handle an interrupt, the host must read the ISR to determine the source. If necessary, the host may read the interrupt status registers of other blocks to ascertain the cause. In some cases, the host may need to take action to clear the root cause of the interrupt. Once the appropriate action is taken, the host can clear the ISR bit by setting the corresponding bit of the interrupt clear register (ICR). If the cause of the interrupt condition has not been cleared, or if there is another interrupt condition from the same source, then the ISR bit clears for a cycle and then goes high again, and the interrupt signal to the host remains high. If the ISR bit is successfully cleared and no other interrupt conditions are present, the controller negates its interrupt signal. If any interrupts are still pending in the ISR, the interrupt remains asserted.

18.6.4 Controller Registers

The controller registers are described in detail in the following sections.

18.6.4.1 EU Assignment Status Register (EUASR)

The EUASR, displayed in [Figure 18-16](#), indicates which EUs are reserved by a particular channel. When an EU is already assigned, it is inaccessible to any other channel.

A 4-bit field (see [Table 18-16](#)) indicates the channel to which an EU is assigned.

Offset 0x3_1028 Access: Read only

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------|----|----|----|----|----|----|----|----|------|----|----|----|----|----|----|----|------|---|---|---|------|---|---|---|-----|---|---|---|-----|---|---|---|
| | 0 | 3 | 4 | 7 | 8 | 11 | 12 | 15 | 16 | 19 | 20 | 23 | 24 | 27 | 28 | 31 | | | | | | | | | | | | | | | | |
| R | — | | | | — | | | | MDEU | | | | — | | | | AESU | | | | — | | | | DEU | | | | | | | |
| W | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Reset | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 |
| | 32 | 35 | 36 | 39 | 40 | 43 | 44 | 47 | 48 | 51 | 52 | 55 | 56 | 59 | 60 | 63 | | | | | | | | | | | | | | | | |
| R | — | | | | — | | | | CRCU | | | | — | | | | — | | | | PKEU | | | | — | | | | RNG | | | |
| W | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Reset | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 |

Figure 18-16. EU Assignment Status Register (EUASR)

Table 18-16. Channel Assignment Value

| Value | Channel |
|---------|---------------------|
| 0x0 | No channel assigned |
| 0x1 | Channel 1 |
| 0x2 | Channel 2 |
| 0x3 | Channel 3 |
| 0x4 | Channel 4 |
| 0xA–0xE | Undefined |
| 0xF | Unavailable |

18.6.4.2 Interrupt Enable Register (IER)

Interrupt sources can be individually enabled by setting the corresponding IER bits (see [Table 18-17](#) for the correspondence between IER bits and interrupt sources). If an IER bit is set, the corresponding interrupt source value is captured in the corresponding interrupt status register (ISR) bit. If an IER bit is cleared, the corresponding ISR bit remains cleared.

At reset, all IER bits are cleared, so all interrupts are disabled.

[Figure 18-17](#) shows the bit positions of each potential interrupt source.

NOTE

For normal operation the IER should be programmed with the value 0x0031_0fff_0000_0000, which enables all channel interrupts and disables interrupts from the EUs. The EU interrupt bits are provided as a convenience during debug: during normal operation, an EU error causes the channel using that EU to generate the appropriate interrupt to the host.




Figure 18-17. Interrupt Enable Register (IER)

Table 18-17 describes the register fields in the interrupt enable register, interrupt status register, and interrupt clear register.

Table 18-17. Fields in Interrupt Enable, Interrupt Status, and Interrupt Clear Registers

| Bits | Name | Description |
|------|--------|---|
| 10 | DI_CNT | Data in count rollover. 0 No rollover. 1 The data in counter rolled over to zero (see Section 18.5.3.1.3, "Data Bytes In Counter"). |
| 11 | DO_CNT | Data out count rollover. 0 No rollover. 1 The data out counter rolled over to zero (see Section 18.5.3.1.4, "Data Bytes Out Counter"). |
| 15 | ITO | Internal time out. 0 No internal time out 1 An internal time out was detected Note: Internal time out is an indication that a channel or EU has failed to respond to a slave read or write within 16 cycles, which would only occur in an impending hang condition. Assertion of this interrupt indicates the SEC Controller has completed the transaction to avoid a hang, however, the completed transaction does not result in a successful read or write, and the interrupt advises the system that the slave transaction was unsuccessful. |

Table 18-17. Fields in Interrupt Enable, Interrupt Status, and Interrupt Clear Registers (continued)

| Bits | Name | Description |
|---|--|---|
| 20–23 | Done Overflow | Overflow. Done overflow (one bit for each channel, CH1–CH4) 0 No done overflow 1 Done overflow error. Indicates that more than 15 done interrupts were queued from the associated channel without an interrupt clear from the host. |
| 24–31 | Err and Dn bits for channels (CHN_1–CHN_4) | Error. 0 No error detected. 1 Error detected. Indicates that channel status register must be read to determine exact cause of the error. Done. 0 Not DONE. 1 DONE bit indicates that the corresponding channel has completed a descriptor. |
| 36–39, 42–43, 46–47, 50–51, 54–55, 58–59, 62–63 | Err and Dn bits for execution units (PKEU, and so on.) | Error. 0 No error detected. 1 Error detected. Indicates that execution unit status register must be read to determine exact cause of the error. Done. 0 Not Done 1 DONE bit indicates that the corresponding EU has completed its operation. This means that final values are available from EU registers. For EUs with output FIFOs, it means that all textdata output has been placed in the output FIFO. For EUs that provide context out through the output FIFO, the EU places the context in the output FIFO after asserting PRI_DONE. |
| 0–9, 16–19, 32–35, 40–41, 44–45, 48–49, 52–53, 56–57, 60–61 | — | Reserved. Clear to zero. |

18.6.4.3 Interrupt Status Register (ISR)

Each ISR bit shows the status of a corresponding interrupt source (see [Table 18-17](#) for the correspondence between ISR bits and interrupt sources). However, if the corresponding IER bit is cleared, then the ISR bit remains cleared.

18.6.4.6 IP Block Revision Register

The read-only IP block revision register, displayed in [Figure 18-21](#), contains a 64-bit value that uniquely identifies the version of the SEC 3.0. The value of this register is 0x0030_0300_0000_0000.




Figure 18-21. IP Block Revision Register

[Table 18-18](#) describes the fields of the IP block revision register.

Table 18-18. IP Block Revision Register Fields

| Bits | Name | Description |
|-------|--------|--------------------------------|
| 0–15 | IP_ID | IP block identifier |
| 16–23 | IP_MJ | IP major revision number |
| 24–31 | IP_MN | IP minor revision number |
| 32–39 | — | Reserved |
| 40–47 | IP_INT | IP block integration options |
| 48–55 | — | Reserved |
| 56–63 | IP_CFG | IP block configuration options |

18.6.4.7 Master Control Register (MCR)

The MCR, shown in [Figure 18-22](#), controls certain functions in the controller and provides a means for software to reset the SEC.




Figure 18-22. Master Control Register (MCR)

[Table 18-19](#) describes the MCR bit fields.

Table 18-19. MCR Bit Descriptions

| Bits | Name | Description |
|-------|----------------|--|
| 0–21 | — | Reserved |
| 22–23 | Priority | Priority on master bus. The setting of these bits determines the transaction priority level the SEC asserts to the platform internal arbiter. The SEC does not dynamically alter its priority level based on system congestion or SEC utilization, however software may change the SEC priority level in realtime. 00 Lowest priority (default) 01 Next lowest priority 10 Next highest priority 11 Highest priority |
| 24–29 | — | Reserved |
| 30 | GIH | Global Inhibit. Writing 1 to this bit indicates that external master bus transfers are defined as not snoopable and results in lowering the snoop attribute of bus requests generated by the external gasket. 0 External master bus transfers are defined as snoopable (default) 1 External master bus transfers are defined as not snoopable |
| 31 | SWR | Software reset. Writing 1 to this bit will cause a global software reset. Upon completion of the reset, this bit will be automatically cleared. 0 Do not reset 1 Global Reset |
| 32–39 | CHN3_EU_PR_CNT | Channel 3 EU priority counter. This counter is used by the controller to determine when Channel 3 has been denied access to a requested EU long enough to warrant immediate elevation to top priority. Note: If set to zero, the CHN4_EU_PR_CTR must also be set to zero, and the controller will assign EU's on a pure round robin basis. |

Table 18-19. MCR Bit Descriptions (continued) (continued)

| Bits | Name | Description |
|-------|-----------------|---|
| 40–47 | CHN4_EU_PR_CNT | Channel 4 EU priority counter. This counter is used by the controller to determine when Channel 4 has been denied access to a requested EU long enough to warrant immediate elevation to top priority. Note: If set to zero, the CHN3_EU_PR_CTR must also be set to zero, and the controller will assign EU's on a pure round robin basis. |
| 48–55 | CHN3_BUS_PR_CNT | Channel 3 Bus priority counter. This counter is used by the controller to determine when Channel 3 has been denied access to the polychannel long enough to warrant immediate elevation to top priority. Note: If set to zero, the CHN4_BUS_PR_CTR must also be set to zero, and the controller will assign access to the polychannel on a pure round robin basis. If set to non-zero, CHN4_BUS_PR_CTR must also be set to a different, non-zero value. |
| 56–63 | CHN4_BUS_PR_CNT | Channel 4 Bus priority counter. This counter is used by the controller to determine when Channel 4 has been denied access to the polychannel long enough to warrant immediate elevation to top priority. Note: If set to zero, the CHN3_BUS_PR_CTR must also be set to zero, and the controller will assign access to the polychannel on a pure round robin basis. If set to non-zero, CHN3_BUS_PR_CTR must also be set to a different, non-zero value. |

NOTE

By default, all SEC memory transactions are snooped by the e300 core. CPU snooping of SEC memory transactions can be disabled by setting the global inhibit bit (GIH) in the MCR.

SEC transactions can be snooped by the cache if defined as global. This definition is programmed in the master control register MCR[GI]. See [Section 18.6.4.7, “Master Control Register \(MCR\),”](#) for more details. Note that SEC transactions are defined as global by default.

18.7 Power Saving Mode

The SEC may be disabled by setting DEVDISR[SEC] in the SoC. The clocks to the SEC are active by default. The SEC should not be enabled/disabled during normal operation.

SEC disablement is delayed if the disable request is made while descriptors are being processed. Once notified of the disable request, the SEC channels complete their current tasks, and then are forced to idle (with no additional reads from the fetch descriptor FIFO). Once all channels are idle, then SEC permits disablement.

18.8 Execution Units

Execution unit (EU) is the term used for a functional block that performs the mathematical manipulations required by protocols used in cryptographic processing. The following execution units are used in the SEC (covered here in alphabetical order):

- Advanced Encryption Standard Execution Unit (AESU) implementing the Rijndael symmetric key cipher.

- Cyclical Redundancy Check Accelerator (CRCU)
- Data Encryption Standard Execution Unit (DEU)
- Message Digest Execution Unit (MDEU)
- Public Key Execution Unit (PKEU)
- Random Number Generator (RNGU)

Working together, the EUs can perform high-level cryptographic tasks, such as IPsec Encapsulating Security Protocol (ESP) and digital signature. The remainder of this chapter provides details about the execution units themselves.

18.8.1 Advanced Encryption Standard Execution Unit (AESU)

This section contains details about the Advanced Encryption Standard Execution Unit (AESU), including modes of operation, status and control registers, and FIFOs.

NOTE

Most of the registers described in this section are not accessed by the host under normal operation. They are documented here mainly for debug purposes. Normally the AESU is used through channel-controlled access, so that most reads and writes of AESU registers are directed by the SEC channels. Driver software performs host-controlled register accesses only on a few registers for initial configuration and error handling.

18.8.1.1 ICV Checking in AESU

For CCM, GCM, CMAC (OMAC1), and XCBC-MAC cipher modes, the AESU includes an ICV checking feature which can generate an ICV and compare it to another supplied ICV.

There are two methods for returning the pass/fail result of ICV checking to the host:

- The ICV check result can be sent to the host by a writeback of EU status fields into host memory. This is enabled as follows:
 - Set either the IWSE or AWSE bit in the channel configuration register (see [Section 18.5.4.1, “Channel Configuration Register \(CCR\)”](#))
 - Set the ICE bit in the interrupt mask register ([Section 18.8.1.8, “AESU Interrupt Mask Register”](#)).

In this case the normal done signaling (by interrupt or writeback) is undisturbed.

- The ICV checking result can be sent to the host by interrupt. This is enabled as follows:
 - Clear the ICE bit in the interrupt mask register
 - Clear both IWSE and AWSE bits in the channel configuration register.

In this case, then the normal done signaling (by interrupt or writeback) occurs if there is no ICV mismatch. If an ICV mismatch occurs, then an error interrupt is sent to the host, but no channel done interrupt or writeback.

Table 18-20. AESU Mode Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 59 | AUX1 | <p>AUX1 Mode has a definition that depends upon the value of the 4 Cipher Mode and Cipher Mode Extend bits:</p> <ul style="list-style-type: none"> • CCM Cipher Mode (CME=10, CM=00): Initialize Mode Bit - Initializes AESA for new message <ul style="list-style-type: none"> 0 = Do not initialize (context will be loaded by host) 1 = Initialize new message with nonce/initialization vector • GCM Cipher Mode (CME=10, CM=01): Generate Final GHASH Bit - Enables completion of GHASH computation by signaling that the last iteration of GHASH should be performed. This last iteration will perform XOR of the current (intermediate) GHASH result with the concatenation of AAD and ciphertext bit lengths in case of GHASH(H, AAD, ciphertext), or with the concatenation of 0^{64} and the bit length of IV in case of GHASH(H, {}, IV). As an exception, this bit should be cleared if the whole message (IV+AAD+textdata) together with the generation of the final MAC is processed with one descriptor since in that case the generation of final GHASH is implied. Incidentally, whenever AUX1=1 in GCM cipher mode, the total bit lengths of AAD, textdata or IV must be provided in context registers 9-10. <ul style="list-style-type: none"> 0 = Do not perform the last iteration in GHASH(H, AAD, ciphertext) or GHASH(H, {}, IV) unless the message is processed and the final MAC computed in 1 descriptor. 1 = Generate the final result of GHASH(H, AAD, ciphertext) or GHASH(H, {}, IV) - implies that the message processing is split into multiple descriptors. • XCBC-MAC Cipher Mode (CME=10, CM=10): Load Keys - Do not compute K1, K2 and K3, but instead use the keys loaded in the Key Data Registers (K1), and Context Registers 5-6 (K2) and 7-8 (K3). <ul style="list-style-type: none"> 0 = Compute $K1=E(K, 16\{01\})$, $K2=E(K, 16\{02\})$, $K3=E(K, \{03\})$ and write K1 to Context Registers 3-4, K2 to 5-6, and K3 to 7-8. 1 = Load keys: $K1= [Key\ Data\ Reg\ 1-2]$, $K2= [Reg\ 5-6]$, $K3=[Reg\ 7-8]$ • CMAC Cipher Mode (CME=01, CM=10): Load Keys - Do not compute $E(K, 0^{128})$ to derive K1 and K2, but instead use the value loaded in Context Registers 3-4. This is useful after a context switch. Deriving K1 and K2 does not incur any timing penalty. <ul style="list-style-type: none"> 0 = Compute $E(K, 0^{128})$ and write it to Context Registers 3-4 1 = Load $E(K, 0^{128})$ and preserve it in Context Registers 3-4 |
| 60 | AUX0 | <p>AUX0 Mode has a definition that depends upon the value of the 4 Cipher Mode and Cipher Mode Extend bits, and Encrypt/Decrypt bit:</p> <ul style="list-style-type: none"> • ECB, CBC, OFB, CFB Cipher Mode: Restore Decrypt Key Bit - Specifies that key data write will contain pre-expanded key (decrypt mode only). See note below on use of RDK bit. <ul style="list-style-type: none"> 0 = Expand user key prior to decrypting first block 1 = Do not expand key. Expanded decryption key will be written following context switch. • GCM Cipher Mode (CME=10, CM=01) and Encrypt (Encrypt/Decrypt=1): Specifies GHASH mode—performs GHASH on AAD and ciphertext <ul style="list-style-type: none"> 0 = Perform GCM encryption 1 = Compute GHASH(H, AAD, ciphertext) • XCBC-MAC, CMAC Cipher Mode (CME=10, 01, CM=10): Do Not Generate Final MAC Bit - Does not generate final MAC tag at the end of message processing (used only when splitting a message into multiple descriptors) <ul style="list-style-type: none"> 0 = Generate final MAC tag that is XOR the final data block with K2/K3 (XCBC-MAC), or K1/K2 (CMAC) before encryption 1 = Do not generate final MAC tag that is does not XOR final data block with K2 (XCBC-MAC) or K1 (CMAC) before encryption. This enables message processing to be interrupted on the block boundary and later continued after a context switch |

Table 18-20. AESU Mode Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 61–62 | CM | Cipher Mode: Used in combination with bits 56:57 (Extended Cipher Mode) to select the cipher mode for AES operation. See Table 18-21 on page 18-56 for mode bit combinations. |
| 63 | ED | Encrypt/Decrypt. If set, AESU operates the encryption algorithm; if not set, AESU operates the decryption algorithm. 0 Perform decryption 1 Perform encryption Note: This bit is ignored in CTR, SRT, CMAC, and XCBC-MAC Cipher Modes. |

Table 18-21. AES Cipher Modes

| Cipher Mode | ECM (56–57) | AUX2 (58) | CM (61–62) |
|------------------|-------------|-----------|------------|
| ECB | 00 | X | 00 |
| CBC | 00 | X | 01 |
| CBC-RBP | 00 | 1 | 01 |
| OFB | 00 | X | 10 |
| CTR | 00 | X | 11 |
| CMAC | 01 | X | 10 |
| CMAC w/ICV | 01 | 1 | 10 |
| SRT ¹ | 01 | X | 11 |
| CCM | 10 | X | 00 |
| GCM | 10 | X | 01 |
| XCBC-MAC | 10 | 0 | 10 |
| XCBC-MAC w/ICV | 10 | 1 | 10 |
| CFB128 | 10 | X | 11 |
| CCM w/ICV | 11 | X | 00 |
| GCM w/ICV | 11 | X | 01 |
| XOR | 11 | X | 11 |
| Reserved | All Others | | |

¹ SRT is not a new AES cipher mode, it is an AESU method of performing AES-CTR mode with reduced context loading overhead specifically for performing SRTP. It should be used with descriptor type 0010_0'srtp'. See Section 18.8.1.11.3 on page 18-66 for more information on how SRT cipher mode reduces context loading overhead.

18.8.1.3 AESU Key Size Register

The AESU key size register is used to specify the number of bytes in the key (16, 24, or 32). Any key data beyond the number of bytes specified in the key size register is ignored. This register is cleared when the AESU is reset or re-initialized. If a key size other than 16, 24, or 32 bytes (or other than 16 bytes, in XCBC-MAC cipher mode) is specified, an illegal key size error is generated. If the key size register is modified during processing, a context error is generated.




Figure 18-24. AESU Key Size Register

18.8.1.4 AESU Data Size Register

The AESU data size register, shown in Figure 18-25, is used to specify the number of bits (not bytes) of plaintext/ciphertext to be processed in the current descriptor. The number of data size register bits used by the SEC, and the acceptable values for these bits, vary depending on the AES cipher mode selected as specified in Table 18-22.

Writing to this register signals the AESU to start processing data from the input FIFO as soon as it is available. If the value of data size is modified during processing, a context error is generated. The register is cleared when the AESU is reset or re-initialized.




Figure 18-25. AESU Data Size Register

Table 18-22. Use of Data Size Register

| AESU Cipher Mode | Register bits used by SEC (others are don't cares) | Legal Values (data size in bits) |
|---------------------------------------|--|----------------------------------|
| ECB, CBC | lowest 7 bits [57:63] | must be a multiple of 128 |
| OFB, CMAC, SRT, CCM, XCBC-MAC, CFB128 | | must be a multiple of 8 |
| GCM | all bits | any value |
| XOR | all bits | must be a multiple of 256 |



Figure 18-27. AESU Status Register

Table 18-24 describes the AESU status register fields.

Table 18-24. AESU Status Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–39 | — | Reserved |
| 40–47 | OFL | The number of dwords currently in the output FIFO |
| 48–55 | IFL | The number of dwords currently in the input FIFO |
| 56–57 | — | Reserved |
| 58 | HALT | Halt. Indicates that the AESU has halted due to an error. 0 AESU not halted 1 AESU halted Note: Because the error causing the AESU to stop operating may be masked before reaching the interrupt status register, the AESU interrupt status register is used to provide a second source of information regarding errors preventing normal operation. |
| 59–60 | ICCR | Integrity check comparison result. 00 No integrity check comparison was performed. 01 The integrity check comparison passed. 10 The integrity check comparison failed. 11 Reserved A passed or failed result is generated only if the cipher mode with ICV checking is selected |
| 61 | EI | Error interrupt. This status bit reflects the state of the error interrupt signal, as sampled by the controller interrupt status register (Section 18.6.4.3, “Interrupt Status Register (ISR)"). 0 AESU is not signaling error 1 AESU is signaling error |
| 62 | DI | Done interrupt. This status bit reflects the state of the done interrupt signal, as sampled by the Controller Interrupt Status Register (Section 18.6.4.3, “Interrupt Status Register (ISR)"). 0 AESU is not signaling done 1 AESU is signaling done |
| 63 | RD | Reset done. This status bit, when high, indicates that AESU has completed its reset sequence, as reflected in the signal sampled by the appropriate channel. 0 Reset in progress 1 Reset done Note: Reset Done resets to 0, but has typically switched to 1 by the time a user checks the register, indicating the EU is ready for operation. |

18.8.1.7 AESU Interrupt Status Register

The AESU interrupt status register indicates which unmasked errors have occurred and have generated error interrupts to the channel. Each bit in this register can only be set if the corresponding bit of the AESU interrupt mask register is zero (see Section 18.8.1.8, “AESU Interrupt Mask Register”).

If the AESU interrupt status register is non-zero, the AESU halts and the AESU error interrupt signal is asserted to the controller (see [Section 18.6.4.3, “Interrupt Status Register \(ISR\)”](#)). In addition, if the AESU is operated through channel-controlled access, an interrupt signal is generated to the channel to which this EU is assigned. The EU error then appears in bit 55 of the channel status register (see [Table 18-13](#)) and generates a channel error interrupt to the controller.

If the AESU interrupt status register is written from the host, ones in the value written are recorded if the corresponding bit is unmasked in the interrupt mask register. All other bits are cleared. This register can also be cleared by setting the RI bit of the AESU reset control register. The definition of each bit in the AESU interrupt status register is shown in [Figure 18-28](#).

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|--------|---------------|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|-------------------|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| Offset | AESU 0x3_4030 | | | | | | | | | | | | | | | | Access: Read only | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| R | — | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | 48 | 49 | 50 | 51 | 52 | 53 | 54 | 55 | 56 | 57 | 58 | 59 | 60 | 61 | 62 | 63 |
| W | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Reset | All zeros | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

Figure 18-28. AESU Interrupt Status Register

[Table 18-25](#) describes AESU interrupt register fields.

Table 18-25. AESU Interrupt Status Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–48 | — | Reserved |
| 49 | ICE | Integrity check error. 0 No error detected 1 Integrity check error detected. An ICV check was performed and the supplied ICV did not match the one computed by the ASEU. |
| 50 | — | Reserved |
| 51 | IE | Internal error. An internal processing error was detected while the AESU was processing. 0 No error detected 1 Internal error Note: This bit will be asserted any time an enabled error condition occurs and can only be cleared by setting the corresponding bit in the Interrupt Mask Register or by resetting the AESU. |
| 52 | ERE | Early read error. An AESU context register was read while the AESU was processing. 0 No error detected 1 Early read error |
| 53 | CE | Context error. An AESU key register, the key size register, data size register, mode register, or IV register was modified while AESU was processing 0 No error detected 1 Context error |
| 54 | KSE | Key size error. An inappropriate value (not 16, 24 or 32bytes) was written to the AESU key size register. 0 No error detected 1 Key size error |

Table 18-25. AESU Interrupt Status Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 55 | DSE | Data size error (DSE): A value was written to the AESU data size register that is not a proper size. See Section 18.8.1.4, “AESU Data Size Register.” 0 No error detected 1 Data size error |
| 56 | ME | Mode error. Indicates that invalid data was written to a register or a reserved mode bit was set. 0 Valid data 1 Reserved or invalid mode selected |
| 57 | AE | Address error. An illegal read or write address was detected within the AESU address space. 0 No error detected 1 Address error |
| 58 | OFE | Output FIFO error. The AESU output FIFO was detected non-empty upon write of AESU data size register. 0 No error detected 1 Output FIFO non-empty error |
| 59 | IFE | Input FIFO error. The AESU input FIFO was detected non-empty upon generation of done interrupt. 0 No error detected 1 Input FIFO non-empty error |
| 60 | — | Reserved |
| 61 | IFO | Input FIFO overflow. The AESU Input fifo was pushed while full. 0 No error detected 1 Input FIFO has overflowed Note: When operated through channel-controlled access, the SEC implements flow control, and FIFO size is not a limit to data input. When operated through host-controlled access, the AESU cannot accept FIFO inputs larger than 256 bytes without overflowing. |
| 62 | OFU | Output FIFO underflow. The AESU Output FIFO was read while empty. 0 No error detected 1 Output FIFO has underflow error |
| 63 | — | Reserved |

18.8.1.8 AESU Interrupt Mask Register

The AESU interrupt mask register, shown in [Figure 18-29](#), controls the result of detected errors. For a given error (as defined in [Section 18.8.1.7, “AESU Interrupt Status Register”](#)), if the corresponding bit in this register is set, the error is ignored; no error interrupt occurs and the interrupt status register is not updated to reflect the error. If the corresponding bit is not set and an error is detected, the interrupt status register is updated to reflect the error, causing assertion of the error interrupt signal, and causing the module to halt processing.




Figure 18-29. AESU Interrupt Mask Register

Table 18-26 describes the AESU interrupt mask register fields.

Table 18-26. AESU Interrupt Mask Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–48 | — | Reserved |
| 49 | ICE | Integrity Check Error. The supplied ICV did not match the one computed by the ASEU. 0 Integrity check error enabled. WARNING: Do not enable this if using EU status writeback (see bits IWSE and AWSE in Section 18.5.4.1, “Channel Configuration Register (CCR”). 1 Integrity check error disabled |
| 50 | — | Reserved |
| 51 | IE | Internal error. An internal processing error was detected while the AESU was processing. 0 Internal error enabled 1 Internal error disabled |
| 52 | ERE | Early read error. The AESU IV register was read while the AESU was processing. 0 Early read error enabled 1 Early read error disabled |
| 53 | CE | Context error. An AESU Key Register, the Key Size Register, Data Size Register, Mode Register, or IV Register was modified while the AESU was processing. 0 Context error enabled 1 Context error disabled |
| 54 | KSE | Key size error. An inappropriate value (non 16, 24 or 32 bytes) was written to the AESU key size register 0 Key size error enabled 1 Key size error disabled |
| 55 | DSE | Data size error. Indicates that the number of bits to process is out of range. 0 Data size error enabled 1 Data size error disabled |
| 56 | ME | Mode error. Indicates that invalid data was written to a register or a reserved mode bit was set. 0 Mode error enabled 1 Mode error disabled |
| 57 | AE | Address error. An illegal read or write address was detected within the AESU address space. 1 Address error disabled 0 Address error enabled |
| 58 | OFE | Output FIFO error. the AESU Output FIFO was detected non-empty upon write of AESU data size register 0 Output FIFO non-empty error enabled 1 Output FIFO non-empty error disabled |

Table 18-26. AESU Interrupt Mask Register Field Descriptions (continued)


| Bits | Name | Description |
|------|------|---|
| 59 | IFE | Input FIFO error. The AESU Input FIFO was detected non-empty upon generation of done interrupt 0 Input FIFO non-empty error enabled 1 Input FIFO non-empty error disabled |
| 60 | — | Reserved |
| 61 | IFO | Input FIFO overflow. The AESU Input FIFO was pushed while full. 0 Input FIFO overflow error enabled 1 Input FIFO overflow error disabled |
| 62 | OFU | Output FIFO underflow. The AESU Output FIFO was read while empty. 0 Output FIFO underflow error enabled 1 Output FIFO underflow error disabled |
| 63 | — | Reserved |

18.8.1.9 ICV Size Register

The ICV size register, shown in [Figure 18-30](#), is used in AES hashing modes CMAC and GCM to specify the number of most significant bytes in the received MAC tag supplied in context registers 3–4. AES truncates the computed MAC in context registers 1–2 to the same number of bytes, and writes zeros in the remaining LSB's. It follows that the received MAC can be padded to 16 bytes with arbitrary data (not necessarily zeros) when written into context registers 3–4. Acceptable values for ICV size are 8, 10, 12, 14 and 16 bytes in CMAC, or 8, 12, and 16 bytes in GCM. All other sizes are interpreted as 16.

In XCBC-MAC cipher mode, the ICV size register is not used. The received MAC (written to context registers 9-10) is always truncated to the most significant 12 bytes, as defined in the XCBC-MAC-96 for IPsec specification. The computed MAC written at the end of processing to Context Registers 1-2 is a full 16-byte MAC.

In CCM mode with ICV, the ICV size register is not used. Instead, the tag size is encoded within one of the CCM formatting flags.


Figure 18-30. AESU ICV Size Register

18.8.1.10 AESU End of Message Register

The AESU end of message register, shown in [Figure 18-31](#), is used to indicate an AES operation may be completed. After the final message block is written to the input FIFO, the end of message register must be written. The value in the data size register is used to determine how many bits of the final message block (always 128) are processed. Writing to this register causes the AESU to process the final block of a message, allowing it to signal done interrupt. A read of this register always returns a zero value.




Figure 18-31. AESU End of Message Register

18.8.1.11 AESU Context Registers

Twelve 64-bit AESU context registers allow the host to read/write the contents of the context used to process a message. The context must be written prior to the key data. If the context registers are written during message processing, a context error is generated. All context registers are cleared when an initialization or a hard or soft reset is performed. If a message is processed through the AESU in two separate operations (for example, using two descriptors), the context must be read out of the SEC at the end of the first operation and then restored at the beginning of the second operation. Context is always read and restored as a contiguous subset of the twelve context registers ending with the highest numbered register used in that cipher mode. That is, if restoring context in CTR cipher mode, seven context registers (1–7) must be written where registers 1–4 must be filled with zeros. The context registers are summarized in [Table 18-27](#) and [Table 18-28](#).

Table 18-27. AESU Context Registers

| Context Register # (byte address) | Cipher Mode | | | | | | |
|-----------------------------------|-------------|--------------------------|----------------------|--|---------------------------|--------------|------------------------------|
| | ECB | CBC/CBC-R BP/OFB/CF B128 | CTR | CCM | GCM | GCM - GHASH | SRT |
| 1 (0x34100) | — | IV ¹ | -- | IV ¹ / MAC | Computed MAC | Computed MAC | Counter ¹ |
| 2 (0x34108) | | | | | | | |
| 3 (0x34110) | — | — | — | Encrypted MAC ² / Decrypted MAC / Encrypted Counter | Received MAC ² | — | Counter Modulus ¹ |
| 4 (0x34118) | | | | | | | — |
| 5 (0x34120) | — | — | Counter ¹ | Counter ¹ | Counter | — | — |
| 6 (0x34128) | | | | | | | |

Table 18-27. AESU Context Registers

| Context Register # (byte address) | Cipher Mode | | | | | | |
|--------------------------------------|-------------|--------------------------------|---|---|-----------------------|-----------------------|-----|
| | ECB | CBC/CBC-R BP/OFB/CF B128 | CTR | CCM | GCM | GCM - GHASH | SRT |
| 7 (0x34130) | — | — | Counter Modulus Exponent ¹ | Counter Modulus Exponent ¹ (header size/ MAC size) ³ | len(AAD) ⁴ | len(AAD) ⁴ | — |
| 8 (0x34138) | | | — | — | len(IV) ⁴ | — | |
| 9 (0x34140) | — | — | — | — | Y ₀ | H | — |
| 10 (0x34148) | | | | | | | |
| 11 (0x34150) | — | — | — | — | len(AAD) ⁵ | len(AAD) ⁵ | — |
| 12 (0x34158) | | | | | len(IV) ⁵ | | |

¹ Must be written at start of new message, except if zero

² Needed only in ICV mode—must be written at start of new CCM decryption

³ The header and MAC sizes are internally constructed by the AES engine; then, that information is included inside context register 7 for context switching purposes

⁴ Length of total data (in bits)

⁵ Length of data processed with current descriptor (in bits)

— Don't care

Table 18-28. AESU Context Registers

| Context Register | Cipher Mode | |
|------------------|--------------------------------|--------------------------------|
| | XCBC-MAC | CMAC |
| 1 | MAC | MAC |
| 2 | | |
| 3 | K1 ¹ | E(K, 0 ²) |
| 4 | | |
| 5 | K2 | MAC ³ (Received) |
| 6 | | |
| 7 | K3 | — |
| 8 | | |
| 9 | MAC ³ (Received) | — |
| 10 | | |

Table 18-28. AESU Context Registers (continued)

| Context Register | Cipher Mode | |
|------------------|-------------|------|
| | XCBC-MAC | CMAC |
| 11 | — | -- |
| 12 | | |

¹ Only as output, as input, K1 should be loaded into Key Data Registers 1–2

² Must be written at the start of a new message, except if zero

³ Needed only in ICV mode

— Don't care

18.8.1.11.1 Context for CBC, CBC-RBP, OFB, and CFB128 Cipher Modes

Within the context registers, for use in CBC, CBC-RBP, OFB, and CFB128 cipher modes, are two 64-bit context data registers that allow the host to read/write the contents of the initialization vector (IV):

- Context Register 1 holds the *most* significant bytes of the initialization vector (bytes 1–8).
- Context Register 2 holds the *least* significant bytes of the initialization vector (bytes 9–16).

The IV must be written prior to the message data. If the IV registers are written during message processing, or the mode is not set, a context error will be generated. The IV registers can be read only after processing has completed, as indicated by the assertion of DI (done interrupt) in the AESU status register as shown in [Section 18.8.1.6, “AESU Status Register.”](#) If the IV registers are read prior to assertion of Interrupt Done, an early read error is generated.

18.8.1.11.2 Context for Counter Cipher Mode

In counter cipher mode, a random 128-bit initial counter value is incremented modulo 2^M with each block processed. The running counter is encrypted and exclusive-ORed with the plaintext to derive the ciphertext, or with the ciphertext to recover the plaintext. The modulus exponent M can be set between 8 and 128 in multiples of 8. The value of the M is specified by writing to context register 7 as described in [Table 18-27](#).

18.8.1.11.3 Context for SRT Cipher Mode

As was noted in the AESU mode register, SRT is not a new AES cipher mode, it is an AESU method of performing AES-CTR cipher mode with reduced context loading overhead specifically for performing SRTP. It should be used with descriptor type 0010_0 “srtp”. As with counter cipher mode, a random 128-bit initial counter value is incremented modulo 2^M with each block processed. The running counter is encrypted and exclusive-ORed with the plaintext to derive the ciphertext, or with the ciphertext to recover the plaintext. The modulus exponent M can be set between 8 and 128 in multiples of 8. The value of M is specified by writing to Context Register 3 as described in [Table 18-27](#).

The only difference between SRT and CTR cipher modes is that in SRT mode, the AES context is loaded and read via context registers 1–3, with no requirement to access context registers 4–7. In CTR mode, context registers 1–4 must be loaded with zeros, with the counter and modulus being loaded into and read from context registers 5–7.

18.8.1.11.4 Context for CCM Cipher Mode

The SEC AESU is capable of performing single pass encryption and MAC generation. The host is required to order the CCM context in such a way that the context can be fetched as a contiguous string into the context registers, prior to encryption/MAC generation or decryption/MAC validation. The context register contents for CCM cipher mode is summarized in Figure 18-32 and further described below.

NOTE

AES-CCM mode does not support zero-length AAD and zero-length payload simultaneously. Either the AAD length or the payload length must be at least 1 byte.

| | | Context Registers | | | | | | |
|-----------------------|---------|-------------------|---|------------------|---|-----------------|---|--------------------------------|
| | | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| Encrypt (outbound) | Inputs | IV | | 0 | | Initial Counter | | Counter Modulus Exponent |
| | Outputs | MAC | 0 | MIC | 0 | | | |
| Decrypt (inbound) | Inputs | IV | | MIC | 0 | Initial Counter | | Counter Modulus Exponent |
| | Outputs | Computed MAC | 0 | Decrypted MAC | 0 | | | |

Figure 18-32. AESU CCM Context Registers

Context for CCM encryption/MAC generation:

- Reg 1–2 session specific 128-bit initialization vector (from memory)
- Reg 3–4 128 bits of zero padding
- Reg 5–6 Session Specific Counter (Initial Counter Value) (from memory)
- Reg 7 Counter Modulus Exponent (msb<--lsb) Should be fixed at 0x0000_0080.

NOTE

The counter modulus for CCM cipher mode is currently defined as 2^{128} making the exponent 128. This value has been made programmable in the SEC in case the final version of 802.11i uses a different counter modulus. Because this is a programmable field, it must be generated and stored along with other session specific information for loading into the AESU Context Register prior to CCM encryption.

CCM encryption processing steps:

With the session specific key and context, the AESU will perform the following operations.

1. Initialize the IV, and encrypt with the symmetric key.

2. In CBC fashion, take the output of step 1, hash with the first block of plaintext, and encrypt with the symmetric key.
3. Continue as in step 2 until the final block of plaintext has been processed. The result of the encryption of the final block of plaintext with the symmetric key is the MAC Tag. The full 128 bits of MAC data is written to context registers 1–2, for use in the next phase of CCM processing. Once the MAC Tag has been generated (step 3), the MAC tag, along with the plaintext is encrypted with the AESU operating in counter cipher mode.
4. The first item to be encrypted in counter cipher mode is the counter (initial counter value) from context registers 5–6. The counter is encrypted with the symmetric key, and the result is hashed with the MAC tag (retrieved from context registers 1–2) to produce the MIC (encrypted MAC), which is then stored in context registers 3–4. At the completion of CCM encrypt processing, this MIC is output to memory (per the descriptor pointer) for the host to append to the 802.11i frame. Note: The MIC written out to memory by the AESU is the full 128 bits. The host must only append the most significant 64 bits to the frame as the MIC.
5. The counter value is incremented, and is then encrypted with the symmetric key. The result is then hashed with the first block of plaintext to produce the first block of cipher text. The ciphertext is placed in the AESU output FIFO.
6. The counter continues to be incremented, and encrypted with the symmetric key, with the result hashed with each successive block of plaintext, until all plaintext has been converted to ciphertext. The SEC controller will manage FIFO reads and writes, fetching plaintext and writing ciphertext per the pointers provided in the descriptor. When all ciphertext and the MIC has been output, the CCM encrypt operation is complete.

Context for CCM decryption/MAC generation:

- Reg 1–2 Session specific 128 bit Initialization Vector (from memory)
- Reg 3–4 MIC (from received frame) + 64 bits of zero padding
- Reg 5–6 Session Specific Counter (Initial Counter Value) (from memory)
- Reg 7 Counter Modulus Exponent (msb<--lsb) Should be fixed at 0x0000_0080.

NOTE

The counter modulus for CCM cipher mode is currently defined as 2^{128} , making the exponent 128. This value has been made programmable in the SEC to in case the final version of 802.11i uses a different counter modulus. Because this is a programmable field, it must be generated and stored along with other session specific information for loading into the AESU Context Register prior to CCM decryption.

CCM decryption processing steps:

With the session specific key and context, the AESU will perform the following operations.

1. Initialize the IV, and encrypt with the symmetric key. Simultaneously, the counter (Initial Counter Value) from Context Registers 5-6 is encrypted with the symmetric key. The result is hashed with the Encrypted MAC (from Context Register 3-4), and the resulting Original MAC is written to Context Reg 3–4, overwriting the Encrypted MAC.

NOTE

Strictly speaking, the Counter is encrypted with the symmetric key, however the AESU should be set for “decrypt” to perform the counter and CBC processes in the correct order.

2. The 802.11 frame header is hashed with the encrypted IV. (The AESU automatically determines the header length.) Simultaneously, the counter is incremented, and is then encrypted with the symmetric key. The result is then hashed with the first block of ciphertext to produce the first block of plaintext. The plaintext is placed in the AESU output FIFO, while simultaneously, in CBC fashion, a copy of the first block of plaintext is hashed with the output of encryption of the 802.11 frame header. The output is encrypted with the symmetric key.
3. As each ciphertext block is converted to plaintext, the plaintext is CBC encrypted. When the final plaintext block has been processed, the CBC MAC (MAC Tag) is written to Context Registers 1–2. The first 64 bits of the MAC Tag are compared to the MAC Tag recovered in step 1.

NOTE

For both encrypt and decrypt operations, if the 802.11 frame is being processed as a whole (not split across multiple descriptors), the “Initialize” (AUX1) and “Final MAC” (AUX2) bits should be set in the AESU mode register.

18.8.1.11.5 Context and Operation for GCM Cipher Mode

Galois counter mode (GCM) uses AES counter mode to achieve data confidentiality. Authentication is achieved by computing a GHASH Message Authentication Code (GMAC) through performing repetitive multiplication-accumulate functions in a Galois Field.

Normally, the IV (which is provided through the input FIFO) is 96 bits. If it is 96 bits, then the IV is padded with the value $(\{0\}^{31}1)^1$. Else, the IV is hashed using the GHASH (H, {}, IV) function, where H is defined as $E(\{0\}^{128}, K)$, E stands for encryption operation, and K represents the key used. The resulting value Y_0 (the padded IV or the GHASH’ed IV) is provided as the counter input to Counter Mode AES. The result of encrypting Y_0 is called $E(Y_0, K)$, and is used to generate the final MAC tag.

Data is encrypted or decrypted by XORing input data with the pseudorandom key stream generated by counter mode AES, starting with the *second* PRK block. Initial counter value Y_0 is incremented modulo 2^{32} .

GCM cipher mode can be used to perform only the authentication part (GHASH (H, AAD, ciphertext)). To do this, set AUX0 and specify encryption operation. This special sub-mode is called GCM-GHASH in this document. The format of the context registers for GCM-GHASH mode is shown in [Table 18-33](#). GCM cipher mode also has option of automatically verifying that the received and computed MAC tags are identical. This cipher mode is called GCM w/ICV and can be specified by setting Mode Register bits 56, 57 and 62 to 1, and bit 61 to 0. GCM w/ICV context format is shown in [Table 18-33](#).

Messages (IV+AAD+textdata) are fed in through the input FIFO, and are always processed in the following order: IV, AAD, textdata, followed by the final MAC computation (where “textdata” refers to plaintext or ciphertext to be operated on). The whole message, however, does not have to be processed in

1. Notation: $\{0\}^{31}1$ is defined to mean a string of thirty-one bits of 0 followed by a single bit of 1.

one GCM execution. It can be split and processed in multiple GCM runs separated by resets of the AESU block that is in multiple descriptors. The boundaries can be set at the end of any full block (16 bytes) of the stream IV+AAD+textdata. Hence, any of the individual components (IV, AAD, or textdata) can be split into multiple descriptors. Refer to [Table 18-29](#) for proper AUX mode specification in this case and to [Table 18-30](#) through [Table 18-33](#) for proper context formatting. It should be noted that in case of a late arrival of the MAC tag on the receiving side, the final MAC can be computed and verified against the received MAC in a separate descriptor after the rest of the message (IV+AAD+textdata) has already been processed.

For host-controlled operation of the AESU in GCM cipher mode, perform the following steps:

1. Reset.
2. Set Cipher Mode to GCM or GCM w/ICV and specify encrypt, decrypt in the Mode Register. To perform GCM-GHASH (only GHASH (H, AAD, ciphertext) is computed) set AUX0 and specify encrypt. Set AUX2 and AUX1 bits according to [Table 18-29](#).
3. Load Key.
4. Load (Restore) Context as needed (see [Table 18-30](#) to [Table 18-33](#)).
5. Set Key size.
6. Set the size of the computed/received MAC (8, 12 or 16 bytes, default is 16).
7. Set Data size.
8. While available:
9. Load IV into the input FIFO (1 or multiple blocks up to 2^{64} bits in total).
10. Load AAD into the input FIFO (0 or multiple blocks up to 2^{64} bits in total).
11. Load plaintext (for encryption) or ciphertext (for decryption) blocks into the input FIFO.
12. Unload ciphertext (for encryption) or plaintext (for decryption) blocks from the output FIFO.
13. Write to the End of Message register.
14. Unload final ciphertext (for encryption) or plaintext (for decryption) blocks.
15. Read (Save) context registers if another segment of the message will be processed later.
16. Read final GCM MAC from Context Registers 3-4, if AUX2 bit was set in mode register.
17. For GCM w/ICV, check ICCR bits in the AESU status register.

Table 18-29. GCM Cipher Mode Auxiliary Bit Definitions

| Auxiliary Bit | Definitions | |
|---------------|--------------------|-------------|
| | 0 | 1 |
| AUX2 (bit 58) | Do not compute MAC | Compute MAC |

Table 18-29. GCM Cipher Mode Auxiliary Bit Definitions (continued)

| Auxiliary Bit | Definitions | |
|---------------------------|--|--|
| | 0 | 1 |
| AUX1 (bit 59) | One of the following cases: <ul style="list-style-type: none"> • Descriptor contains the whole message (IV+AAD+textdata) • Descriptor contains the whole IV and no or part of AAD or textdata • Descriptor contains a non-final part of IV, AAD, textdata (IV, AAD or textdata split between descriptors) • Descriptor contains the final part of AAD or textdata but no MAC is computed | One of the following cases: <ul style="list-style-type: none"> • Descriptor contains the final part of IV (IV split between descriptors) - len(IV)^T needed • Descriptor contains the final part of textdata and the final MAC is computed, that is $\text{AUX2}=1$ (textdata split between descriptors) - $\text{len(AAD)}^T, \text{len(textdata)}^T$ needed • Descriptor contains the whole textdata but no or part of AAD and the final MAC is computed - $\text{len(AAD)}^T, \text{len(textdata)}^T$ needed • Descriptor contains the final part of AAD and the final MAC is computed - $\text{len(AAD)}^T, \text{len(textdata)}^T$ needed • Descriptor computes only MAC (based on restored context) but does not contain either IV, AAD or textdata - $\text{len(AAD)}^T, \text{len(textdata)}^T$ needed |
| AUX0 (bit 60) and Encrypt | — | GHASH-only mode |
| AUX0 (bit 60) and Decrypt | The key is to be unrolled | The key is already unrolled |

AUX0 has different use depending on whether encryption or decryption is specified. For decryption, it determines whether the provided key should be first unrolled before processing starts, while in case of encryption it should generally be set to 0 unless GCM-GHASH cipher mode is desired. AUX2 bit determines whether the final MAC tag is to be computed or not. If AUX2 is set to 1, $E(K, Y_0)$ and the last iteration of the GHASH(H, AAD, ciphertext) is going to be performed and then XOR'ed to give the MAC tag. Hence, if the message is split into multiple descriptors, only the last one should have $\text{AUX2}=1$ for proper MAC tag computation. AUX1 is used to resolve the issues related to the splitting of messages into multiple descriptors. Table 18-29 shows the proper settings of AUX1 for several scenarios of message splitting. In general, whenever the final GHASH iteration needs to be computed (either for GHASH(H, {}, IV) or GHASH(H, AAD, ciphertext)), and the current length is not equal to total length for either IV, AAD, or textdata, AUX1 should be set to 1. Consequently, an AUX1 value of 1 also indicates that the context registers 9-10 need to provide the total length of IV, AAD, or textdata for this to be accomplished.

Table 18-30 to Table 18-33 describe the proper usage of context registers in case of encryption, decryption, GCM w/ICV and GCM-GHASH cipher mode settings. The context is in each case described in terms of the input context required for starting new GCM processing or continuation of processing after context switch, and in terms of the results (output) stored in context registers after GCM execution run is completed.

Description of context registers:

- Reg 1–2

- Initial counter value Y_0 . Normally, this value is a result of the IV stream processing and needs to be provided only if the message is split into multiple descriptors and for those descriptors that come after IV processing is complete. Otherwise, value provided here is ignored and overwritten with computed Y_0 .
- In case of GCM-GHASH cipher mode setting, constant H from GHASH(H, AAD, ciphertext) should be provided in these registers. Note that this, in the general case, does not have to be equal to $E(K, \{0\}^{128})$ where K is a key as defined for GCM.
- Reg 3–4
 - Intermediate MAC value. This needs to be provided only when switching context during AAD and/or textdata processing (AAD+textdata stream split into multiple descriptors).
 - On the output side, these registers contain either intermediate MAC tag in case of context switching (requires AUX2=0) or the final MAC tag at the end of processing if AUX2=1. If AUX2=0 on the last descriptor processing a particular message, these registers will contain partially computed GHASH(H, AAD, ciphertext) where the last GHASH iteration is not computed.
 - In case of GCM w/ICV, the final MAC tag written here as the result of GCM processing will be truncated to 8, 12, or 16 (no truncation) bytes as defined in ICV Size register. Note, that any size from 1 to 16 bytes can be specified in ICV Size register but any value other than 8 or 12 will be defaulted to 16 bytes automatically.
- Reg 5–6
 - Counter value Y_i is required only if restoring context to continue processing a message. Note that the same value read when saving context should be written to these registers when restoring context since it is automatically incremented after every processed block.
 - In case of GCM-GHASH, these registers are not used.
- Reg 7
 - Length of AAD in bits. This pertains to the length of the AAD part processed in the current descriptor. If the current descriptor does not process AAD, 0 should be written. If AAD is not split into multiple descriptors, this field should contain the total AAD length. The value written here should be divisible by 128 for all AAD segments except for the last one that can be any number of bits. Note, however, that the actual AAD stream supplied to the AES engine through FIFOs has to be zero-padded to an integral number of 16-byte blocks.
 - Reg 8
 - Length of IV in bits or a part of it processed in the current descriptor. Similar remarks apply like in the case of AAD.
 - In case of GCM-GHASH, this register is not used.
- Reg 9
 - Total length of AAD in bits. This is the total AAD length irrespective of whether AAD is split in multiple descriptors. It is required when AUX1=1 and the current descriptor processes the last segment of AAD or textdata. It is also required if the whole message is already processed and the current descriptor only computes the final MAC tag.
- Reg 10

- Total length of the plaintext/ciphertext or IV in bits. Required only when AUX1=1 (see [Table 18-29](#)). If the current descriptor processes the last segment of the IV then total IV length should be provided, otherwise total length of textdata should be provided.
- Reg 11–12
 - Received MAC tag used only in case of inbound processing with GCM with ICV. This can be a 8-, 12-, or 16-byte block as specified by writing to the ICV size register.

Table 18-30. GCM Encryption Context

| Context Register | GCM Encrypt (Outbound) | | | |
|---|---|----------------------------|----------------------|----------|
| | Mode Register (ECM = 10, AUX0 = 0, CM = 01, ED = 1) | | | |
| | AUX1 Value | | AUX2 Value | |
| | Inputs | | Outputs | |
| | AUX1 = 0 | AUX1 = 1 | | AUX2 = 0 |
| last AAD or textdata segment, or MAC only | | last IV segment | | |
| 1 | Y_0 (Initial Counter) | | | |
| 2 | | | | |
| 3 | MAC | | — | MAC |
| 4 | | | | |
| 5 | Y_i (Counter) | | | |
| 6 | | | | |
| 7 | len(AAD) ^{1,2} | | | |
| 8 | len(IV) ^{1,2} | | | |
| 9 | — | len(AAD) ³ | — | |
| 10 | — | len(textdata) ³ | len(IV) ³ | |
| 11 | — | | | |
| 12 | | | | |

¹ Length of data processed with current descriptor (in bits)

² Must be written at the start of a new message, except if zero

³ Length of total data (in bits)

— Don't care

Table 18-31. GCM Decryption Context

| Context Register | GCM Decrypt (Inbound) | | | |
|---|--|----------------------------|----------------------|-------------------|
| | Mode Register (ECM = 10, AUX0 = 0 or 1, CM = 01, ED = 0) | | | |
| | AUX1 Value | | AUX2 Value | |
| | Inputs | | Outputs | |
| | AUX1 = 0 | AUX1 = 1 | | AUX2 = 0 |
| Last AAD or textdata segment, or MAC only | | Last IV Segment | | |
| 1 | Y ₀ (Initial Counter) | | | |
| 2 | | | | |
| 3 | MAC (Computed) | | — | MAC (Computed) |
| 4 | | | | |
| 5 | Y _i (Counter) | | | |
| 6 | | | | |
| 7 | len(AAD) ^{1,2} | | | |
| 8 | len(IV) ^{1,2} | | | |
| 9 | — | len(AAD) ³ | — | |
| 10 | — | len(textdata) ³ | len(IV) ³ | |
| 11 | — | | | |
| 12 | | | | |

¹ Length of data processed with current descriptor (in bits)

² Must be written at the start of a new message, except if zero

³ Length of total data (in bits)

— Don't care

Table 18-32. GCM with ICV Context

| Context Register | GCM with ICV (Inbound) | | | |
|---|--|----------------------------|----------------------|----------|
| | Mode Register (ECM = 11, AUX0 = 0 or 1, CM = 01, ED = 0) | | | |
| | AUX1 Value | | AUX2 Value | |
| | Inputs | | Outputs | |
| | AUX1 = 0 | AUX1 = 1 | | AUX2 = 0 |
| Last AAD or textdata segment, or MAC only | | Last IV Segment | | |
| 1 | Y_0 (Initial Counter) | | | |
| 2 | | | | |
| 3 | MAC (Computed) | | | — |
| 4 | | | | |
| 5 | Y_i (Counter) | | | |
| 6 | | | | |
| 7 | len(AAD) ^{1,2} | | | |
| 8 | len(IV) ^{1,2} | | | |
| 9 | — | len(AAD) ³ | — | |
| 10 | — | len(textdata) ³ | len(IV) ³ | |
| 11 | MAC (Received) | | | |
| 12 | | | | |

¹ Length of data processed with current descriptor (in bits)

² Must be written at the start of a new message, except if zero

³ Length of total data (in bits)

— Don't care

Table 18-33. GCM-GHASH Context

| Context Register | GCM-GHASH (Only GHASH Computed) | | | |
|---|---|----------------------------|----------------------|----------------|
| | Mode Register (ECM = 10, AUX0 = 1, CM = 01, ED = 1) | | | |
| | AUX1 Value | | AUX2 Value | |
| | Inputs | | Outputs | |
| | AUX1 = 0 | AUX1 = 1 | | AUX2 = 0 |
| Last AAD or textdata segment, or MAC only | | Last IV Segment | | |
| 1 | H ¹ | | | |
| 2 | | | | |
| 3 | MAC (Computed) | | — | MAC (Computed) |
| 4 | | | | |
| 5 | — | | | |
| 6 | | | | |
| 7 | len(AAD) ^{1,2} | | | |
| 8 | -- | | | |
| 9 | — | len(AAD) ³ | — | |
| 10 | — | len(textdata) ³ | len(IV) ³ | |
| 11 | — | | | |
| 12 | | | | |

¹ Must be written at the start of a new message, except if zero

² Length of data processed with current descriptor (in bits)

³ Length of total data (in bits)

— Don't care

As an example, let's consider the case of GCM encrypt operation that generates the final MAC tag and where the whole message is small enough that it can be processed with one descriptor. The mode bits 56–63 (ECM, AUX, CM, and ED) should be set to 10_100_01_1. Only context registers 7–8 need to be written with the bit lengths of AAD and IV, respectively. The bit length of textdata should be written to the data size register up to the size of 2¹⁹ bits. IV, AAD, and textdata in that order should be sent through the input FIFO and the result (textdata) read from the output FIFO as available. At the end, the final MAC tag is read from the context registers 3–4.

18.8.1.11.6 Context and Operation for XCBC-MAC Cipher Mode

XCBC-MAC cipher mode is an authentication only mode of AES. Normal CBC-MAC runs AES in CBC cipher mode and assigns the final ciphertext result as the MAC. XCBC-MAC supports only 16-byte keys and extends the normal CBC-MAC as follows:

1. 3 keys are precomputed
 - $K1 = \text{AES-Encrypt}(K, \{01\}^{16})$.¹
 - $K2 = \text{AES-Encrypt}(K, \{02\}^{16})$.
 - $K3 = \text{AES-Encrypt}(K, \{03\}^{16})$.
2. Compute $C_{n-1} = \text{AES-CBC}(P_1, 0, K1) \dots \text{AES-CBC}(P_{n-1}, C_{n-2}, K1)$
3. If $|P_n| = \text{block size (128 bits)}$
 then: $\text{MAC} = \text{AES-CBC}(P_n \oplus K2, C_{n-1}, K1)$
 else: $\text{MAC} = \text{AES-CBC}(P_n \parallel 10^i \oplus K3, C_{n-1}, K1)$

In XCBC-MAC cipher mode, AUX0=1 means that the final data block will not be XOR'ed with K2 or K3, so that message processing can be interrupted and later continued after a context switch. AUX1=1 disables computation of keys K1, K2, and K3, and instead expects these keys to be placed in key data registers 1-2 (K1), context registers 5-6 (K2) and 7-8 (K3). If AUX1=0, computed keys will be placed into context registers 3-4, 5-6, and 7-8. AUX2=1 will enable XCBC-MAC with ICV. In XCBC-MAC with ICV, the transmitted MAC is supplied in context registers 9-10 and compared to the computed MAC in context registers 1-2.

For host-controlled operation of the AESU in XCBC-MAC cipher mode, the following steps must be performed:

1. Reset
2. Write the mode register to:
 - a) Set cipher mode to XCBC-MAC (enc/dec bit is ignored)
 - b) Set AUX0 = 1 if processing of the message is going to be interrupted and later continued after a context switch. Set AUX0 = 0 if this is the last (or only) part of the message so that the final MAC can be generated.
 - c) Set AUX1 = 1 if keys K1, K2 and K3 are loaded to Key Data Registers 1-2, Context Registers 5-6 and 7-8, respectively. Otherwise, set AUX1=0, and put K into Key Data Registers 1-2, so that keys K1, K2, and K3 can be computed and written to Context Registers 3-4, 5-6 and 7-8.
 - d) Set AUX2 = 1 if desire XCBC-MAC w/ICV.
3. Load key (K or K1 - see above)
4. Load context
5. Set key size
6. Set data size
7. While available:
8. Load data blocks
9. Write to the End of Message register
10. Read MAC from Context Registers 1-2
11. For XCBC-MAC w/ICV, check ICCR bits in the Status Register

¹.Notation: $\{01\}^{16}$ means the byte 01 repeated 16 times

18.8.1.11.7 Context and Operation for CMAC Cipher Mode

CMAC cipher mode is an authentication only mode of AES. The following is a specification of CMAC:

1. 2 keys are precomputed
 - a) $K1 = \text{xtime}(\text{AES-Encrypt}(K, \{0\}^{128}))$.¹
 - b) $K2 = \text{xtime}(K1)$.
2. Compute $C_{n-1} = \text{AES-CBC}(P_1, 0, K) \dots \text{AES-CBC}(P_{n-1}, C_{n-2}, K)$
3. If $|P_n| = \text{block size (128 bits)}$
 then: $\text{MAC} = \text{AES-CBC}(P_n \oplus K1, C_{n-1}, K)$
 else: $\text{MAC} = \text{AES-CBC}((P_n \parallel 10^i) \oplus K2, C_{n-1}, K)$

Assuming L is a 128-bit vector, L[127] is its most significant bit, and << is a bitwise shift, $\text{xtime}(L)$ is defined as:

```
If L[127] = 0, then xtime(L) = L << 1;
Else xtime(L) = (L << 1) XOR 0x87;
```

CMAC cipher mode requires transmitted MAC to be placed in Context Registers 5-6, whereas computed MAC will be in Registers 1-2 if AUX0=1. Context registers 3-4 are used to provide $E(K, \{0\}^{128})$ if AUX1=1 so that K1 and K2 can be computed after context switch without a time penalty. Computed value of $E(K, \{0\}^{128})$ will always be stored in Context Registers 3-4 to be available for saving context in case of context switching.

For host-controlled operation of the AESU in CMAC cipher mode, the following steps must be performed:

1. Reset
2. Write the mode register to:
 - a) Set Cipher Mode to CMAC (enc/dec bit is ignored)
 - b) Set AUX0 = 1 if processing of the message is going to be interrupted and later continued after a context switch. Set AUX0 = 0 if this is the last (or only) part of the message so that the final MAC can be generated.
 - c) Set AUX1 = 1 for keys K1 and K2 to be derived from $E(K, \{0\}^{128})$ that is loaded into Context Registers 3-4. Otherwise, set AUX1=0 and CMAC will compute $E(K, \{0\}^{128})$.
 - d) Set AUX2 = 1 if desire CMAC with ICV.
3. Load key
4. Load context
5. Set key size
6. Set ICV size for computed/received MAC (8, 10, 12, 14 or 16 bytes, default is 16)—ignored if AUX0 = 1
7. Set data size
8. While available, load data blocks
9. Write to the end of message register
10. Read MAC from context registers 1–2
11. For CMAC with ICV, check ICCR bits in the status register

¹.Notation: $\{0\}^{128}$ means the bit 0 repeated 128 times

18.8.1.11.8 AESU Key Registers

The AESU key registers hold from 16, 24, or 32 bytes of key data, with the first 8 bytes of key data written to key 1. Any key data written to bytes beyond the value written to the key size register will be ignored. The key data registers are cleared when the AESU is reset or re-initialized. If these registers are modified during message processing, a context error will be generated.

The key data registers may be read when changing context in decrypt mode. To resume processing, the value read must be written back to the key registers and the “restore decrypt key” bit must be set in the mode register. This eliminates the overhead of expanding the key prior to starting decryption when switching context.

18.8.1.11.9 AESU FIFOs

AESU uses an input FIFO/output FIFO pair to hold data before and after the encryption process. Normally, the channels control all access to these FIFOs. For host-controlled operation, a write to anywhere in the AESU FIFO address space enqueues data to the AESU input FIFO, and a read from anywhere in the AESU FIFO address space dequeues data from the AESU output FIFO.

Writes to the input FIFO go first to a staging register which can be written by byte, word (4 bytes), or dword (8 bytes). When all 8 bytes of the staging register have been written, the entire dword is automatically enqueued into the FIFO. If any byte is written twice between enqueues, it causes an error interrupt of type AE from the EU. When writing the last portion of data, it is not necessary to write all 8 bytes. Any last bytes remaining in the staging register are automatically padded with zeros and forced into the input FIFO when the AESU End of Message Register is written.

The output FIFO is readable by byte, word, or dword. When all 8 bytes of the head dword have been read, that dword is automatically dequeued from the FIFO so that the next dword (if any) becomes available for reading. If any byte is read twice between dequeues, it causes an error interrupt of type AE from the EU.

Overflows and underflows caused by reading or writing the AESU FIFOs are reflected in the AESU interrupt status register.

The AESU fetches data 128 bits at a time from the input FIFO. During processing, the input data is encrypted or decrypted and the results are placed in the output FIFO. The output size is the same as the input size.

The input FIFO may be written any time the number of dwords currently in the input FIFO (as indicated by the IFL field of the AESU status register) is less than 32. There is no limit on the total number of bytes in a message. The number of bits in the final message block must be set in the data size register.

The output FIFO may be read any time the OFR signal is asserted (as indicated in the AESU status register). This will indicate that the number of bytes in the output FIFO is at or above the threshold specified in the mode register.

18.8.2 Cyclical Redundancy Check Unit (CRCU)


This section contains details about the Cyclical Redundancy Check Accelerator (CRCU), including modes of operation, status and control registers, and FIFO.

Table 18-34. CRCU Mode Register Bit Definitions (continued)

| Bits | Name | Description |
|-------|------|--|
| 62 | CICV | ICV (CRC) check. Selects whether a comparison between the computed CRC and the residue should be performed. 0 Disabled 1 Enable ICV check. Compares a CRC calculated across the message and received ICV against the stored residue. The comparison is performed prior to the bit manipulations controlled by the RAW bit. |
| 59–61 | — | Reserved |
| 63 | ALG | Algorithm. Selects the CRC algorithm mode for the CRCU. 00 IEEE 802 mode. The CRC32 algorithm will be performed using the polynomial 0x04C11DB7 and the residue 0xC704DD7B will be used for ICV checking. 01 iSCSI mode. The CRC32c algorithm will be performed using the polynomial 0x1EDC6F41 and the residue 0x1C2D19ED will be used for ICV checking. 10 Static custom mode. The CRC remainder will be computed using the Control Register bits 32-63 as the polynomial and bits 0–31 as the residue. 11 Dynamic custom mode. The CRC remainder will be computed using the Key Register bits 32-63 as the polynomial and bits 0–31 as the residue. |


18.8.2.3 CRCU Key Size Register

The CRCU key size register, shown in [Figure 18-34](#), stores the number of valid bytes in the dynamic polynomial written to the key register ([Section 18.8.2.14, “CRCU Key Register”](#)). This value is reset to zero, and if a value other than zero or four is written to it, an illegal key size error is generated. It is not necessary to write to this register, as the polynomial size is clearly fixed by the algorithm. A context error is generated if this register is written after processing begins.


Figure 18-34. CRCU Key Size Register

18.8.2.4 CRCU Data Size Register

The CRCU data size register, shown in [Figure 18-35](#), is written with the number of bits of data to be processed. Writing to this register puts the CRCU module into a busy state and starts data processing. This register can be written multiple times while data processing is in progress. The actual values written are ignored, although an error is generated if the value is not a multiple of 8 bits.


Figure 18-35. CRCU Data Size Register

18.8.2.5 CRCU Reset Control Register

The reset control register, shown in [Figure 18-36](#), controls the reset/re-initialization of the block.




Figure 18-36. CRCU Reset Control Register

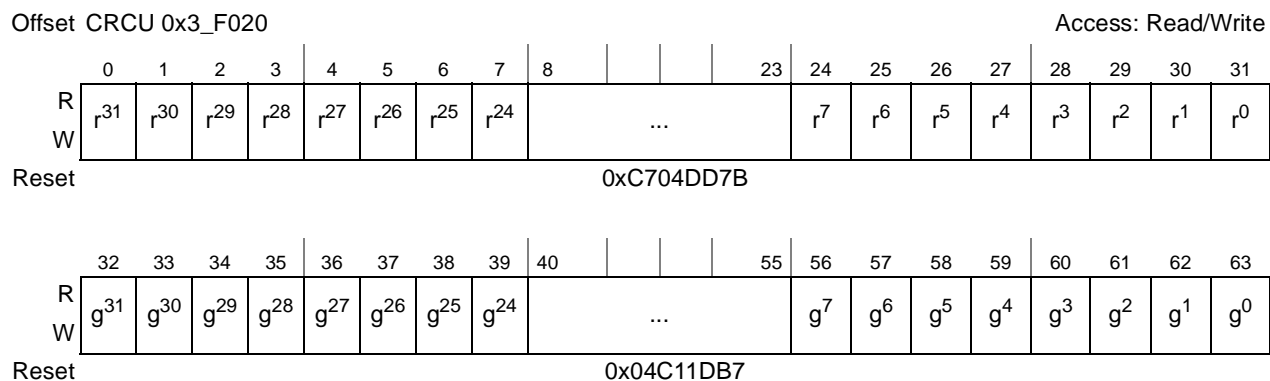
[Table 18-35](#) describes the CRCU reset control register fields.

Table 18-35. CRCU Reset Control Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–60 | — | Reserved |
| 61 | RI | Reset Interrupt. Writing this bit active high causes CRCU interrupts signalling done and error to be reset. It further resets the state of the CRCU interrupt status register. 0 Do not reset 1 Reset interrupt logic |
| 62 | MI | Module initialization is nearly the same as software reset, except that the interrupt mask register remains unchanged. 0 Do not reset 1 Reset most of CRCU NOTE: Statically programmed registers (interrupt/error mask and control) are not reset by asserting this bit. |
| 63 | SR | Software reset is functionally equivalent to hardware reset (the RESET# pin), but only for CRCU. All registers and internal state are returned to their defined reset state. On negation of SW_RESET, the CRCU will enter a routine to perform proper initialization of the S-box. 0 Do not reset 1 Full CRCU reset |

18.8.2.6 CRCU Control Register


The CRCU control register stores the static polynomial used in custom CRC computations. [Figure 18-37](#) shows the bit position of each coefficients in the polynomial.


Figure 18-37. CRCU Control Register

The reset value of this register is the IEEE Std. 802 CRC32 residue and polynomial. This register is static, in that it is only reset by performing a software reset, not by a CHA re-initialize. This allows a platform specific custom polynomial to be written to the register once and used many times. A context error is generated if this register is written after processing has begun. A polynomial error is generated if a value is written to this register which does not have a one in bit position 0.

18.8.2.7 CRCU Status Register

The CRCU status register, shown in [Figure 18-38](#), provides general information on the operational status of the CRCU. A read of the status register captures a snapshot of CRCU operating state. Three interrupt flags (irq_error, irq_done, and irq_reset) provide visibility to EU ports ipi_error_int, ipi_done_int, and ipi_reset_int, respectively. Writes to this register are ignored.


Figure 18-38. CRCU Status Register
Table 18-36. CRCU Status Register Bit Definitions

| Bits | Signal | Description |
|-------|--------|--|
| 0–48 | — | Reserved |
| 49–55 | IFL | Input FIFO level. The number of dwords currently in the input FIFO. |
| 56–57 | — | Reserved |
| 58 | HALT | Indicates when the CRCU core has halted due to an error. 0 CRCU not halted 1 CRCU core halted (Must be reset/re-initialized) Because the error causing the CRCU to stop operating may be masked to the Interrupt Status register, the status register is used to provide a second source of information regarding errors preventing normal operation. |

Table 18-37. CRCU Interrupt Status Register Bit Definitions

| Bits | Name | Description |
|-------|------|---|
| 0–48 | — | Reserved |
| 49 | ICE | Integrity check (CRC) error. 0 No error detected 1 Integrity check error detected. An ICV (CRC) check was performed and the supplied ICV (CRC) did not match the one computed by the CRCU. |
| 50 | PE | Polynomial error. 0 No error. 1 An invalid polynomial has been written to the key or control register. |
| 51 | IE | Internal error. Indicates the CRCU has been locked up and requires a reset before use. 0 No internal error detected 1 Internal error detected Note: This bit will be asserted any time an enabled error condition occurs and can only be cleared by resetting the CRCU. |
| 52 | ERE | Early read error. The CRCU context was read before the CRCU completed the requested operation. 0 No error detected 1 Early read error |
| 53 | CE | Context error. The CRCU mode, key size, control, context, or key register was modified after processing had begun. 0 No error detected 1 Context error |
| 54 | KSE | Key size error. A value other than 0 or 4 has been written to the CRCU key size register. 0 No error detected 1 Key size error |
| 55 | DSE | Data size error. An non-byte-multiple size has been written to the CRCU data size register. 0 No error detected 1 Data size error |
| 56 | ME | Mode error. Will be set if any of these error conditions is detected: <ul style="list-style-type: none"> • Any reserved bit of the mode register is set. • The ALG field of the mode register contains an illegal value. 0 No error detected 1 Mode error |
| 57 | AE | Address error. An illegal read or write address was detected within the CRCU address space. 0 No error detected 1 Address error |
| 58–60 | — | Reserved |
| 61 | IPO | Input FIFO overflow. The CRCU input FIFO was pushed while full. 0 No overflow detected 1 Input FIFO has overflowed Note: When operated through channel-controlled access, the SEC implements flow control, and FIFO size is not a limit to data input size. When operated through host-controlled access, the CRCU cannot accept FIFO inputs larger than 256 bytes without overflowing. |
| 62–63 | — | Reserved |

Table 18-38. CRCU Interrupt Mask Register Bit Definitions (continued)

| Bits | Name | Description |
|-------|------|--|
| 57 | AE | Address error. An illegal read or write address was detected within the CRCU address space. 0 Address error enabled 1 Address error disabled |
| 58–60 | — | Reserved |
| 61 | IFO | Input FIFO overflow. The CRCU input FIFO was pushed while full. 0 Input FIFO overflow error enabled 1 Input FIFO overflow error disabled |
| 62–63 | — | Reserved |

18.8.2.10 CRCU ICV Size Register

The CRCU ICV size register is word readable/writeable for compatibility with the MDEU. Values written to this location are always ignored and values read from this location always return zero. A context error is generated if this register is written after processing begins.

18.8.2.11 CRCU End of Message Register

The CRCU end of message register indicates that all data has been written to the CRCU. A write to this register is required to complete a CRC32 operation. The CRCU starts processing message data as soon as the data size register is written and data becomes available in the FIFO, but it does not process a remaining partial word or perform an ICV check until a write to this register. The value written is not used. Reading this register returns a zero value.

18.8.2.12 CRCU Received ICV Register

The CRCU received ICV register is written with the ICV, a 32-bit CRC value computed over the original data. When the CICV feature is enabled in the mode register, the value written to this register is compared with the computed CRC result available in the data out register. During this direct comparison of the two registers, the bit position of each term depends on the mode in the same way as the bits of the data out register. If the two values differ, a CICV fail error is generated. This register can be written only before an end of message is written.

18.8.2.13 CRCU Context Register

The CRCU context register can be written with an intermediate CRC result or desired initial state prior to processing any data. When processing completes, the CRC result is available from this register. The reset state of this register is all ones because this allows the CRC32 algorithm to detect bit errors in the leading zeros of a message. [Figure 18-41](#) shows the bit position of each term in the written context value. A context

error is generated if this register is written after processing begins. An early read error is generated if this register is read while the module is busy.




Figure 18-41. CRCU Context Register (Write)

If the CRCU is in the default output mode (RAW bit is 0), this register when read holds the CRC remainder after it is bit swapped, byte swapped, and complemented. This sequence of operations is described in the protocol specifications and generates a result that can be written directly to the end of a frame or command. [Figure 18-42](#) shows the value in each bit position.




Figure 18-42. CRCU Context Register (Read-Default Mode)

If the CRCU is in the raw output mode (RAW bit is 1), this register when read holds the unmodified CRC remainder. This form is the one used internally to match against the polynomial specific residue during ICV checking. [Figure 18-43](#) shows the value in each bit position.




Figure 18-43. CRCU Context Register (Read-Raw Mode)

18.8.2.14 CRCU Key Register

The CRCU key register stores the polynomial and residue for the dynamic custom mode as set in the mode register (see [Section 18.8.2.2, “CRCU Mode Register”](#)). [Figure 18-44](#) shows the bit position of each coefficients. The reset value of this register is all zeros with a one in bit position 0. This register is dynamic in that it is reset by performing a re-initialize or a software reset. This allows a custom polynomial to be used for specific processing without changing the platform-specific static custom polynomial stored in the control register (see [Section 18.8.2.6, “CRCU Control Register”](#)). A residue does not need to be programmed unless ICV checking is performed. A context error is generated if this register is written after

processing begins. A polynomial error is generated if a value is written to this register that does not have a one in bit position 0.




Figure 18-44. CRCU Key Register

18.8.2.15 CRCU FIFO

Words written to the CRCU FIFO are pushed onto the CRCU input FIFO, thereby buffering them for processing. Partial words and misaligned data can be written to this address and it is automatically realigned based on a big endian byte order.

18.8.3 Data Encryption Standard Execution Unit (DEU)

This section contains details on the data encryption standard execution unit (DEU), including modes of operation, status and control registers, and FIFOs. Most of the registers described here would not normally be accessed by the host. They are documented here mainly for debug purposes. In typical operation, the DEU is used through channel-controlled access, which means that most reads and writes of DEU registers are directed by the SEC channels. Driver software would perform host-controlled register accesses only on a few registers for initial configuration and error handling.

18.8.3.1 DEU Mode Register

The DEU mode register contains 3 bits which are used to program DEU operation. The mode register is cleared when the DEU is reset or reinitialized. Setting a reserved mode bit generates a data error. If the mode register is modified during processing, a context error is generated.

Table 18-39 describes DEU Mode Register fields.

Table 18-39. DEU Mode Register Field Descriptions

| Bits | Name | Description |
|--|------|-------------|
| The following bits are described for information only. They are not under direct user control. | | |
| 0–54 | — | Reserved |
| 55 | | Reserved |
| The following bits are controlled through the MODE0 field of the descriptor header. | | |
| 56–59 | — | Reserved |

Table 18-39. DEU Mode Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 60-61 | CM | Cipher Mode: Used to define the mode of DEU operation. 00 ECB 01 CBC 10 CFB-64 11 OFB-64 |
| 62 | TS | Triple/Single DES. If set, DEU operates the Triple DES algorithm; if not set, DEU operates the single DES algorithm. 0 Single DES 1 Triple DES |
| 63 | ED | Encrypt/decrypt. If set, DEU operates the encryption algorithm; if not set, DEU operates the decryption algorithm. 0 Perform decryption 1 Perform encryption |

18.8.3.2 DEU Key Size Register

The DEU key size register indicates the number of bytes of key memory that should be used in encrypting or decrypting. If the DEU mode register is set for single DES, any value other than 8 bytes automatically generates a key size error in the DEU interrupt status register. If the mode bit is set for triple DES, any value other than 16 bytes (112 bits for 2-key triple DES (K1=K3) or 24 bytes (168 bits for 3-key triple DES) generates an error. Triple DES always uses K1 to encrypt, K2 to decrypt, K3 to encrypt.




Figure 18-45. DEU Key Size Register

Table 18-40 shows the legal values for DEU key size.

Table 18-40. DEU Key Size Register Field Descriptions

| Bits | Name | Description |
|-------|----------|---|
| 0-51 | — | Reserved |
| 52-63 | Key Size | 8 bytes = 0x08 (only legal value if mode is single DES.) 16 bytes = 0x10 (for 2 key 3DES, K1 = K3) 24 bytes = 0x18 (for 3 key 3DES) |

18.8.3.3 DEU Data Size Register

The DEU data size register, shown in Figure 18-46, specifies the number of bits in the final message, with an upper bound of 4096. Any number written (and whatever truncated value is stored) must be a multiple of 64. All data to be processed by the DEU must be a multiple of the DES algorithm block size of 64 bits; the DEU does not automatically pad messages out to 64-bit blocks. If a data size that is not a multiple of

64 bits is written, a data size error is generated. Only bits 58–63 are checked to determine if there is a data size error. Because all upper bits are ignored, the entire message length (in bits) can be written to this register. This register is cleared when the DEU is reset or re-initialized.




Figure 18-46. DEU Data Size Register

18.8.3.4 DEU Reset Control Register

The DEU reset control register, shown in Figure 18-47, allows three levels reset of just DEU, as defined by the three self-clearing bits:




Figure 18-47. DEU Reset Control Register

Table 18-41 describes DEU reset control register fields.

Table 18-41. DEU Reset Control Register Field Descriptions

| Bits | Names | Description |
|------|-------|---|
| 0–60 | — | Reserved |
| 61 | RI | Reset interrupt. Writing this bit active high causes DEU interrupts signalling done and error to be reset. It further resets the state of the DEU interrupt status register. 0 Don't reset 1 Reset interrupt logic |
| 62 | MI | Module initialization is nearly the same as software reset, except that the interrupt mask register remains unchanged. this module initialization includes execution of an initialization routine, completion of which is indicated by the RESET_DONE bit in the DEU status register 0 Don't reset 1 Reset most of DEU |
| 63 | SR | Software reset is functionally equivalent to hardware reset (the RESET# pin), but only for DEU. All registers and internal state are returned to their defined reset state. Upon negation of SW_RESET, the DEU will enter a routine to perform proper initialization of the parameter memories. The RESET_DONE bit in the DEU status register will indicate when this initialization routine is complete 0 Don't reset 1 Full DEU reset |

18.8.3.5 DEU Status Register

The DEU status register, displayed in [Figure 18-48](#), contains fields that reflect the state of DEU internal signals. Writing to this location results in an address error reflected in the DEU interrupt status register.




Figure 18-48. DEU Status Register

[Table 18-42](#) describes the DEU status register bits.

Table 18-42. DEU Status Register

| Bits | Name | Description |
|-------|------|---|
| 0–39 | — | Reserved |
| 40–47 | OFL | The number of dwords currently in the output FIFO. |
| 48–55 | IFL | The number of dwords currently in the input FIFO. |
| 56–57 | — | Reserved |
| 58 | HALT | Halt. Indicates that the DEU has halted due to an error. 0 DEU not halted 1 DEU halted Note: Because the error causing the DEU to stop operating may be masked before reaching the interrupt status register, the DEU interrupt status register is used to provide a second source of information regarding errors preventing normal operation. |
| 59–60 | — | Reserved |
| 61 | EI | Error interrupt. This status bit reflects the state of the error interrupt signal, as sampled by the controller interrupt status register (Section 18.6.4.3, “Interrupt Status Register (ISR)”). 0 DEU is not signaling error 1 DEU is signaling error |
| 62 | DI | Done interrupt: This status bit reflects the state of the done interrupt signal, as sampled by the controller interrupt status register (Section 18.6.4.3, “Interrupt Status Register (ISR)”). 0 DEU is not signaling done 1 DEU is signaling done |
| 63 | RD | Reset done. This status bit, when high, indicates that DEU has completed its reset sequence, as reflected in the signal sampled by the appropriate channel. 0 Reset in progress 1 Reset done Note: Reset Done resets to 0, but has typically switched to 1 by the time a user checks the register, indicating the EU is ready for operation. |

18.8.3.6 DEU Interrupt Status Register

The DEU interrupt status register indicates which unmasked errors have generated error interrupts to the channel. Each bit in this register can only be set if the value of the corresponding bit in the DEU interrupt mask register is zero (see [Section 18.8.3.7, “DEU Interrupt Mask Register”](#)). If the DEU interrupt status

register is non-zero, the DEU halts and the DEU error interrupt signal is asserted to the controller (see [Section 18.6.4.3, “Interrupt Status Register \(ISR\)”](#)). In addition, if the DEU is operating through channel-controlled access, an interrupt signal is generated to the channel to which this EU is assigned. The EU error then appears in bit 55 of the channel status register (see [Table 18-13](#)) and generates a channel error interrupt to the controller.

If the interrupt status register is written from the host, ones in the value written are recorded in the interrupt status register if the corresponding bit is unmasked in the interrupt mask register. All other bits are cleared. This register can also be cleared by setting the RI bit of the DEU reset control register. The bits in the DEU interrupt status register are shown in [Figure 18-49](#).




Figure 18-49. DEU Interrupt Status Register

[Table 18-43](#) describes the DEU interrupt status register fields.

Table 18-43. DEU Interrupt Status Register Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0-49 | — | Reserved |
| 50 | KPE | Key parity error. Defined parity bits in the keys written to the key registers did not reflect odd parity correctly. (Note that key register 2 and key register 3 are checked for parity only if the appropriate DEU mode register bit indicates triple DES. Also, key register 3 is checked only if key size reg = 24. Key register 2 is checked only if key size reg = 16 or 24.) 0 No error detected 1 Key parity error |
| 51 | IE | Internal error. An internal processing error was detected while performing encryption. 0 No error detected 1 Internal error Note: This bit will be asserted any time an enabled error condition occurs and can only be cleared by setting the corresponding bit in the Interrupt Mask Register or by resetting the DEU. |
| 52 | ERE | Early read error. The DEU IV register was read while the DEU was performing encryption. 0 No error detected 1 Early read error |
| 53 | CE | Context error. A DEU key register, the key size register, data size register, mode register, or IV register was modified while DEU was performing encryption. 0 No error detected 1 Context error |
| 54 | KSE | Key size error. An inappropriate value (8 being appropriate for single DES, and 16 and 24 being appropriate for triple DES) was written to the DEU key size register 0 No error detected 1 Key size error |
| 55 | DSE | Data size error. A value was written to the DEU data size register that is not a multiple of 64 bits. 0 No error detected 1 Data size error |

Table 18-43. DEU Interrupt Status Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 56 | ME | Mode error. An illegal value was detected in the mode register. Note: writing to reserved bits in mode register is likely source of error. 0 No error detected 1 Mode error |
| 57 | AE | Address error. An illegal read or write address was detected within the DEU address space. 0 No error detected 1 Address error |
| 58 | OFE | Output FIFO error. The DEU output FIFO was detected non-empty upon write of DEU data size register. 0 No error detected 1 Output FIFO non-empty error |
| 59 | IFE | Input FIFO error. The DEU input FIFO was detected non-empty upon generation of done interrupt. 0 No error detected 1 Input FIFO non-empty error |
| 60 | IFU | Input FIFO underflow. The DEU input FIFO was read while empty. 0 No error detected 1 Input FIFO has had underflow error |
| 61 | IFO | Input FIFO Overflow. The DEU input FIFO was pushed while full. 0 No error detected 1 Input FIFO has overflowed Note: When operated through channel-controlled access, the SEC implements flow control, and FIFO size is not a limit to data input. When operated through host-controlled access, the DEU cannot accept FIFO inputs larger than 256 bytes without overflowing. |
| 62 | OFU | Output FIFO underflow. The DEU output FIFO was read while empty. 0 No error detected 1 Output FIFO has underflow error |
| 63 | OFO | Output FIFO overflow. The DEU output FIFO was pushed while full. 0 No error detected 1 Output FIFO has overflowed |

18.8.3.7 DEU Interrupt Mask Register

The DEU interrupt mask register controls the result of detected errors. For a given error (as defined in [Section 18.8.3.6, “DEU Interrupt Status Register”](#)), if the corresponding bit in this register is set, the error is ignored; no bit is set in the DEU interrupt status register, and no error interrupt occurs. If the corresponding bit is not set and an error is detected, the interrupt status register is updated to reflect the error, causing assertion of the error interrupt signal and causing the module to halt processing.




Figure 18-50. DEU Interrupt Mask Register

Table 18-44. DEU Interrupt Mask Register Field Descriptions


| Bits | Name | Description |
|------|------|---|
| 0–49 | — | Reserved |
| 50 | KPE | Key parity error. The defined parity bits in the keys written to the key registers did not reflect odd parity correctly. (Note that key register 2 and key register 3 are only checked for parity if the appropriate DEU mode register bit indicates triple DES. 0 Key parity error enabled 1 Key parity error disabled |
| 51 | IE | Internal error. An internal processing error was detected while performing encryption. 0 Internal error enabled 1 Internal error disabled |
| 52 | ERE | Early read error. The DEU IV register was read while the DEU was performing encryption. 0 Early read error enabled 1 Early read error disabled |
| 53 | CE | Context error. A DEU key register, the key size register, the data size register, the mode register, or IV register was modified while DEU was performing encryption. 0 Context error enabled 1 Context error disabled |
| 54 | KSE | Key size error. An inappropriate value (8 being appropriate for single DES, and 16 and 24 being appropriate for Triple DES) was written to the DEU key size register. 0 Key size error enabled 1 Key size error disabled |
| 55 | DSE | Data size error. A value that is not a multiple of 64 bits was written to the DEU data size register. 0 Data size error enabled 1 Data size error disabled |
| 56 | ME | Mode error. An illegal value was detected in the mode register. 0 Mode error enabled 1 Mode error disabled |
| 57 | AE | Address error. An illegal read or write address was detected within the DEU address space. 0 Address error enabled 1 Address error disabled |
| 58 | OFE | Output FIFO error. The DEU output FIFO was detected non-empty upon write of DEU data size register 0 Output FIFO non-empty error enabled 1 Output FIFO non-empty error disabled |
| 59 | IFE | Input FIFO error. The DEU input FIFO was detected non-empty upon generation of done interrupt 0 Input FIFO non-empty error enabled 1 Input FIFO non-empty error disabled |
| 60 | IFU | Input FIFO underflow. The DEU input FIFO was read while empty. 0 Input FIFO underflow error enabled 1 Input FIFO underflow error disabled |

Table 18-44. DEU Interrupt Mask Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 61 | IFO | Input FIFO overflow. The DEU input FIFO was pushed while full. 0 Input FIFO overflow error enabled 1 Input FIFO overflow error disabled Note: When operated through channel-controlled access, the SEC implements flow control, and FIFO size is not a limit to data input. When operated through host-controlled access, the DEU cannot accept FIFO inputs larger than 256 bytes without overflowing. |
| 62 | OFU | Output FIFO underflow. The DEU output FIFO was read while empty. 0 Output FIFO underflow error enabled 1 Output FIFO underflow error disabled |
| 63 | OFO | Output FIFO overflow. The DEU output FIFO was pushed while full. 0 Output FIFO Overflow error enabled 1 Output FIFO Overflow error disabled |

18.8.3.8 DEU End_of_message Register

Writing the DEU end-of-message register is a handshake mechanism signalling that no more data is to be written to the input FIFO. DESA signals a done interrupt after the input FIFO is detected empty any time following the write to end-of-message. After the final message block is written to the input FIFO, the DEU end_of_message register must be written. The value in the data size register is used to determine the number of bits of the final message block (always 64) to be processed. Note that the end_of_ register has no data size, and during the write operation, the host data bus is not read. Hence, any data value is accepted. Normally, a write operation with a zero data value is performed. Reading from this register is not meaningful, but a zero value is always returned and no error is generated. Writing to this register is merely a trigger causing the DEU to process the final block of a message, allowing it to signal done interrupt.


Figure 18-51. DEU End_of_Message Register

18.8.3.9 DEU IV Register

For CBC mode, the initialization vector is written to and read from the DEU IV register. The value of this register changes as a result of the encryption process and reflects the context of DEU. Reading this memory location while the module is processing data generates an error interrupt.

18.8.3.10 DEU Key Registers

The DEU uses three write-only key registers, K1, K2, and K3, to perform encryption and decryption. In Single DES mode, only K1 may be written. The value written to K1 is simultaneously written to K3, auto-enabling the DEU for 112-bit Triple DES if the key size register indicates 2 key 3DES is to be

performed (key size = 16 bytes). To operate in 168-bit Triple DES, K1 must be written first, followed by the write of K2, then K3. Reading any of these memory locations generates an address error interrupt.

18.8.3.11 DEU FIFOs

DEU uses an input FIFO/output FIFO pair to hold data before and after the encryption process. Normally, the channels control all access to these FIFOs. For host-controlled operation, a write to anywhere in the DEU FIFO address space enqueues data to the DEU input FIFO, and a read from anywhere in the DEU FIFO address space dequeues data from the DEU output FIFO.

Writes to the input FIFO go first to a staging register which can be written by byte, word (4 bytes), or dword (8 bytes). When all 8 bytes of the staging register have been written, the entire dword is automatically enqueued into the FIFO. If any byte is written twice between enqueues, it causes an error interrupt of type AE from the EU. Since the DEU data length should always be a multiple of 8 bytes, the last write should complete a dword. However, if there is any partial dword in the staging register when the DEU End of Message Register is written, the partial dword is automatically padded with zeros to a full 8 bytes and enqueued to the input FIFO.

The output FIFO is readable by byte, word, or dword. When all 8 bytes of the head dword have been read, that dword is automatically dequeued from the FIFO so that the next dword (if any) becomes available for reading. If any byte is read twice between dequeues, it causes an error interrupt of type AE from the EU.

Overflows and underflows caused by reading or writing the DEU FIFOs are reflected in the DEU interrupt status register.

18.8.4 Message Digest Execution Unit (MDEU)

This section contains details about the Message Digest Execution Unit (MDEU), including modes of operation, status and control registers, and FIFO.

Most of the registers described here would not normally be accessed by the host. They are documented here mainly for debug purposes. In typical operation, the MDEU is used through channel-controlled access, which means that most reads and writes of MDEU registers are directed by the SEC channels. Driver software would perform host-controlled register accesses only on a few registers for initial configuration and error handling.

18.8.4.1 ICV Checking

This EU includes an ICV checking feature, that is, it can generate an ICV and compare it to another supplied ICV. The pass/fail result of this ICV check can be returned to the host either via interrupt by a writeback of EU status fields into host memory, but not by both methods at once.

To signal the ICV checking result by status writeback, turn on either the IWSE bit or AWSE bit in the Channel Configuration Register (see [Section 18.5.4.1, “Channel Configuration Register \(CCR\)”](#)), and mask the ICE bit in the Interrupt Mask Register ([Section 18.8.4.9, “MDEU Interrupt Mask Register”](#)). In this case the normal done signalling (by interrupt or writeback) is undisturbed.

To signal the ICV checking result by interrupt, unmask the ICE bit in the Interrupt Mask Register and turn off the IWSE and AWSE bits in the channel configuration register. If there is no ICV mismatch, then the

normal done signalling (by interrupt or writeback) will occur. When there is an ICV mismatch, there will be an error interrupt to the host, but no channel done interrupt or writeback.

18.8.4.2 MDEU Mode Register

The MDEU Mode Register is used to program the function of the MDEU. For normal operation through a channel, bits 56–63 of this register are specified by the user through the MODE0 or MODE1 field of the descriptor header. The remaining two bits are supplied by the channel and thus are not under direct user control.

The two bits supplied by the channel are bits that control the meanings of other mode register fields. They are the MDEU_B bit, and the NEW bit.

The MDEU_B bit determines which of two sets of algorithms will be available through the ALG bits. The two sets of algorithms are referred to as the MDEU-A set (MD5, SHA-1, SHA-224, and SHA-256) and the MDEU-B set (SHA-224, SHA-256, SHA-384, and SHA-512). MDEU_B = 0 selects the MDEU-A set, and MDEU_B = 1 selects the MDEU-B set. In channel-controlled operation, the MDEU_B mode bit is supplied by the channel, based on the EU_SEL field of the descriptor header, where the user can choose MDEU-A or MDEU-B (see Table 18-6).

The NEW bit determines the configuration of other mode register fields as shown in Figure 18-52 and Figure 18-53. The “new” configuration (NEW=1) is used only by TLS/SSL descriptor types (1000_1, 1001_1). The old configuration (NEW=0) is used by all other descriptor types. The old configuration is the same as the one used in SEC 2.0, except for the CICV and SMAC bits. When MDEU is configured by the Polychannel, the value of NEW is determined by the descriptor type field of the descriptor header.

The mode register is cleared when the MDEU is reset or re-initialized. Setting a reserved mode bit will generate a data error. If the mode register is modified during processing, a context error will be generated.




Figure 18-52. MDEU Mode Register in Old Configuration (NEW = 0)

Table 18-45 describes MDEU mode register fields in old configuration.

Table 18-45. MDEU Mode Register in Old Configuration

| Bits | Name | Description |
|--|--------|--|
| The following bits are described for information only. They are not under direct user control. | | |
| 0–50 | — | Reserved |
| 51 | MDEU_B | Selects which algorithms are enabled by the ALG bits. 0 MDEU-A enables selection between SHA-1, SHA-256, MD5, and SHA-224 1 MDEU-B enables selection between SHA-384, SHA-256, SHA-512, and SHA-224. |
| 52–53 | — | Reserved, must be set to zero |

Table 18-45. MDEU Mode Register in Old Configuration (continued)

| Bits | Name | Description |
|---|----------|--|
| 54 | NEW (=0) | Determines the configuration of the MDEU Mode Register. This table shows the configuration for NEW=0. |
| 55 | — | Reserved, must be set to zero |
| The following bits are controlled through the MODE0 or MODE1 fields of the descriptor header. | | |
| 56 | CONT | Continue: Most operations will require this bit to be cleared. It is set only when the data to be hashed is spread across multiple descriptors. The value programmed in PD must be opposite to the value in this bit. 0 Do autopadding and complete the message digest. Used when the entire hash is performed with one descriptor, or on the last of a sequence of descriptors. 1 This hash will be continued in a subsequent descriptor. Do not autopad and do not complete the message digest. |
| 57 | CICV | Compare integrity check values. 0 Normal operation; no ICV checking. 1 After the message digest (ICV) is computed, compare it to the data in the MDEU's input FIFO. If the ICVs do not match, send an error interrupt to the channel. The number of bytes to be compared is given by the ICV Size Register. Only applicable to descriptor types that provide for reading an ICV in value. |
| 58 | SMAC | Specifies whether to perform an SSL-MAC operation. 0 Normal operation 1 Perform an SSL3.0 MAC operation. This requires a key and key length. If this is set then the HMAC bit should be 0. |
| 59 | INIT | Initialization bit. Most operations will require this bit to be set. Cleared only for operations that load context from a known intermediate hash value. 0 Do not initialize digest registers. In this case the registers must be loaded from a hash context pointer in the descriptor. When the data to be hashed is spread across multiple descriptors, this bit must be 0 on all but the first descriptor. 1 Do an algorithm-specific initialization of the digest registers. |
| 60 | HMAC | Specifies whether to perform an HMAC operation: 0 Normal operation 1 Perform an HMAC operation. This requires a key and key length. If this is set then the SMAC bit should be 0. |
| 61 | PD | This bit must be programmed opposite to the CONT bit. |
| 62–63 | ALG | Message digest algorithm selection 00 if MDHA-B, then SHA-384. If MDHA-A, then SHA-160 algorithm (full name for SHA-1) 01 SHA-256 algorithm 10 if MDHA-B, then SHA-512. If MDHA-A, then MD5 algorithm 11 SHA-224 algorithm |



Figure 18-53. MDEU Mode Register in New Configuration (NEW = 1)

Table 18-46 describes MDEU Mode Register fields in new configuration.

Table 18-46. MDEU Mode Register in New Configuration

| Bits | Name | Description |
|--|----------|---|
| The following bits are described for information only. They are not under direct user control. | | |
| 0–50 | — | Reserved |
| 51 | MDEU_B | Selects which algorithms are enabled by the ALG bits. 0 MDEU-A enables selection between SHA-1, SHA-256, MD5, and SHA-224 1 MDEU-B enables selection between SHA-384, SHA-256, SHA-512, and SHA-224. |
| 52 | — | Reserved, must be set to zero |
| 53 | STIB | SSL/TLS inbound, block cipher. 0 Normal operation. 1 Special operation only for SSL/TLS inbound, block cipher. Upon receiving End_of_message, the MDEU performs a calculation involving the last valid byte of data written into its input FIFO (which is Pad Length) to compute a final data size. The MDEU then processes the amount of data specified by this data size, and completes the message digest. |
| 54 | NEW (=1) | Determines the configuration of the MDEU mode register. This table shows the configuration for NEW = 1. |
| 55 | — | Reserved, must be set to zero. |
| The following bits are controlled through the MODE0 or MODE1 fields of the descriptor header. | | |
| 56 | CONT | Continue. Most operations will require this bit to be cleared. Set only when the data to be hashed is spread across multiple descriptors. 0 Do autopadding and complete the message digest. Used when the entire hash is performed with one descriptor, or on the last of a sequence of descriptors. 1 This hash will be continued in a subsequent descriptor. Do not autopad and do not complete the message digest. |
| 57 | CICV | Compare integrity check values. 0 Normal operation; no ICV checking. 1 After the message digest (ICV) is computed, compare it to the data in the MDEU's input FIFO. If the ICVs do not match, send an error interrupt to the channel. The number of bytes to be compared is given by the ICV Size Register. |
| 58 | SMAC | Specifies whether to perform an SSL-MAC operation. 0 Normal operation 1 Perform an SSL3.0 MAC operation. This requires a key and key length. If this is set then the HMAC bit should be 0. |
| 59 | INIT | Initialization bit. Most operations will require this bit to be set. Cleared only for operations that load context from a known intermediate hash value. 0 Do not initialize digest registers. In this case the registers must be loaded from a hash context pointer in the descriptor. When the data to be hashed is spread across multiple descriptors, this bit is set on all but the first descriptor. 1 Do an algorithm-specific initialization of the digest registers. |
| 60 | HMAC | Specifies whether to perform an HMAC operation. 0 Normal operation 1 Perform an HMAC operation. This requires a key and key length. If this is set then the SMAC bit should be 0. |

Table 18-46. MDEU Mode Register in New Configuration (continued)

| Bits | Name | Description |
|-------|------|--|
| 61 | EALG | The EALG (extended algorithm bit) and ALG (algorithm) bits together specify the message digest algorithm, as follows: |
| 62–63 | ALG | 000 if MDHA-B, then SHA-384. If MDHA-A, then SHA-160 algorithm (full name for SHA-1) 001 SHA-256 algorithm 010 if MDHA-B, then SHA-512. If MDHA-A, then MD5 algorithm 011 SHA-224 algorithm Others: Reserved |

18.8.4.3 Recommended Settings for MDEU Mode Register

The most common task likely to be executed via the MDEU is HMAC generation. HMACs are used to provide message integrity within a number of security protocols, including IPsec, and TLS. The SSL 3.0 protocol uses a slightly different SSL-MAC. If an HMAC or SSL-MAC is to be performed using a single descriptor (with the MDEU acting as sole or secondary EU), the following mode register bit settings should be used:

Table 18-47. Mode Register—HMAC or SSL-MAC Generated by Single Descriptor

| Bits | Field | Value | |
|------|-------|----------|-------------|
| | | for HMAC | for SSL-MAC |
| 56 | CONT | 0 (off) | 0 (off) |
| 58 | SMAC | 0 (on) | 1 (on) |
| 59 | INIT | 1 (on) | 1 (on) |
| 60 | HMAC | 1 (on) | 0 (on) |

To generate an HMAC for a message that is spread across a sequence of descriptors, the following mode register bit settings should be used:

Table 18-48. Mode Register—HMAC Generated across a Sequence of Descriptors

| Bits | Field | Value | | |
|------|-------|------------------|----------------------|------------------|
| | | First Descriptor | Middle Descriptor(s) | Final Descriptor |
| 56 | CONT | 1 (on) | 1 (on) | 0 (off) |
| 59 | INIT | 1 (on) | 0 (off) | 0 (off) |
| 60 | HMAC | 1 (on) | 0 (off) | 1 (on) |

All descriptors other than the final descriptor must output the intermediate message digest for the next descriptor to reload as MDEU context. SSL-MAC operations cannot be spread across a sequence of descriptors. For additional information on descriptors, see [Section 18.3, “Descriptors.”](#)

18.8.4.4 MDEU Key Size Register

The MDEU key size register, shown in Figure 18-54, indicates the number of bytes of key memory that should be used in HMAC generation. MDEU supports at most one block of key. MDEU generates a key size error if the value written to this register exceeds 64 bytes for MD5, SHA-1, SHA-224, or SHA-256. If algorithms SHA-384 or SHA-512 are selected, MDEU generates a key size error if the value written to this register exceeds 128 bytes.




Figure 18-54. MDEU Key Size Register

18.8.4.5 MDEU Data Size Register

The MDEU data size register, shown in Figure 18-55, specifies the number of bits of data to be processed. The Data Size field is a 21-bit signed number. Values written to this register are added to the current register value. Multiple writes are allowed. The MDEU processes data when there is a positive value in this register and there is data available in the MDEU input FIFO. (Negative values can arise in inbound processing, when it is necessary to hold back data from the MDEU until the pad length is decrypted.)

Because the MDEU does not support bit offsets, bits 61–63 must be written as 0 and are always read as zero. Furthermore, when the CONT bit of the MDEU mode register is set, the data size must be a multiple of the block size (512 bits for MD5, SHA-1, SHA-224 and SHA-256; 1024 bits for SHA-384 and SHA-512). Violating either of these conditions causes a data size error (DSE in the MDEU interrupt status register). This register is cleared when the MDEU is reset or re-initialized. At the end of processing, its contents are decremented to zero (unless there is an error interrupt).

NOTE

Writing to the data size register allows the MDEU to enter auto-start mode. Therefore, the required context registers must be written prior to writing the data size.




Figure 18-55. MDEU Data Size Register

18.8.4.6 MDEU Reset Control Register

The MDEU reset control register, shown in Figure 18-56, allows three levels of reset of just the MDEU, as defined by the three self-clearing bits.

Table 18-50 describes the MDEU status register fields.

Table 18-50. MDEU Status Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–57 | — | Reserved |
| 58 | HALT | Halt. Indicates that the MDEU has halted due to an error. 0 MDEU not halted 1 MDEU halted Note: Because the error causing the MDEU to stop operating may be masked before reaching the interrupt status register, the MDEU interrupt status register is used to provide a second source of information regarding errors preventing normal operation. |
| 59–60 | ICCR | Integrity check comparison result. 00 No integrity check comparison was performed. 01 The integrity check comparison passed. 10 The integrity check comparison failed. 11 Reserved Note: A passed or failed result is generated only if ICV checking is enabled. |
| 61 | EI | Error interrupt. This status bit reflects the state of the error interrupt signal, as sampled by the controller interrupt status register (Section 18.6.4.3, “Interrupt Status Register (ISR)”). 0 MDEU is not signaling error 1 MDEU is signaling error |
| 62 | DI | Done interrupt. This status bit reflects the state of the done interrupt signal, as sampled by the controller interrupt status register (Section 18.6.4.3, “Interrupt Status Register (ISR)”). 0 MDEU is not signaling done 1 MDEU is signaling done |
| 63 | RD | Reset done. This status bit, when high, indicates that MDEU has completed its reset sequence, as reflected in the signal sampled by the appropriate channel. 0 Reset in progress 1 Reset done Note: Reset Done resets to 0, but has typically switched to 1 by the time a user checks the register, indicating the EU is ready for operation. |

18.8.4.8 MDEU Interrupt Status Register

The MDEU interrupt status register indicates which unmasked errors have generated error interrupts to the channel. Each bit in this register can be set only if the corresponding bit of the MDEU interrupt mask register is zero (see Section 18.8.4.9, “MDEU Interrupt Mask Register”). If the value in the MDEU interrupt status register is non-zero, the MDEU halts and the MDEU error interrupt signal is asserted to the controller (see Section 18.6.4.3, “Interrupt Status Register (ISR)”). In addition, if the MDEU is operating through channel-controlled access, an interrupt signal is generated to the channel to which this EU is assigned. The EU error then appears in bit 55 of the channel status register (see Table 18-13) and generates a channel error interrupt to the controller. If the interrupt status register is written from the host, ones in the value written are recorded in the interrupt status register if the corresponding bit is unmasked in the interrupt mask register. All other bits are cleared. This register can also be cleared by setting the RI bit of the MDEU reset control register.

Table 18-51. MDEU Interrupt Status Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 57 | AE | Address error. An illegal read or write address was detected within the MDEU address space. 0 No error detected 1 Address error |
| 58–60 | — | Reserved |
| 61 | IFO | Input FIFO Overflow. the MDEU Input FIFO was pushed while full. 0 No overflow detected 1 Input FIFO has overflowed Note: When operated through channel-controlled access, the SEC implements flow control, and FIFO size is not a limit to data input size. When operated through host-controlled access, the MDEU cannot accept FIFO inputs larger than 256 bytes without overflowing. |
| 62–63 | — | Reserved |

18.8.4.9 MDEU Interrupt Mask Register

The MDEU interrupt mask register, shown in [Figure 18-59](#), controls the result of detected errors. For a given error (as defined in [Section 18.8.4.8, “MDEU Interrupt Status Register”](#)), if the corresponding bit in this register is set, the error is disabled; no error interrupt occurs and the interrupt status register is not updated to reflect the error. If the corresponding bit is not set and an error is detected, the interrupt status register is updated to reflect the error, causing assertion of the error interrupt signal and causing the module to halt processing.




Figure 18-59. MDEU Interrupt Mask Register

[Table 18-52](#) describes MDEU interrupt status register fields.

Table 18-52. MDEU Interrupt Mask Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–48 | — | Reserved |
| 49 | ICE | Integrity check error. The supplied ICV did not match the one computed by the MDEU. 0 Integrity check error enabled. WARNING: Do not enable this if using EU status writeback (see bits IWSE and AWSE in Section 18.5.4.1, “Channel Configuration Register (CCR)”). 1 Integrity check error disabled |
| 50 | — | Reserved |
| 51 | IE | Internal error. An internal processing error was detected while performing hashing. 0 Internal error enabled 1 Internal error disabled |

block (always 512) are to be processed. Note that this register has no data size, and during the write operation, the host data bus is not read. Hence, any data value is accepted. Normally, a write operation with a zero data value is performed. Reading from this register is not meaningful, but a zero value is always returned, and no error is generated. Writing to this register merely triggers the MDEU to process the final block of a message, allowing it to signal a done interrupt.




Figure 18-61. MDEU End_of_message Register

18.8.4.12 MDEU Context Registers

For MDEU, context consists of the hash plus the message length count. Write access to this register block allows continuation of a previous hash. Reading these registers provide the resulting message digest or HMAC, along with an aggregate bit count.

NOTE

All SHA algorithms are big-endian. MD5 is little endian. The MDEU module internally reverses the endianness of the five registers A, B, C, D, and E upon writing to or reading from the MDEU context if the MDEU mode register indicates MD5 is the hash of choice. Most other endian considerations are performed as 8-byte swaps. In this case, 4-byte endianness swapping is performed within the A, B, C, D, and E fields as individual registers. Reading this memory location while the module is not done generates an error interrupt.

After a power-on reset, all the MDEU context register values are cleared. [Figure 18-62](#) shows how the MDEU context registers are initialized if the INIT bit is set in the MDEU mode register. All registers are initialized, regardless of mode selected. However, only the appropriate context register values are used in hash generation per the mode selected. The user typically does not care about the MDEU Context Register initialization values; they are documented for completeness in the event the user reads these registers using host-controlled access. MDEU reset via the MDEU reset control register ([Figure 18-56](#)) or SEC global software reset ([Figure 18-22](#)) does not clear these registers.

| | 0 | 31 | 32 | 63 | |
|------------------|------------------------|----|----------------|----|-------------------------------|
| Algorithm | | | | | Context offset 0x3_6100 |
| MD-5 | A = 0x01234567 | | B = 0x89ABCDEF | | |
| SHA-1 | A = 0x67452301 | | B = 0xEFCDAB89 | | |
| SHA-224 | A = 0xC1059ED8 | | B = 0x367CD507 | | |
| SHA-256 | A = 0x6A09E667 | | B = 0xBB67AE85 | | |
| SHA-384 | A = 0xCBBB9D5DC1059ED8 | | | | |
| SHA-512 | A = 0x6A09E667F3BCC908 | | | | |
| Algorithm | | | | | Context offset 0x3_6108 |
| MD-5 | C = 0xFEDCBA98 | | D = 0x76543210 | | |
| SHA-1 | C = 0x98BADCFE | | D = 0x10325476 | | |
| SHA-224 | C = 0x3070DD17 | | D = 0xF70E5939 | | |
| SHA-256 | C = 0x3C6EF372 | | D = 0xA54FF53A | | |
| SHA-384 | B = 0x629A292A367CD507 | | | | |
| SHA-512 | B = 0xBB67AE8584CAA73B | | | | |
| Algorithm | | | | | Context offset 0x3_6110 |
| MD-5 | E = 0xF0E1D2C3 | | F = 0x8C68059B | | |
| SHA-1 | E = 0xC3D2E1F0 | | F = 0x9B05688C | | |
| SHA-224 | E = 0xFFC00B31 | | F = 0x68581511 | | |
| SHA-256 | E = 0x510E527F | | F = 0x9B05688C | | |
| SHA-384 | C = 0x9159015A3070DD17 | | | | |
| SHA-512 | C = 0x3C6EF372FE94F82B | | | | |

Figure 18-62. MDEU Context Register

| | | |
|-----------------------------|--------------------------|----------------|
| Algorithm | | Context offset |
| MD-5 | G = 0xABD9831F | 0x3_6118 |
| | H = 0x19CDE05B | |
| SHA-1 | G = 0x1F83D9AB | |
| | H = 0x5BE0CD19 | |
| SHA-224 | G = 0x64F98FA7 | |
| | H = 0xBEFA4FA4 | |
| SHA-256 | G = 0x1F83D9AB | |
| | H = 0x5BE0CD19 | |
| SHA-384 | D = 0xh152FEC88HF70E5939 | |
| SHA-512 | D = 0xhA54FF53AH5F1D36F1 | |
| Algorithm | | Context offset |
| MD5, SHA1, SHA-224, SHA-256 | Message Length Count = 0 | 0x3_6120 |
| SHA-384 | E = 0x67332667FFC00B31 | |
| SHA-512 | E = 0x510E527FADE682D1 | |
| Algorithm | | Context offset |
| MD5, SHA1, SHA-224, SHA-256 | Reserved | 0x3_6128 |
| SHA-384 | F = 0x8EB44A8768581511 | |
| SHA-512 | F = 0x9B05688C2B3E6C1F | |
| Algorithm | | Context offset |
| MD5, SHA1, SHA-224, SHA-256 | Reserved | 0x3_6130 |
| SHA-384 | G = 0xDB0C2E0D64F98FA7 | |
| SHA-512 | G = 0x1F83D9ABFB41BD6B | |
| Algorithm | | Context offset |
| MD5, SHA1, SHA-224, SHA-256 | Reserved | 0x3_6138 |
| SHA-384 | H = 0x47B5481DBEFA4FA4 | |
| SHA-512 | H = 0x5BE0CD19137E2179 | |
| Algorithm | | Context offset |
| MD5, SHA1, SHA-224, SHA-256 | Reserved | 0x3_6140 |
| SHA-384, SHA-512 | Message Length Count = 0 | |

Figure 18-62. MDEU Context Register

If SHA-384 or SHA-512 are selected, then each of the registers A, B, C, D, E, F, G, H are 64-bits (instead of 32 bits for other hash algorithms). As a result, the base address for each context register is shifted to adjust.

18.8.4.13 MDEU Key Registers

The MDEU maintains sixteen 64-bit registers for writing an HMAC key. Only the first eight are used for MD5, SHA-1, SHA-224, or SHA-256. The IPAD and OPAD operations are performed automatically on the key data when required.

NOTE

All SHA algorithms are big endian. MD5 is little endian. The MDEU module internally reverses the endianness of the key upon writing to or reading from the MDEU key registers if the MDEU mode register indicates MD5 is the hash of choice.

18.8.4.14 MDEU FIFOs

MDEU uses an input FIFO to hold data to be hashed (followed in some cases by an ICV value for ICV checking). Normally, the channels control all access to this FIFO. For host-controlled operation, a write to anywhere in the MDEU FIFO address space enqueues data to the MDEU input FIFO, and a read from anywhere in this address space returns all zeros.

When the host writes to the MDEU FIFO (using host-controlled access), it can write to any FIFO address by byte, word (4 bytes), or dword (8 bytes). The MDEU assembles these bytes from left to right, that is, the first bytes written are placed in the most significant bit-positions. Whenever the MDEU accumulates 8 bytes, this dword is automatically enqueued into the FIFO, and any remaining bytes are left-justified in preparation for assembling the next dword. It is not necessary to fill all bytes of the final dword. Any last bytes remaining in the staging register are automatically padded with zeros and forced into the input FIFO when the MDEU end of message register is written. Overflows caused by writing the MDEU FIFO are reflected in the MDEU interrupt status register.

18.8.5 Public Key Execution Units (PKEU)

This section contains details on the public key execution unit (PKEU), including modes of operation, status and control registers, and parameter RAMs. Most of the registers described here would not normally be accessed by the host. They are documented here mainly for debug purposes. In typical operation, the PKEU is used through channel-controlled access, which means that most reads and writes of PKEU registers are directed by the SEC channels. Driver software would perform host-controlled register accesses only on a few registers for initial configuration and error handling.

18.8.5.1 PKEU Mode Register

This register specifies the internal PKEU routine to be executed. The mode register is cleared when the PKEU is reset or re-initialized. Setting a reserved mode bit will generate a data error. If the mode register is modified during processing, a context error will be generated.

Figure 18-63 shows the PKEU Mode Register, and Table 18-53 lists the possible values for the ROUTINE field. For channel-controlled access to the PKEU, the descriptor type is determined by the ROUTINE to be used. The descriptor type used with each ROUTINE is listed in Table 18-53.




Figure 18-63. PKEU Mode Register

18.8.5.2 PKEU Key Size Register

The PKEU key size register, shown in Figure 18-64, specifies the number of significant bytes to be used from PKEU parameter memory E in performing modular exponentiation or elliptic curve point multiplication. The range of values for this register, when performing either modular exponentiation or elliptic curve point multiplication, is from 1 to 512. Specifying a key size outside of this range causes a key size error (KSE) in the PKEU interrupt status register.




Figure 18-64. PKEU Key Size Register

Table 18-53. ROUTINE Field Description

| Mode [56-63] | Routine Name | Routine Description | Descriptor Type |
|--------------|--------------------|---|-----------------|
| 0x00 | RESERVED | Reserved | NA |
| 0x01 | CLEARMEMORY | Clear memory | pkeu_mm |
| 0x02 | MOD_EXP | FP: Exponentiate mod N and deconvert from Montgomery format | pkeu_mm |
| 0x03 | MOD_R2MODN | FP: Compute Montgomery converter ($R^2 \text{ mod } N$) | pkeu_mm |
| 0x04 | MOD_RRMODP | FP: Compute Montgomery converter for Chinese Remainder Theorem ($R_n R_p \text{ mod } N$) | pkeu_mm |
| 0x05 | EC_FP_AFF_PTMULT | FP EC: Multiply scalar times point in affine coordinates | pkeu_ptmul |
| 0x06 | EC_F2M_AFF_PTMULT | F2m EC: Multiply scalar times point in affine coordinates | pkeu_ptmul |
| 0x07 | EC_FP_PROJ_PTMULT | FP EC: Multiply scalar times point in projective coordinates | pkeu_ptmul |
| 0x08 | EC_F2M_PROJ_PTMULT | F2m EC: Multiply scalar times point in projective coordinates | pkeu_ptmul |
| 0x09 | EC_FP_ADD | FP EC: Add two points in projective coordinates | pkeu_ptadd_dbl |
| 0x0A | EC_FP_DOUBLE | FP EC: Double a point in projective coordinates | pkeu_ptadd_dbl |

Table 18-53. ROUTINE Field Description (continued)

| Mode [56-63] | Routine Name | Routine Description | Descriptor Type |
|--------------|------------------|---|-----------------|
| 0x0B | EC_F2M_ADD | F2m EC: Add two points in projective coordinates | pkeu_ptadd_dbl |
| 0x0C | EC_F2M_DOUBLE | F2m EC: Double a point in projective coordinates | pkeu_ptadd_dbl |
| 0x0D | F2M_R2 | F2m: Compute Montgomery converter ($R^2 \bmod N$) | pkeu_mm |
| 0x0E | F2M_INV | F2m: Invert mod N | pkeu_mm |
| 0x0F | MOD_INV | FP: Invert mod N | pkeu_mm |
| 0x10 | MOD_ADD | FP: Add mod N | pkeu_mm |
| 0x20 | MOD_SUB | FP: Subtract mod N | pkeu_mm |
| 0x30 | MOD_MULT1_MONT | FP: Multiply mod N in Montgomery format | pkeu_mm |
| 0x40 | MOD_MULT2_DECONV | FP: Multiply mod N and deconvert from Montgomery format | pkeu_mm |
| 0x50 | F2M_ADD | F2m: Add mod N | pkeu_mm |
| 0x60 | F2M_MULT1_MONT | F2m: Multiply mod N in Montgomery format | pkeu_mm |
| 0x70 | F2M_MULT2_DECONV | F2m: Multiply mod N and deconvert from Montgomery format | pkeu_mm |
| 0x80 | RSA_SSTEP | FP: Exponentiate mod N (combines MOD_R2MODN, POLY_F2M_MULT1_MONT, and MOD_EXP) | pkeu_mm |
| 0x1D | MOD_EXP_TEQ | FP: Exponentiate mod N and deconvert from Montgomery format with timing equalization | pkeu_mm |
| 0x1E | RSA_SSTEP_TEQ | FP: Exponentiate mod N with timing equalization (combines MOD_R2MODN, EC_F2M_MULT1_MONT, and MOD_EXP_TEQ) | pkeu_mm |
| 0xFF | SPK_BUILD | Build PK data structure (data structure used by all elliptic curve routines) | pkeu_build |

18.8.5.3 PKEU AB Size Register

The PKEU AB size register, shown in [Figure 18-65](#), represents the size of each operand written into parameter memory A and parameter memory B in bits. An exact size in bits must be provided since a big to little endian realignment is performed based on this value. No error checking is performed as to whether the operand sizes are greater than the prime modulus or the field size written in N-ram. It is assumed that operands are modular reduced before they are written into the PKEU module. This register must be written to before each write to parameter memory A or parameter memory B and must be written before each read of parameter memory A and parameter memory B if the amount of data being taken out is different than the amount of data put in A or B. The value written to the AB_Size register must also adhere to the constraints on parameters A and B, set by the chosen routine (see [Table 18-10](#)). The AB_Size register should not be read or written to while the PKEU is processing as this will cause a ‘data modify during processing error.’ This register is cleared when the PKEU is reset. The maximum size acceptable is 4096 bits.




Figure 18-65. PKEU AB Size Register

18.8.5.4 PKEU Data Size Register

The PKEU data size register, shown in [Figure 18-66](#), specifies, in bits, the size of the significant portion of the modulus or irreducible polynomial. The minimum size valid for all routines to operate properly is 33 bits. The maximum size to operate properly is 4096 bits. A value in bits larger than 4096 results in a data size error (see the DSE bit in [Table 18-56](#)).



Figure 18-66. PKEU Data Size Register

18.8.5.5 PKEU Reset Control Register

The PKEU reset control register, shown in [Figure 18-67](#), contains three reset options specific to the PKEU.




Figure 18-67. PKEU Reset Control Register

[Table 18-54](#) describes the PKEU reset control register fields.

Table 18-54. PKEU Reset Control Register Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–60 | — | Reserved |
| 61 | RI | Reset interrupt. Writing this bit active high causes PKEU interrupts signalling done and error to be reset. It further resets the state of the PKEU interrupt status register. 0 Do not reset 1 Reset interrupt logic |

Table 18-55. PKEU Status Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 61 | EI | Error interrupt: This status bit reflects the state of the error interrupt signal, as sampled by the controller interrupt status register (Section 18.6.4.3, “Interrupt Status Register (ISR)”). 0 PKEU is not signaling error 1 PKEU is signaling error |
| 62 | DI | Done interrupt: This status bit reflects the state of the done interrupt signal, as sampled by the Controller Interrupt Status Register (Section 18.6.4.3, “Interrupt Status Register (ISR)”). 0 PKEU is not signaling done 1 PKEU is signaling done |
| 63 | RD | Reset done. This status bit, when high, indicates that PKEU has completed its reset sequence, as reflected in the signal sampled by the appropriate channel. 0 Reset in progress 1 Reset done Note: Reset Done resets to 0, but has typically switched to 1 by the time a user checks the register, indicating the EU is ready for operation. |

18.8.5.7 PKEU Interrupt Status Register

The PKEU interrupt status register, shown in Figure 18-69, indicates which unmasked errors have generated error interrupts to the channel. Each bit in this register can be set only if the corresponding bit of the PKEU interrupt mask register has a value of zero (see Section 18.8.5.8, “PKEU Interrupt Mask Register”). If the PKEU interrupt status register is non-zero, the PKEU halts and the PKEU error interrupt signal is asserted to the controller (see Section 18.6.4.3, “Interrupt Status Register (ISR)”). In addition, if the PKEU operates through channel-controlled access, an interrupt signal is generated to the channel to which this EU is assigned. The EU error then appears in bit 55 of the channel status register (see Table 18-13) and generates a channel error interrupt to the controller. If the interrupt status register is written from the host, ones in the value written are recorded in the interrupt status register if the corresponding bit is unmasked in the interrupt mask register. All other bits are cleared. This register can also be cleared by setting the RI bit of the PKEU reset control register.




Figure 18-69. PKEU Interrupt Status Register

Table 18-56 describes PKEU interrupt status register signals.

Table 18-56. PKEU Interrupt Status Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–48 | — | Reserved |
| 49 | EVM | Even modulus error. Indicates that an even modulus was supplied for a PK operation that requires an odd modulus. 0 No even modulus error detected 1 Even modulus error detected |
| 50 | INV | Inversion error. Indicates that the inversion routine has a zero operand. 0 No inversion error detected 1 Inversion error detected |
| 51 | IE | Internal error. An internal processing error was detected while the PKEU was operating. 0 No error detected 1 Internal error Note: This bit will be asserted any time an enabled error condition occurs and can only be cleared by setting the corresponding bit in the Interrupt Mask Register or by resetting the PKEU. |
| 52 | BE | Boot error. Indicates that either at boot (reset) sequence, RAM locations are not reset correctly. 0 No boot error detected 1 Boot error |
| 53 | CE | Context error. A PKEU key register, the key size register, the data size register, or mode register was modified while the PKEU was operating. 0 No error detected 1 Context error |
| 54 | KSE | Key size error. Value outside the bounds of 1–256 bytes was written to the PKEU key size register 0 No error detected 1 Key size error detected |
| 55 | DSE | Data size error. Value outside the bounds 97– 2048 bits was written to the PKEU data size register 0 No error detected 1 Data size error detected |
| 56 | ME | Mode error. An illegal value was detected in the mode register. 0 No error detected 1 Mode error Note: Writing to reserved bits in a mode register is a likely source of error. |
| 57 | AE | Address error. Illegal read or write address was detected within the PKEU address space. 0 No error detected 1 Address error |
| 58–63 | — | Reserved |

18.8.5.8 PKEU Interrupt Mask Register

The PKEU interrupt mask register controls the result of detected errors. For a given error (as defined in [Section 18.8.5.7, “PKEU Interrupt Status Register”](#)), if the corresponding bit in this register is set, then the error is disabled; no error interrupt occurs and the interrupt status register is not updated to reflect the error. If the corresponding bit is not set, then upon detection of an error, the PKEU interrupt status register is

updated to reflect the error, causing assertion of the error interrupt signal and causing the module to halt processing.

Table 18-57 describes PKEU interrupt mask register fields.

Table 18-57. PKEU Interrupt Mask Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–48 | — | Reserved |
| 49 | EVM | Even modulus error 0 Even modulus error enabled 1 Even modulus error disabled |
| 50 | INV | Inversion error 0 Inversion error enabled 1 Inversion error disabled |
| 51 | IE | Internal error 0 Internal error enabled 1 Internal error disabled |
| 52 | BE | Boot error 0 Boot error enabled 1 Boot error disabled |
| 53 | CE | Context error 0 Context error enabled 1 Context error disabled |
| 54 | KSE | Key size error 0 Key size error enabled 1 Key size error disabled |
| 55 | DSE | Data size error 0 Data size error enabled 1 Data size error disabled |
| 56 | ME | Mode error 0 Mode error enabled 1 Mode error disabled |
| 57 | AE | Address error 0 Address error enabled 1 Address error disabled |
| 58–63 | — | Reserved |

18.8.5.9 PKEU End_Of_Message Register

The PKEU end_of_message register indicates the start of a new computation. Writing to this register causes the PKEU to execute the function requested by the ROUTINE field, per the contents of the parameter memories. This register has no data size, and during the write operation, the host data bus is not read. Hence, any data value is accepted. Normally, a write operation with a zero data value is performed. Reading from this register is not meaningful, but a zero value is always returned, and no error is generated.

18.8.5.10 PKEU Parameter Memories

The PKEU uses four 4096-bit memories to receive and store operands for PKEU arithmetic operations. In addition, results are stored in one particular parameter memory. Data addressing within these memories is big-endian; that is, the most significant byte is stored in the lowest address.

18.8.5.10.1 PKEU Parameter Memory A

The 4096-bit memory of PKEU parameter memory A is typically used as an input parameter memory space. For modular arithmetic routines, this memory operates as one of the operands of the desired function. For elliptic curve routines, this memory is segmented into four 1024 bit memories and is used to specify particular curve parameters and input values.

18.8.5.10.2 PKEU Parameter Memory B

The 4096-bit memory of is PKEU parameter memory B is typically used as an input parameter memory space, as well as the result memory space. For modular arithmetic routines, this memory operates as one of the operands of the desired function, as well as the result memory space. For elliptic curve routines, this memory is segmented in to four 1024 bit memories, and is used to specify particular curve parameters and input values, as well as to store result values.

18.8.5.10.3 PKEU Parameter Memory E

The 4096-bit memory of PKEU parameter memory E is non-segmentable and specifies the exponent for modular exponentiation or the multiplier k for elliptic curve point multiplication. This memory space is write only; a read of this memory space causes an address error to be reflected in the PKEU interrupt status register.

18.8.5.10.4 PKEU Parameter Memory N

The 4096-bit memory of PKEU parameter memory N is non-segmentable, and specifies the modulus for modular arithmetic and F_p elliptic curve routines. For F_{2^m} elliptic curve routines, this memory specifies the irreducible polynomial.

18.8.6 Random Number Generator (RNGU)

This section contains details about the random number generator, including modes of operation, status and control registers, and FIFO.

The RNGU is an execution unit capable of generating 64-bit random numbers. It combines a True Random Number Generator (TRNG) and an implementation of a NIST-approved Pseudo-Random Number Generator (PRNG) (as described in Annex C of FIPS 140-2 and ANSI X9.62). The RNGU is designed to comply with the FIPS-140 standard for randomness and non-determinism.

The RNGU consists of five major functional blocks:

- 64-bit Skyblue interface, registers, and FIFO
- True Random Number Generator (ring oscillator, LFSRs, Statistical Checker)
- Xseed Generator

- Pseudo-Random Number Generator (XKEY, SHA-1, FSM)
- Simultaneous Reseed LFSR

The states of the LFSRs in the TRNG are advanced at an unknown frequency determined by the ring oscillator clock. The entropy generated by this structure is then added into the XKEY structure of the PRNG during seed generation. Seed generation takes approximately 2,000,000 cycles as 20,000 bits of entropy are sampled from the output of the LFSRs of the TRNG.

After the initial seeding, the RNGU turns off the TRNG and uses solely the PRNG to generate random data. After 1,000,000 times through the algorithm the RNGU is once again seeded. This second seed occurs the next time through the algorithm by using data from the Simultaneous Reseed LFSR to modify the algorithm. The data in the simultaneous reseed LFSR comes directly from the TRNG as well and was being generated during the first 20,000 times through the PRNG algorithm after the initial seed was completed.

Most of the registers described here would not normally be accessed by the host. They are documented here mainly for debug purposes. In typical operation, the RNGU is used through channel-controlled access, which means that most reads and writes of RNGU registers are directed by the SEC channels. Driver software would perform host-controlled register accesses only on a few registers for initial configuration and error handling.

18.8.6.1 RNGU Mode Register

The RNGU mode register is a writable location but all mode bits are reserved. It is documented for the sake of consistency with the other EUs. The RNGU mode register is shown in [Figure 18-70](#).




Figure 18-70. RNGU Mode Register

18.8.6.2 RNGU Data Size Register

The RNGU data size register is used to tell the RNGU to begin generating random data. The actual contents of the data size register do not affect the operation of the RNGU. After a reset and prior to the first write of data size, the RNGU builds entropy without pushing data onto the FIFO. After the data size register is written, the RNGU begins pushing data onto the FIFO. One dword (64 bits) of data is pushed onto the FIFO every 112 cycles until the FIFO is full. The RNGU then attempts to keep the FIFO full.




Figure 18-71. RNGU Data Size Register

18.8.6.8 RNGU ENTROPY

RNGU allows the user to input Entropy into the PRNG algorithm to modify the randomness of the RNGU. This group of registers are write only and all writes to these registers will be ignored when the RNGU is busy. However, when the RNGU is IDLE (FIFO is full or RNGU has not yet been started), all data written to these registers are used to modify the internal XKEY structure. These registers cannot be written back to back, there must be a clock cycle in between writes, so that the RNGU can process all 64-bits of data, as the RNGU processes only 32-bits per cycle.

18.8.6.9 RNGU FIFO

RNGU uses an output FIFO to collect periodically sampled random 64-bit-words, with the intent that random data always be available for reading. Normally, the channels control all access to this FIFO. For host-controlled operation, a read from anywhere in the RNGU FIFO address space dequeues data from the RNGU output FIFO.

The output FIFO is readable by byte, word, or dword. When all 8 bytes of the head dword have been read, that dword is automatically dequeued from the FIFO so that the next dword (if any) becomes available for reading. If any byte is read twice between dequeues, it causes an error interrupt of type AE from the EU.

Underflows caused by reading or writing the RNGU output FIFO are reflected in the RNGU interrupt status register. Also, a write to the RNGU output FIFO space will be reflected as an addressing error in the RNGU interrupt status register.

NOTE

Host reads of the RNGB FIFO should be performed on an 8-byte basis, regardless of how many bits of random number is actually required. Partial host reads can leave the RNGB FIFO in a state that will result in a channel error.



Chapter 19

Enhanced Three-Speed Ethernet Controllers

19.1 Overview

The enhanced three-speed Ethernet controllers (eTSECs) of the device interface to 10 Mbps, 100 Mbps, and 1 Gbps Ethernet/IEEE Std 802.3™ networks. For Ethernet, an external PHY or SerDes device is required to complete the interface to the media. Each eTSEC supports multiple standard media-independent interfaces. Multiple eTSECs are available, providing flexible options for connectivity and control access at different speeds.

The eTSEC provides the flexibility to accelerate the identification and retrieval of standard and non-standard protocols carried over Ethernet, including both IP versions 4 and 6 and TCP/UDP. CPU-intensive parsing and checksum operations can be optionally off-loaded to an eTSEC to accelerate existing TCP/IP stacks. On transmission, varying fractions of link bandwidth can be allocated to each of multiple transmit queues through a modified weighted round-robin scheduler. On receive, an arbitrary set of queue selection rules can be programmed into each eTSEC to implement flexible quality of service or firewall strategies based on high-level protocol identification. Without enabling these advanced features, each eTSEC emulates a PowerQUICC II Pro TSEC, allowing existing driver software to be re-used with minimal change. Each eTSEC is organized as shown in [Figure 19-1](#).

NOTE

The eTSECs do not support TBI, GMII, and FIFO operating modes, so all references to these interfaces and features should be ignored for this device.




Figure 19-1. eTSEC Block Diagram

19.2 Features

The eTSECs of the device include these distinctive features:

- IEEE 802.3, 802.3u, 802.3x, 802.3z, 802.3ac, 802.3ab compatible
- Support for different Ethernet physical interfaces:
 - 10/100 Mbps IEEE 802.3 MII and RMII
 - 10/100 Mbps RGMII
 - 1000 Mbps full-duplex RGMII and RTBI

- 10/100 Mbps SGMII
- 1000 Mbps full-duplex SGMII
- TCP/IP off-load
 - IP v4 and IP v6 header recognition on receive
 - IP v4 header checksum verification and generation
 - TCP and UDP checksum verification and generation
 - Per-packet configurable off-load
 - Recognition of VLAN, stacked-VLAN, 802.2, PPPoE session, MPLS stacks, and ESP/AH IP-Security headers
- Quality of service (QoS) support
 - Transmission from up to eight queues
 - Priority-based queue selection
 - Modified weighted round-robin queue selection with fair bandwidth allocation
 - Reception to up to eight physical queues
 - 64 virtual receive queues overlaid on 8 physical buffer descriptor rings
 - Table-oriented queue filing strategy based on 16 header fields or flags
 - Frame rejection support for filtering applications
 - Filing based on Ethernet, IP, and TCP/UDP properties, including VLAN fields, Ether-type, IP protocol type, IP TOS or differentiated services, IP source and destination addresses, TCP/UDP port numbers
- Interrupt coalescing
 - Packet-count-based thresholds for both receive and transmit
 - Timer-based thresholds
- Full- and half-duplex Ethernet support (1000 Mbps supports only full duplex):
 - IEEE 802.3 full-duplex flow control (automatic PAUSE frame generation or software programmed PAUSE frame generation and recognition)
 - Programmable maximum frame length supports jumbo frames (up to 9.6 Kbytes) and IEEE 802.1 virtual local area network (VLAN) tags and priority
 - VLAN insertion and deletion
 - Per-frame VLAN control word or default VLAN for each eTSEC
 - Extracted VLAN control word passed to software separately
 - Programmable VLAN tag to support metropolitan bridging
 - Retransmission following a collision
 - Support for CRC generation and verification of inbound/outbound packets
 - Programmable Ethernet preamble insertion and extraction of up to 7 bytes
- MAC address recognition:
 - Exact match on primary and virtual 48-bit unicast addresses
 - VRRP and HSRP support for seamless router fail-over

- In addition to primary station address, up to fifteen additional exact-match MAC addresses supported
 - Broadcast address (accept/reject)
 - Hash table match on up to 256 unicast/multicast or 512 multicast-only addresses
 - Promiscuous mode
- Remote network monitoring (RMON) statistics support
 - 32-bit byte counters
 - Carry/Overflow of counter interrupts
- Backward compatibility with MPC8349E (PowerQUICC II Pro) TSEC
 - PowerQUICC II Pro buffer descriptor (BD) format and rings supported
 - Common register memory map, with specific exceptions:
 - Out-of-sequence transmit BD not supported
 - Internal DMA BD pointers and data counts not visible
 - MINFLR register not supported
 - Reset state of eTSEC defaults to common PowerQUICC II Pro TSEC subset
 - TSEC_ID register permits TSEC versus enhanced TSEC differentiation
- Hardware assist for 1588-compliant timestamping (1588 not supported in conjunction with SGMII 10/100)
 - Per packet timestamp tag for Receive
 - Programmable timestamp capture for Transmit
 - Recognition of PTP packet
 - Periodic Pulse Generation
 - Self-correcting precision timer with nano-second resolution
 - Phase aligned adjustable (divide by N) clock output
 - Two 64-bit alarm (future time) registers for future time comparison

19.3 Modes of Operation

The eTSEC's primary operational modes are the following:

- Full- and half-duplex operation

This is determined by the MACCFG2 register's full-duplex bit (MACCFG2[Full Duplex]). Full-duplex mode is intended for use on point-to-point links between switches or end node to switch. Half-duplex mode is used in connections between an end node and a repeater or between repeaters.

If configured in half-duplex mode (10- and 100-Mbps operation; MACCFG2[Full Duplex] is cleared), the MAC complies with the IEEE CSMA/CD access method.

If configured in full-duplex mode (10/100/1000 Mbps operation; MACCFG2[Full Duplex] is set), the MAC supports flow control. If flow control is enabled, it allows the MAC to receive or send PAUSE frames.

- 10- and 100-Mbps MII interface operation

The MAC-PHY interface operates in MII mode by setting `MACCFG2[I/F Mode] = 01`. The MII is the media-independent interface defined by the 802.3 standard for 10/100 Mbps operation. The speed of operation is determined by the `TSECn_TX_CLK` and `TSECn_RX_CLK` signals, which are driven by the transceiver. The transceiver either auto-negotiates the speed, or it may be controlled by software using the serial management interface (MDC/MDIO signals) to the transceiver.

Clause 22.2.4 of the IEEE 802.3 specification describes the MII management interface.
- 10- and 100-Mbps RMI interface operation

The RMI is the reduced media-independent interface defined by the RMI Consortium (March 1998) for 10/100 Mbps operation. The speed of operation is determined by the `TSECn_TX_CLK` signal, which is driven by the transceiver.
- MAC address recognition options

The options supported are promiscuous, broadcast, exact unicast address match, exact unicast virtual address match to support router redundancy, and multicast hash match. For detailed descriptions refer to [Section 19.6.2.7, “Frame Recognition.”](#)

eTSEC supports automatic LAN-initiated wake-up during power management through the AMD Magic Packet protocol, as described in [Section 19.6.2.8, “Magic Packet Mode.”](#)
- Receive frame parsing options

Frame parsing options are to disable parsing (no TCP/IP off-load), IP header parsing, and TCP or UDP parsing. Parsing must be enabled to make use of receive queue filing algorithms. The options are detailed in [Section 19.6.3, “TCP/IP Off-Load.”](#)
- Receive queue selection options

Received frames are by default sent to a single buffer descriptor ring. If multiple receive queues are enabled, a receive queue filer can be programmed with selection criteria to differentiate received frames and file them to different buffer descriptor rings. See [Section 19.6.4, “Quality of Service \(QoS\) Provision,”](#) for detailed descriptions.
- TCP/IP transmit options

Frames for transmission may be sent as-is, with IP header processing, or TCP header processing. The transmit buffer descriptors, described in [Section 19.6.7.2, “Transmit Data Buffer Descriptors \(TxBD\),”](#) enable these options and operate with parameters prepended to frame buffers, as described in [Section 19.6.3, “TCP/IP Off-Load.”](#)
- Transmit queue selection options

The options supported are single transmit queue, priority-based queue selection, and modified weighted round-robin queueing. These options are described further in [Section 19.5.3.2.1, “Transmit Control Register \(TCTRL\).”](#)
- RMON support

Standard Ethernet interface management information base (MIBs) can be generated through the RMON MIB counters.
- Internal loop back supported for all interfaces except when configured for half-duplex operation

Internal loop back mode is selected through the loop back bit in the MACCFG1 register. See [Section 19.7.1, “Interface Mode Configuration,”](#) for details.

19.4 External Signals Description

This section defines the eTSEC interface signals. The buses are described using the bus convention used in IEEE 802.3 because the PHY follows this same convention. (That is, TxD[3:0] means 0 is the lsb.) Note that except for external physical interfaces the buses and registers follow a big-endian format, where 0 denotes the msb.

Each eTSEC network interface supports multiple options:

- The MII option requires 18 I/O signals (including the MDIO and MDC MII management interface) and supports both a data and a management interface to the PHY (transceiver) device. The MII option supports both 10- and 100-Mbps Ethernet rates.
- The RGMII, RTBI, and RMII options are reduced-pin implementations of the GMII, TBI, and MII interfaces, respectively.
- SGMII interfaces are offered via the SerDes interface signals.
- 1588 timer signals

[Table 19-1](#) lists the network interface signals.

Table 19-1. eTSEC_n Network Interface Signal Properties

| Signal Name | Function | Reset State |
|-----------------------------|---|--------------|
| TSEC _n _COL | MII—collision, input | — |
| TSEC _n _CRS | MII—carrier sense, input | — |
| TSEC _n _GTX_CLK | RTBI, RGMII— <i>inverted</i> transmit clock feedback, output MII, RMII—transmit clock feedback when transmission is enabled, zero otherwise, output | 0 |
| EC_GTX_CLK125 | Oscillator source for RGMII, RTBI transmit clock, input, shared by all eTSECs | — |
| EC_MDC | Management clock, output. | 0 |
| EC_MDIO | Management data, bidirectional. | Hi-Z (input) |
| TSEC _n _RX_CLK | MII, RGMII—receive clock, input | — |
| TSEC _n _RX_DV | MII—receive data valid, input RGMII (RX_CLK rising)—receive data valid, input RGMII (RX_CLK falling)—receive error, input RTBI (RX_CLK rising)—receive code group (RCG) bit 4, input RTBI (RX_CLK falling)—receive code group (RCG) bit 9, input RMII—CRS_DV carrier sense/data valid, input | — |
| TSEC _n _RXD[3:0] | MII—Receive data bits 3:0, input RGMII (RX_CLK rising) —Receive data bits 3:0, input RGMII (RX_CLK falling)—Receive data bits 7:4, input RTBI (RX_CLK rising)—RCG bits 3:0, input RTBI (RX_CLK falling)—RCG bits 8:5, input RMII—RXD[1:0] receive data bits, input RMII—RXD[3:2] are unused | — |
| TSEC _n _RX_ER | MII, RMII—Receive error, input RGMII, RTBI—Unused | — |

Table 19-1. eTSEC_n Network Interface Signal Properties (continued)

| Signal Name | Function | Reset State |
|--|--|-------------|
| TSEC _n _TX_CLK | MII—transmit clock, input RMII—reference transmit and receive clock, input RGMII, RTBI—unused | — |
| TSEC _n _TXD[3:0] | MII—Transmit data bits 3:0, output RGMII (TX_CLK rising)—Transmit data bits 3:0, output RGMII (TX_CLK falling)—Transmit data bits 7:4, output RTBI (TX_CLK rising)—TCG bits 3:0, output RTBI (TX_CLK falling)—TCG bits 8:5, output RMII—TXD[1:0] transmit data bits, output RMII—TXD[3:2] unused | 0000 |
| TSEC _n _TX_ER | MII—transmit error, output RGMII, RTBI, RMII—unused, output driven zero | 0 |
| TSEC _n _TX_EN | MII, RMII—Transmit data valid, output RGMII (TX_CLK rising)—Transmit data enabled, output RGMII (TX_CLK falling)—Transmit error, output RTBI (TX_CLK rising)—TCG bit 4, output RTBI (TX_CLK falling)—TCG bit 9, output | 0 |
| TSEC_TMR_CLK | 1588—Clock input External high precision timer reference clock input (chip external input pin). | — |
| TSEC_TMR_GCLK | 1588—Clock output Phase aligned timer clock divider output (chip external output pin). | 0 |
| TSEC_TMR_TRIG1 | 1588—Trigger in 1 External timer trigger input 1. This is an asynchronous general purpose input (chip external input pin). | — |
| TSEC_TMR_TRIG2 | 1588—Trigger in 2 External timer trigger input 2. This is an asynchronous general purpose input (chip external input pin). | — |
| TSEC_TMR_PP1 | 1588—Pulse out 1 Timer pulse per period 1. It is phase aligned with 1588 timer clock (chip external output pin). | 0 |
| TSEC_TMR_PP2 | 1588—Pulse out 2 Timer pulse per period 2. It is phase aligned with 1588 timer clock (chip external output pin). | 0 |
| TSEC_TMR_PP3 | 1588—Pulse out 3 Timer pulse per period 3. It is phase aligned with 1588 timer clock (chip external output pin). | 0 |
| TSEC_TMR_ALARM1 | 1588—Timer alarm 1 Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_TMR_ALARM _n _H/L register to deactivate this output (chip external output pin). | 0 |
| TSEC_TMR_ALARM2 | 1588—Timer alarm 2 Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_TMR_ALARM _n _H/L register to deactivate this output (chip external output pin). | 0 |
| SD2_TX[<i>n</i> -1] SD2_TX[<i>n</i> -1] | SGMII transmit data (and complement) | — |

Table 19-1. eTSEC_n Network Interface Signal Properties (continued)

| Signal Name | Function | Reset State |
|----------------------------|--|-------------|
| SD2_RX[n-1] SD2_RX[n-1] | SGMII receive data (and complement) | — |
| SD2_REF_CLK SD2_REF_CLK | SGMII SerDes2 PLL reference clock (and complement) | — |

19.4.1 Detailed Signal Descriptions

Below is a description of the eTSEC interface signals. For RGMII mode details, refer to the Hewlett-Packard reduced gigabit media-independent interface (RGMII) specification version 1.2a, dated 9/22/2000. RMII mode details follow the RMII Consortium Specification, dated 3/20/1998. All other modes follow the IEEE 802.3 standard, 2000 Edition. Input signals not used are internally disabled. Except for TSEC_n_GTX_CLK, output signals not used are driven low.

Table 19-2. eTSEC Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|----------------------------|-----|---|
| TSEC _n _COL | I | Collision input. The behavior of this signal is not specified while in full-duplex mode. |
| | | State Meaning Asserted/Negated—In MII mode, this signal is asserted upon detection of a collision, and must remain asserted while the collision persists. This signal is not used in the following modes: <ul style="list-style-type: none"> • RMII • RTBI • RGMII |
| | | Timing Asserted/Negated—This signal is not required to transition synchronously with TSEC _n _TX_CLK or TSEC _n _RX_CLK. |
| TSEC _n _CRS | I | Carrier sense input. In RTBI mode, this signal is used as SDET (signal detect). In RTBI mode SDET is tied high internally. This signal is not used in the following modes: <ul style="list-style-type: none"> • RMII • RGMII |
| | | State Meaning Asserted/Negated—In MII mode, TSEC _n _CRS is asserted while the transmit or receive medium is not idle. In the event of a collision, TSEC _n _CRS must remain asserted for the duration of the collision. |
| | | Timing Asserted/Negated—This signal is not required to transition synchronously with TSEC _n _TX_CLK or TSEC _n _RX_CLK. |
| TSEC _n _GTX_CLK | O | Gigabit transmit clock. This signal is an output from the eTSEC into the PHY. TSEC _n _GTX_CLK is a 125-MHz clock that provides a timing reference for TX_EN, TXD, and TX_ER in the following modes: <ul style="list-style-type: none"> • RTBI In RGMII mode, TSEC _n _GTX_CLK becomes the transmit clock and provides timing reference during 1000Base-T (125 MHz), 100Base-T (25 MHz) and 10Base-T (2.5 MHz) transmissions. This signal feeds back the uninverted transmit clock in MII mode, but feeds back an inverted transmit clock in RTBI or RGMII modes. This signal is driven low unless transmission is enabled. |

Table 19-2. eTSEC Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|-----------------------------|-----|---|
| EC_GTX_CLK125 | I | Gigabit transmit 125-MHz source. This signal must be generated externally with a crystal or oscillator, or is sometimes provided by the PHY. EC_GTX_CLK125 is a 125-MHz input into the eTSEC and is used to generate all 125-MHz related signals and clocks in the following modes: <ul style="list-style-type: none"> • RTBI • RGMII This input is not used in these modes: <ul style="list-style-type: none"> • RMII • SGMII • MII |
| EC_MDC | O | Management data clock. This signal is a clock (typically 2.5 MHz) supplied by the MAC (IEEE set minimum period of 400 ns or a frequency of 2.5 MHz, but the device may be configured up to 12.5 MHz if supported by the PHY at that speed.) The frequency can be modified by writing to MIIMCFG[28:31] of the eTSEC1 controller. |
| EC_MDIO | I/O | Management data input/output. |
| | | State Meaning Asserted/Negated—EC_MDIO is a bidirectional signal to input PHY-supplied status during management read cycles and output control during MII management write cycles. Addressed using eTSEC1 memory-mapped registers. |
| | | Timing Asserted/Negated—This signal is required to be synchronous with the EC_MDC signal. |
| TSEC _n _RX_CLK | I | Receive clock. In MII or RGMII mode, the receive clock TSEC _n _RX_CLK is a continuous clock (2.5, 25, or 125 MHz) that provides a timing reference for TSEC _n _RX_DV, TSEC _n _RXD, and TSEC _n _RX_ER. In RTBI mode it is a 125-MHz receive clock. In RMII mode this clock is not used for the receive clock, as RMII uses a shared reference clock. |
| TSEC _n _RX_DV | I | Receive data valid. In MII mode, if TSEC _n _RX_DV is asserted, the PHY is indicating that valid data is present on the MII interface. In RGMII mode, TSEC _n _RX_DV becomes RX_CTL. The RX_DV and RX_ERR are received on this signal on the rising and falling edges of TSEC _n _RX_CLK. In RTBI mode, TSEC _n _RX_DV represents receive code group (RCG) bit 4 and 9. On the positive edge of the TSEC _n _RX_CLK, RCG[4] and RCG[3:0] represent the first half of the 10-bit encoded symbol. On the negative edge of the TSEC _n _RX_CLK, RCG[9] and RCG[8:5] represent the second half of the 10-bit encoded symbol. In RMII mode the PHY asserts TSEC _n _RX_DV (CRS_DV) when the receive medium is non-idle. This signal asserts asynchronously with respect to the RMII reference clock, but negates synchronously to indicate loss of carrier. |
| TSEC _n _RXD[3:0] | I | Receive data in. In MII mode, TSEC _n _RXD[3:0] represents a nibble of data to be transferred from the PHY to the MAC when TSEC _n _RX_DV is asserted. A completely formed SFD must be passed across the MII. While TSEC _n _RX_DV is not asserted, TSEC _n _RXD has no meaning. In RGMII mode, data bits 3:0 are received on the rising edge of TSEC _n _RX_CLK and data bits 7:4 are received on the falling edge of TSEC _n _RX_CLK. In RTBI mode, TSEC _n _RXD[3:0] represents RCG[3:0] on the rising edge of TSEC _n _RX_CLK and RCG[8:5] are received on the falling edge of TSEC _n _RX_CLK. In RMII mode, TSEC _n _RXD[1:0] represents RXD[1:0], which is considered valid when TSEC _n _RX_DV (CRS_DV) is asserted, or invalid otherwise. |

Table 19-2. eTSEC Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|-----------------------------|-----|--|
| TSEC _n _RX_ER | I | Receive error |
| | | State Meaning |
| TSEC _n _TX_CLK | I | Transmit clock in. In MII mode, TSEC _n _TX_CLK is a continuous clock (2.5 or 25 MHz) that provides a timing reference for the TSEC _n _TX_EN, TSEC _n _TXD, and TSEC _n _TX_ER signals. In RMII mode this signal is the reference clock shared between transmit and receive, and is supplied by the PHY. This signal is not used in the eTSEC RTBI or RGMII modes. |
| TSEC _n _TXD[3:0] | O | Transmit data out. DVIn MII mode, TSEC _n _TXD[3:0] represent a nibble of data to be sent from the MAC to the PHY when TSEC _n _TX_EN is asserted and have no meaning while TSEC _n _TX_EN is negated. In RGMII or RTBI mode, data bits 3:0 are transmitted on the rising edge of TSEC _n _GTX_CLK, and data bits 7:4 are transmitted on the falling edge of TSEC _n _GTX_CLK. In RMII mode TSEC _n _TXD[1:0] represents TXD[1:0], which is valid data sent to the PHY when TSEC _n _TX_EN is asserted, or undefined otherwise. Note that some of these signals are also used during reset to configure the eTSEC interface mode. |
| TSEC _n _TX_EN | O | Transmit data valid. In MII, or RMII mode, if TSEC _n _TX_EN is asserted, the MAC is indicating that valid data is present on the MII's TSEC _n _TXD signals. In RGMII mode, TSEC _n _TX_EN becomes TX_CTL. TX_EN and TX_ERR are asserted on this signal on rising and falling edges of the TSEC _n _GTX_CLK, respectively. In RTBI mode, TSEC _n _TX_EN represents TCG[4] on the rising edge and TCG[9] on the falling edge of TSEC _n _GTX_CLK, respectively. Together with TCG[3:0] and TCG[8:5], they represent the 10-bit encoded symbol. |
| TSEC _n _TX_ER | O | Transmit error. In MII mode, assertion of TSEC _n _TX_ER for one or more clock cycles while TSEC _n _TX_EN is asserted causes the PHY to transmit one or more illegal symbols. Asserting TSEC _n _TX_ER has no effect while operating at 10 Mbps or while TSEC _n _TX_EN is negated. This signal transitions synchronously with respect to TSEC _n _TX_CLK. This signal is not used in the eTSEC RMII, RTBI, or RGMII modes and is driven low. |
| TSEC_TMR_CLK | I | 1588 clock in. External high precision timer reference clock input (chip external input pin). |
| TSEC_TMR_GCLK | O | 1588 clock out. Phase aligned timer clock divider output (chip external output pin). |
| TSEC_TMR_TRIG1 | I | 1588 trigger in 1. External timer trigger input 1. This is an asynchronous general purpose input (chip external input pin). |
| TSEC_TMR_TRIG2 | i | 1588 trigger in 2. External timer trigger input 2. This is an asynchronous general purpose input (chip external input pin). |
| TSEC_TMR_PP1 | O | 1588 pulse out 1. Timer pulse per period 1. It is phase aligned with 1588 timer clock (chip external output pin) |
| TSEC_TMR_PP2 | O | 1588 pulse out 2. Timer pulse per period 2. It is phase aligned with 1588 timer clock (chip external output pin) |
| TSEC_TMR_PP3 | O | 1588 pulse out 3. Timer pulse per period 3. It is phase aligned with 1588 timer clock (chip external output pin) |
| TSEC_TMR_ALARM1 | O | 1588 timer alarm 1. Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_TMR_ALARM _n _H/L register to deactivate this output (chip external output pin) |

Table 19-2. eTSEC Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|--|-----|---|
| TSEC_TMR_ALARM2 | O | 1588 timer alarm 2. Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_TMR_ALARM n _H/L register to deactivate this output (chip external output pin) |
| $\overline{\text{SD2_TX}}[n-1]$ SD2_TX[n-1] | O | SGMII transmit data (and complement) When in SGMII interface mode: <ul style="list-style-type: none"> eTSEC1 utilizes SD2_TX[0] and $\overline{\text{SD2_TX}}[0]$ eTSEC2 utilizes SD2_TX[1] and $\overline{\text{SD2_TX}}[1]$ eTSEC3 utilizes SD2_TX[2] and $\overline{\text{SD2_TX}}[2]$ |
| $\overline{\text{SD2_RX}}[n-1]$ SD2_RX[n-1] | I | SGMII receive data (and complement) When in SGMII interface mode: <ul style="list-style-type: none"> eTSEC1 utilizes SD2_RX[0] and $\overline{\text{SD2_RX}}[0]$ eTSEC2 utilizes SD2_RX[1] and $\overline{\text{SD2_RX}}[1]$ eTSEC3 utilizes SD2_RX[2] and $\overline{\text{SD2_RX}}[2]$ |
| $\overline{\text{SD2_REF_CLK}}$ SD2_REF_CLK | I | SGMII SerDes2 PLL reference clock (and complement) |

19.5 Memory Map/Register Definition

The eTSECs use a software model that is a superset of the PowerQUICC II Pro TSEC functionality and is similar to that employed by the Fast Ethernet function supported on the Freescale MPC8260 CPM FCC and in the FEC of the MPC860T.

The eTSEC device is programmed by a combination of control/status registers (CSRs) and buffer descriptors. The CSRs are used for mode control, interrupts, and to extract status information. The descriptors are used to pass data buffers and related buffer status or frame information between the hardware and software.

All accesses to and from the registers must be made as 32-bit accesses. There is no support for accesses of sizes other than 32 bits. Writes to reserved register bits must always store 0, as writing 1 to reserved bits may have unintended side-effects. Reads from unmapped register addresses return zero. Unless otherwise specified, the read value of reserved bits in mapped registers is not defined, and must not be assumed to be 0.

This section of the document defines the memory map and describes the registers in detail. The buffer descriptor is described in [Section 19.6.7, “Buffer Descriptors.”](#)

The ten-bit interface (TBI) and reduced ten-bit interface (RTBI) module MII registers are also described in this section. The TBI/RTBI registers are defined like PHY registers and, as such, are accessed through the MII management interface in the same way the PHYs are accessed. For detailed descriptions of the TBI/RTBI registers (the MII register set for the ten-bit interface) refer to [Section 19.5.4, “Ten-Bit Interface \(TBI\).”](#)

19.5.1 Top-Level Module Memory Map

Each of the eTSECs is allocated 4 Kbytes of memory-mapped space. The space for each eTSEC is divided as indicated in [Table 19-3.](#)

Table 19-3. Module Memory Map Summary

| Address Offset | Function |
|----------------|---|
| 000–0FF | eTSEC general control/status registers |
| 100–2FF | eTSEC transmit control/status registers |
| 300–4FF | eTSEC receive control/status registers |
| 500–5FF | MAC registers |
| 600–7FF | RMON MIB registers |
| 800–8FF | Hash table registers |
| 900–AFF | — |
| B00–BFF | DMA system registers |
| C00–C3F | Lossless Flow Control registers |
| C40–DFF | — |
| E00–EFF | 1588 Hardware Assist |

19.5.2 Detailed Memory Map

The eTSEC memory-mapped registers are accessed by reading and writing to an address comprised of the base address (specified in IMMRBAR as defined in [Chapter 3, “Memory Map.”](#)) plus the block base address, plus the offset of the specific register to be accessed. Note that all memory-mapped registers must only be accessed as 32-bit quantities.

[Table 19-4](#) lists the offset, name, and a cross-reference to the complete description of each register. The offsets to the memory map table are applicable to each eTSEC. Block base addresses are as follows:

- eTSEC1 starts at 0x2_4000 address offset
- eTSEC2 starts at 0x2_5000 address offset

In this table and in the register figures and field descriptions, the following access definitions apply:

- Reserved fields are always ignored for the purposes of determining access type.
- R/W, R, and W (read/write, read only, and write only) indicate that all the non-reserved fields in a register have the same access type.
- w1c indicates that all of the non-reserved fields in a register are cleared by writing ones to them.
- Mixed indicates a combination of access types.
- Special is used when no other category applies. In this case the register figure and field description table should be read carefully.

Table 19-4. Module Memory Map

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|--|---------------------------------|---------------------|-------------|----------------------------------|
| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
| eTSEC General Control and Status Registers | | | | |
| 0x2_4000 | TSEC_ID*—Controller ID register | R | 0x0124_0000 | 19.5.3.1.1/19-22 |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|--|---|---------------------|-------------|----------------------------------|
| 0x2_4004 | TSEC_ID2*—Controller ID register | R | 0x0030_00F0 | 19.5.3.1.2/19-23 |
| 0x2_4008– 0x2_400C | Reserved | — | — | — |
| 0x2_4010 | IEVENT—Interrupt event register | w1c | All zeros | 19.5.3.1.3/19-24 |
| 0x2_4014 | IMASK—Interrupt mask register | R/W | All zeros | 19.5.3.1.4/19-28 |
| 0x2_4018 | EDIS—Error disabled register | R/W | All zeros | 19.5.3.1.5/19-30 |
| 0x2_401C | Reserved | — | — | — |
| 0x2_4020 | ECNTRL—Ethernet control register | R/W | All zeros | 19.5.3.1.6/19-31 |
| 0x2_4024 | Reserved | — | — | — |
| 0x2_4028 | PTV—Pause time value register | R/W | All zeros | 19.5.3.1.7/19-33 |
| 0x2_402C | DMACTRL—DMA control register | R/W | All zeros | 19.5.3.1.8/19-34 |
| 0x2_4030 | TBIPA—TBI PHY address register | R/W | All zeros | 19.5.3.1.9/19-35 |
| 0x2_4034– 0x2_40FC | Reserved | — | — | — |
| eTSEC Transmit Control and Status Registers | | | | |
| 0x2_4100 | TCTRL—Transmit control register | R/W | All zeros | 19.5.3.2.1/19-36 |
| 0x2_4104 | TSTAT—Transmit status register | w1c | All zeros | 19.5.3.2.2/19-38 |
| 0x2_4108 | DFVLAN*—Default VLAN control word | R/W | 0x8100_0000 | 19.5.3.2.3/19-42 |
| 0x2_410C | Reserved | — | — | — |
| 0x2_4110 | TXIC—Transmit interrupt coalescing register | R/W | All zeros | 19.5.3.2.4/19-43 |
| 0x2_4114 | TQUEUE*—Transmit queue control register | R/W | 0x0000_8000 | 19.5.3.2.5/19-44 |
| 0x2_4118– 0x2_413C | Reserved | — | — | — |
| 0x2_4140 | TR03WT*—TxBD Rings 0–3 round-robin weightings | R/W | All zeros | 19.5.3.2.6/19-44 |
| 0x2_4144 | TR47WT*—TxBD Rings 4–7 round-robin weightings | R/W | All zeros | 19.5.3.2.7/19-45 |
| 0x2_4148– 0x2_4180 | Reserved | — | — | — |
| 0x2_4184 | TBPTR0—TxBD pointer for ring 0 | R/W | All zeros | 19.5.3.2.8/19-46 |
| 0x2_4188 | Reserved | — | — | — |
| 0x2_418C | TBPTR1*—TxBD pointer for ring 1 | R/W | All zeros | 19.5.3.2.8/19-46 |
| 0x2_4190 | Reserved | — | — | — |
| 0x2_4194 | TBPTR2*—TxBD pointer for ring 2 | R/W | All zeros | 19.5.3.2.8/19-46 |
| 0x2_4198 | Reserved | — | — | — |
| 0x2_419C | TBPTR3*—TxBD pointer for ring 3 | R/W | All zeros | 19.5.3.2.8/19-46 |
| 0x2_41A0 | Reserved | — | — | — |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|-----------------------|--|---------------------|-----------|-----------------------------------|
| 0x2_41A4 | TBPTR4*—TxBD pointer for ring 4 | R/W | All zeros | 19.5.3.2.8/19-46 |
| 0x2_41A8 | Reserved | — | — | — |
| 0x2_41AC | TBPTR5*—TxBD pointer for ring 5 | R/W | All zeros | 19.5.3.2.8/19-46 |
| 0x2_41B0 | Reserved | — | — | — |
| 0x2_41B4 | TBPTR6*—TxBD pointer for ring 6 | R/W | All zeros | 19.5.3.2.8/19-46 |
| 0x2_41B8 | Reserved | — | — | — |
| 0x2_41BC | TBPTR7*—TxBD pointer for ring 7 | R/W | All zeros | 19.5.3.2.8/19-46 |
| 0x2_41C0– 0x2_4200 | Reserved | — | — | — |
| 0x2_4204 | TBASE0—TxBD base address of ring 0 | R/W | All zeros | 19.5.3.2.9/19-47 |
| 0x2_4208 | Reserved | — | — | — |
| 0x2_420C | TBASE1*—TxBD base address of ring 1 | R/W | All zeros | 19.5.3.2.9/19-47 |
| 0x2_4210 | Reserved | — | — | — |
| 0x2_4214 | TBASE2*—TxBD base address of ring 2 | R/W | All zeros | 19.5.3.2.9/19-47 |
| 0x2_4218 | Reserved | — | — | — |
| 0x2_421C | TBASE3*—TxBD base address of ring 3 | R/W | All zeros | 19.5.3.2.9/19-47 |
| 0x2_4220 | Reserved | — | — | — |
| 0x2_4224 | TBASE4*—TxBD base address of ring 4 | R/W | All zeros | 19.5.3.2.9/19-47 |
| 0x2_4228 | Reserved | — | — | — |
| 0x2_422C | TBASE5*—TxBD base address of ring 5 | R/W | All zeros | 19.5.3.2.9/19-47 |
| 0x2_4230 | Reserved | — | — | — |
| 0x2_4234 | TBASE6*—TxBD base address of ring 6 | R/W | All zeros | 19.5.3.2.9/19-47 |
| 0x2_4238 | Reserved | — | — | — |
| 0x2_423C | TBASE7*—TxBD base address of ring 7 | R/W | All zeros | 19.5.3.2.9/19-47 |
| 0x2_4240– 0x2_427C | Reserved | — | — | — |
| 0x2_4280 | TMR_TXTS1_ID* - Tx time stamp identification tag (set 1) | R/W | All zeros | 19.5.3.2.10/19-47 |
| 0x2_4284 | TMR_TXTS2_ID* - Tx time stamp identification tag (set 2) | R/W | All zeros | 19.5.3.2.10/19-47 |
| 0x2_4288– 0x2_42BC | Reserved | — | — | — |
| 0x2_42C0 | TMR_TXTS1_H* - Tx time stamp high (set 1) | R/W | All zeros | 19.5.3.2.11/19-48 |
| 0x2_42C4 | TMR_TXTS1_L* - Tx time stamp high (set 1) | R/W | All zeros | 19.5.3.2.11/19-48 |
| 0x2_42C8 | TMR_TXTS2_H* - Tx time stamp high (set 2) | R/W | All zeros | 19.5.3.2.11/19-48 |
| 0x2_42CC | TMR_TXTS2_L* - Tx time stamp high (set 2) | R/W | All zeros | 19.5.3.2.11/19-48 |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|---|---|---------------------|--|-----------------------------------|
| 0x2_42D0–0x2_42FC | Reserved | — | — | — |
| eTSEC Receive Control and Status Registers | | | | |
| 0x2_4300 | RCTRL—Receive control register | R/W | All zeros | 19.5.3.3.1/19-48 |
| 0x2_4304 | RSTAT—Receive status register | w1c | All zeros | 19.5.3.3.2/19-50 |
| 0x2_4308–0x2_430C | Reserved | — | — | — |
| 0x2_4310 | RXIC—Receive interrupt coalescing register | R/W | All zeros | 19.5.3.3.3/19-52 |
| 0x2_4314 | RQUEUE*—Receive queue control register. | R/W | 0x0080_0080 | 19.5.3.3.4/19-53 |
| 0x2_4318–0x2_432C | Reserved | — | — | — |
| 0x2_4330 | RBIFX*—Receive bit field extract control register | R/W | All zeros | 19.5.3.3.5/19-54 |
| 0x2_4334 | RQFAR*—Receive queue filing table address register | R/W | All zeros | 19.5.3.3.6/19-56 |
| 0x2_4338 | RQFCR*—Receive queue filing table control register | R/W | 0xn _{nnnn} _n _{nnnn} | 19.5.3.3.7/19-56 |
| 0x2_433C | RQFPR*—Receive queue filing table property register | R/W | 0xn _{nnnn} _n _{nnnn} | 19.5.3.3.8/19-57 |
| 0x2_4340 | MRBLR—Maximum receive buffer length register | R/W | All zeros | 19.5.3.3.9/19-61 |
| 0x2_4344–0x2_4380 | Reserved | — | — | — |
| 0x2_4384 | RBPTR0—RxB _D pointer for ring 0 | R/W | All zeros | 19.5.3.3.10/19-61 |
| 0x2_4388 | Reserved | — | — | — |
| 0x2_438C | RBPTR1*—RxB _D pointer for ring 1 | R/W | All zeros | 19.5.3.3.10/19-61 |
| 0x2_4390 | Reserved | — | — | — |
| 0x2_4394 | RBPTR2*—RxB _D pointer for ring 2 | R/W | All zeros | 19.5.3.3.10/19-61 |
| 0x2_4398 | Reserved | — | — | — |
| 0x2_439C | RBPTR3*—RxB _D pointer for ring 3 | R/W | All zeros | 19.5.3.3.10/19-61 |
| 0x2_43A0 | Reserved | — | — | — |
| 0x2_43A4 | RBPTR4*—RxB _D pointer for ring 4 | R/W | All zeros | 19.5.3.3.10/19-61 |
| 0x2_43A8 | Reserved | — | — | — |
| 0x2_43AC | RBPTR5*—RxB _D pointer for ring 5 | R/W | All zeros | 19.5.3.3.10/19-61 |
| 0x2_43B0 | Reserved | — | — | — |
| 0x2_43B4 | RBPTR6*—RxB _D pointer for ring 6 | R/W | All zeros | 19.5.3.3.10/19-61 |
| 0x2_43B8 | Reserved | — | — | — |
| 0x2_43BC | RBPTR7*—RxB _D pointer for ring 7 | R/W | All zeros | 19.5.3.3.10/19-61 |
| 0x2_43C0–0x2_44400 | Reserved | — | — | — |
| 0x2_4404 | RBASE0—RxB _D base address of ring 0 | R/W | All zeros | 19.5.3.3.11/19-62 |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|----------------------------|---|---------------------|-------------|-----------------------------------|
| 0x2_4408 | Reserved | — | — | — |
| 0x2_440C | RBASE1*—RxBD base address of ring 1 | R/W | All zeros | 19.5.3.3.11/19-62 |
| 0x2_4410 | Reserved | — | — | — |
| 0x2_4414 | RBASE2*—RxBD base address of ring 2 | R/W | All zeros | 19.5.3.3.11/19-62 |
| 0x2_4418 | Reserved | — | — | — |
| 0x2_441C | RBASE3*—RxBD base address of ring 3 | R/W | All zeros | 19.5.3.3.11/19-62 |
| 0x2_4420 | Reserved | — | — | — |
| 0x2_4424 | RBASE4*—RxBD base address of ring 4 | R/W | All zeros | 19.5.3.3.11/19-62 |
| 0x2_4428 | Reserved | — | — | — |
| 0x2_442C | RBASE5*—RxBD base address of ring 5 | R/W | All zeros | 19.5.3.3.11/19-62 |
| 0x2_4430 | Reserved | — | — | — |
| 0x2_4434 | RBASE6*—RxBD base address of ring 6 | R/W | All zeros | 19.5.3.3.11/19-62 |
| 0x2_4438 | Reserved | — | — | — |
| 0x2_443C | RBASE7*—RxBD base address of ring 7 | R/W | All zeros | 19.5.3.3.11/19-62 |
| 0x2_4440– 0x2_44BC | Reserved | — | — | — |
| 0x2_44C0 | TMR_RXTS_H* - Rx timer time stamp register high | R/W | All zeros | 19.5.3.3.12/19-62 |
| 0x2_44C4 | TMR_RXTS_L* - Rx timer time stamp register low | R/W | All zeros | 19.5.3.3.12/19-62 |
| 0x2_44C8– 0x2_44FC | Reserved | — | — | — |
| eTSEC MAC Registers | | | | |
| 0x2_4500 | MACCFG1—MAC configuration register 1 | R/W | All zeros | 19.5.3.5.1/19-66 |
| 0x2_4504 | MACCFG2—MAC configuration register 2 | R/W | 0x0000_7000 | 19.5.3.5.2/19-67 |
| 0x2_4508 | IPGIFG—Inter-packet/inter-frame gap register | R/W | 0x4060_5060 | 19.5.3.5.3/19-69 |
| 0x2_450C | HAFDUP—Half-duplex control | R/W | 0x00A1_F037 | 19.5.3.5.4/19-70 |
| 0x2_4510 | MAXFRM—Maximum frame length | R/W | 0x0000_0600 | 19.5.3.5.5/19-71 |
| 0x2_4514– 0x2_451C | Reserved | — | — | — |
| 0x2_4520 | MIIMCFG—MII management configuration | R/W | 0x0000_0007 | 19.5.3.5.6/19-71 |
| 0x2_4524 | MIIMCOM—MII management command | R/W | All zeros | 19.5.3.5.7/19-72 |
| 0x2_4528 | MIIMADD—MII management address | R/W | All zeros | 19.5.3.5.8/19-73 |
| 0x2_452C | MIIMCON—MII management control | WO | All zeros | 19.5.3.5.9/19-74 |
| 0x2_4530 | MIIMSTAT—MII management status | R | All zeros | 19.5.3.5.10/19-74 |
| 0x2_4534 | MIIMIND—MII management indicator | R | All zeros | 19.5.3.5.11/19-75 |
| 0x2_4538 | Reserved | — | — | — |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page | |
|--|--|---------------------|-----------|--|--|
| 0x2_453C | IFSTAT—Interface status | R | All zeros | 19.5.3.5.12/19-75 | |
| 0x2_4540 | MACSTNADDR1—MAC station address register 1 | R/W | All zeros | 19.5.3.5.13/19-76 | |
| 0x2_4544 | MACSTNADDR2—MAC station address register 2 | R/W | All zeros | 19.5.3.5.14/19-77 | |
| 0x2_4548 | MAC01ADDR1*—MAC exact match address 1, part 1 | R/W | All zeros | 19.5.3.5.15/19-77 19.5.3.5.16/19-78 | |
| 0x2_454C | MAC01ADDR2*—MAC exact match address 1, part 2 | R/W | All zeros | | |
| 0x2_4550 | MAC02ADDR1*—MAC exact match address 2, part 1 | R/W | All zeros | | |
| 0x2_4554 | MAC02ADDR2*—MAC exact match address 2, part 2 | R/W | All zeros | | |
| 0x2_4558 | MAC03ADDR1*—MAC exact match address 3, part 1 | R/W | All zeros | | |
| 0x2_455C | MAC03ADDR2*—MAC exact match address 3, part 2 | R/W | All zeros | | |
| 0x2_4560 | MAC04ADDR1*—MAC exact match address 4, part 1 | R/W | All zeros | | |
| 0x2_4564 | MAC04ADDR2*—MAC exact match address 4, part 2 | R/W | All zeros | | |
| 0x2_4568 | MAC05ADDR1*—MAC exact match address 5, part 1 | R/W | All zeros | | |
| 0x2_456C | MAC05ADDR2*—MAC exact match address 5, part 2 | R/W | All zeros | | |
| 0x2_4570 | MAC06ADDR1*—MAC exact match address 6, part 1 | R/W | All zeros | | 19.5.3.5.15/19-77 19.5.3.5.16/19-78 |
| 0x2_4574 | MAC06ADDR2*—MAC exact match address 6, part 2 | R/W | All zeros | | |
| 0x2_4578 | MAC07ADDR1*—MAC exact match address 7, part 1 | R/W | All zeros | | |
| 0x2_457C | MAC07ADDR2*—MAC exact match address 7, part 2 | R/W | All zeros | | |
| 0x2_4580 | MAC08ADDR1*—MAC exact match address 8, part 1 | R/W | All zeros | | |
| 0x2_4584 | MAC08ADDR2*—MAC exact match address 8, part 2 | R/W | All zeros | | |
| 0x2_4588 | MAC09ADDR1*—MAC exact match address 9, part 1 | R/W | All zeros | | |
| 0x2_458C | MAC09ADDR2*—MAC exact match address 9, part 2 | R/W | All zeros | | |
| 0x2_4590 | MAC10ADDR1*—MAC exact match address 10, part 1 | R/W | All zeros | | |
| 0x2_4594 | MAC10ADDR2*—MAC exact match address 10, part 2 | R/W | All zeros | | |
| 0x2_4598 | MAC11ADDR1*—MAC exact match address 11, part 1 | R/W | All zeros | | |
| 0x2_459C | MAC11ADDR2*—MAC exact match address 11, part 2 | R/W | All zeros | | |
| 0x2_45A0 | MAC12ADDR1*—MAC exact match address 12, part 1 | R/W | All zeros | | |
| 0x2_45A4 | MAC12ADDR2*—MAC exact match address 12, part 2 | R/W | All zeros | | |
| 0x2_45A8 | MAC13ADDR1*—MAC exact match address 13, part 1 | R/W | All zeros | | |
| 0x2_45AC | MAC13ADDR2*—MAC exact match address 13, part 2 | R/W | All zeros | | |
| 0x2_45B0 | MAC14ADDR1*—MAC exact match address 14, part 1 | R/W | All zeros | | |
| 0x2_45B4 | MAC14ADDR2*—MAC exact match address 14, part 2 | R/W | All zeros | | |
| 0x2_45B8 | MAC15ADDR1*—MAC exact match address 15, part 1 | R/W | All zeros | | |
| 0x2_45BC | MAC15ADDR2*—MAC exact match address 15, part 2 | R/W | All zeros | | |
| 0x2_45C0– 0x2_467C | Reserved | — | — | — | |
| eTSEC Transmit and Receive Counters | | | | | |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|--------------------------------|---|---------------------|-----------|-----------------------------------|
| 0x2_4680 | TR64—Transmit and receive 64-byte frame counter | R/W | All zeros | 19.5.3.6.1/19-79 |
| 0x2_4684 | TR127—Transmit and receive 65- to 127-byte frame counter | R/W | All zeros | 19.5.3.6.2/19-80 |
| 0x2_4688 | TR255—Transmit and receive 128- to 255-byte frame counter | R/W | All zeros | 19.5.3.6.3/19-80 |
| 0x2_468C | TR511—Transmit and receive 256- to 511-byte frame counter | R/W | All zeros | 19.5.3.6.4/19-81 |
| 0x2_4690 | TR1K—Transmit and receive 512- to 1023-byte frame counter | R/W | All zeros | 19.5.3.6.5/19-81 |
| 0x2_4694 | TRMAX—Transmit and receive 1024- to 1518-byte frame counter | R/W | All zeros | 19.5.3.6.6/19-82 |
| 0x2_4698 | TRMGV—Transmit and receive 1519- to 1522-byte good VLAN frame count | R/W | All zeros | 19.5.3.6.7/19-82 |
| eTSEC Receive Counters | | | | |
| 0x2_469C | RBYT—Receive byte counter | R/W | All zeros | 19.5.3.6.8/19-83 |
| 0x2_46A0 | RPKT—Receive packet counter | R/W | All zeros | 19.5.3.6.9/19-83 |
| 0x2_46A4 | RFCS—Receive FCS error counter | R/W | All zeros | 19.5.3.6.10/19-83 |
| 0x2_46A8 | RMCA—Receive multicast packet counter | R/W | All zeros | 19.5.3.6.11/19-84 |
| 0x2_46AC | RBCA—Receive broadcast packet counter | R/W | All zeros | 19.5.3.6.12/19-84 |
| 0x2_46B0 | RXCF—Receive control frame packet counter | R/W | All zeros | 19.5.3.6.13/19-85 |
| 0x2_46B4 | RXPf—Receive PAUSE frame packet counter | R/W | All zeros | 19.5.3.6.14/19-85 |
| 0x2_46B8 | RXUO—Receive unknown OP code counter | R/W | All zeros | 19.5.3.6.15/19-86 |
| 0x2_46BC | RALN—Receive alignment error counter | R/W | All zeros | 19.5.3.6.16/19-86 |
| 0x2_46C0 | RFLR—Receive frame length error counter | R/W | All zeros | 19.5.3.6.17/19-87 |
| 0x2_46C4 | RCDE—Receive code error counter | R/W | All zeros | 19.5.3.6.18/19-87 |
| 0x2_46C8 | RCSE—Receive carrier sense error counter | R/W | All zeros | 19.5.3.6.19/19-88 |
| 0x2_46CC | RUND—Receive undersize packet counter | R/W | All zeros | 19.5.3.6.20/19-88 |
| 0x2_46D0 | ROVR—Receive oversize packet counter | R/W | All zeros | 19.5.3.6.21/19-89 |
| 0x2_46D4 | RFRG—Receive fragments counter | R/W | All zeros | 19.5.3.6.22/19-89 |
| 0x2_46D8 | RJBR—Receive jabber counter | R/W | All zeros | 19.5.3.6.23/19-90 |
| 0x2_46DC | RDRP—Receive drop counter | R/W | All zeros | 19.5.3.6.24/19-90 |
| eTSEC Transmit Counters | | | | |
| 0x2_46E0 | TBYT—Transmit byte counter | R/W | All zeros | 19.5.3.6.25/19-91 |
| 0x2_46E4 | TPKT—Transmit packet counter | R/W | All zeros | 19.5.3.6.26/19-91 |
| 0x2_46E8 | TMCA—Transmit multicast packet counter | R/W | All zeros | 19.5.3.6.27/19-92 |
| 0x2_46EC | TBCA—Transmit broadcast packet counter | R/W | All zeros | 19.5.3.6.28/19-92 |
| 0x2_46F0 | TXPF—Transmit PAUSE control frame counter | R/W | All zeros | 19.5.3.6.29/19-93 |
| 0x2_46F4 | TDFR—Transmit deferral packet counter | R/W | All zeros | 19.5.3.6.30/19-93 |
| 0x2_46F8 | TEDF—Transmit excessive deferral packet counter | R/W | All zeros | 19.5.3.6.31/19-94 |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|---|--|---------------------|-------------|------------------------------------|
| 0x2_46FC | TSCL—Transmit single collision packet counter | R/W | All zeros | 19.5.3.6.32/19-94 |
| 0x2_4700 | TMCL—Transmit multiple collision packet counter | R/W | All zeros | 19.5.3.6.33/19-95 |
| 0x2_4704 | TLCL—Transmit late collision packet counter | R/W | All zeros | 19.5.3.6.34/19-95 |
| 0x2_4708 | TXCL—Transmit excessive collision packet counter | R/W | All zeros | 19.5.3.6.35/19-96 |
| 0x2_470C | TNCL—Transmit total collision counter | R/W | All zeros | 19.5.3.6.36/19-96 |
| 0x2_4710 | Reserved | — | — | — |
| 0x2_4714 | TDRP—Transmit drop frame counter | R/W | All zeros | 19.5.3.6.37/19-97 |
| 0x2_4718 | TJBR—Transmit jabber frame counter | R/W | All zeros | 19.5.3.6.38/19-97 |
| 0x2_471C | TFCS—Transmit FCS error counter | R/W | All zeros | 19.5.3.6.39/19-98 |
| 0x2_4720 | TXCF—Transmit control frame counter | R/W | All zeros | 19.5.3.6.40/19-98 |
| 0x2_4724 | TOVR—Transmit oversize frame counter | R/W | All zeros | 19.5.3.6.41/19-99 |
| 0x2_4728 | TUND—Transmit undersize frame counter | R/W | All zeros | 19.5.3.6.42/19-99 |
| 0x2_472C | TFRG—Transmit fragments frame counter | R/W | All zeros | 19.5.3.6.43/19-100 |
| eTSEC Counter Control and TOE Statistics Registers | | | | |
| 0x2_4730 | CAR1—Carry register one register ³ | w1c | All zeros | 19.5.3.6.44/19-100 |
| 0x2_4734 | CAR2—Carry register two register ³ | w1c | All zeros | 19.5.3.6.45/19-102 |
| 0x2_4738 | CAM1—Carry register one mask register | R/W | 0xFE03_FFFF | 19.5.3.6.46/19-103 |
| 0x2_473C | CAM2—Carry register two mask register | R/W | 0x000F_FFFD | 19.5.3.6.47/19-104 |
| 0x2_4740 | RREJ*—Receive filer rejected packet counter | R/W | All zeros | 19.5.3.6.48/19-105 |
| 0x2_4744– 0x2_47FC | Reserved | — | — | — |
| Hash Function Registers | | | | |
| 0x2_4800 | IGADDR0—Individual/group address register 0 | R/W | All zeros | 19.5.3.7.1/19-106 |
| 0x2_4804 | IGADDR1—Individual/group address register 1 | R/W | All zeros | |
| 0x2_4808 | IGADDR2—Individual/group address register 2 | R/W | All zeros | |
| 0x2_480C | IGADDR3—Individual/group address register 3 | R/W | All zeros | |
| 0x2_4810 | IGADDR4—Individual/group address register 4 | R/W | All zeros | |
| 0x2_4814 | IGADDR5—Individual/group address register 5 | R/W | All zeros | |
| 0x2_4818 | IGADDR6—Individual/group address register 6 | R/W | All zeros | |
| 0x2_481C | IGADDR7—Individual/group address register 7 | R/W | All zeros | |
| 0x2_4820– 0x2_487C | Reserved | — | — | — |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|--|---|---------------------|-----------|-------------------|
| 0x2_4880 | GADDR0—Group address register 0 | R/W | All zeros | 19.5.3.7.2/19-106 |
| 0x2_4884 | GADDR1—Group address register 1 | R/W | All zeros | |
| 0x2_4888 | GADDR2—Group address register 2 | R/W | All zeros | |
| 0x2_488C | GADDR3—Group address register 3 | R/W | All zeros | |
| 0x2_4890 | GADDR4—Group address register 4 | R/W | All zeros | |
| 0x2_4894 | GADDR5—Group address register 5 | R/W | All zeros | |
| 0x2_4898 | GADDR6—Group address register 6 | R/W | All zeros | |
| 0x2_489C | GADDR7—Group address register 7 | R/W | All zeros | |
| 0x2_48A0–0x2_4AFC | Reserved | — | — | — |
| eTSEC DMA Attribute Registers | | | | |
| 0x2_4B00–0x2_4BF4 | Reserved | — | — | — |
| 0x2_4BF8 | ATTR—Attribute register | R/W | All zeros | 19.5.3.8.1/19-107 |
| eTSEC Lossless Flow Control Registers | | | | |
| 0x2_4C00 | RQPRM0*—Receive Queue Parameters register 0 | R/W | All zeros | 19.5.3.9.1/19-108 |
| 0x2_4C04 | RQPRM1*—Receive Queue Parameters register 1 | R/W | All zeros | |
| 0x2_4C08 | RQPRM2*—Receive Queue Parameters register 2 | R/W | All zeros | |
| 0x2_4C0C | RQPRM3*—Receive Queue Parameters register 3 | R/W | All zeros | |
| 0x2_4C10 | RQPRM4*—Receive Queue Parameters register 4 | R/W | All zeros | |
| 0x2_4C14 | RQPRM5*—Receive Queue Parameters register 5 | R/W | All zeros | |
| 0x2_4C18 | RQPRM6*—Receive Queue Parameters register 6 | R/W | All zeros | |
| 0x2_4C1C | RQPRM7*—Receive Queue Parameters register 7 | R/W | All zeros | |
| 0x2_4C20–0x2_4C40 | Reserved | — | — | — |
| 0x2_4C44 | RFBPTR0*—Last Free RxBD pointer for ring 0 | R/W | All zeros | 19.5.3.9.2/19-109 |
| 0x2_4C48 | Reserved | — | — | — |
| 0x2_4C4C | RFBPTR1*—Last Free RxBD pointer for ring 1 | R/W | All zeros | 19.5.3.9.2/19-109 |
| 0x2_4C50 | Reserved | — | — | — |
| 0x2_4C54 | RFBPTR2*—Last Free RxBD pointer for ring 2 | R/W | All zeros | 19.5.3.9.2/19-109 |
| 0x2_4C58 | Reserved | — | — | — |
| 0x2_4C5C | RFBPTR3*—Last Free RxBD pointer for ring 3 | R/W | All zeros | 19.5.3.9.2/19-109 |
| 0x2_4C60 | Reserved | — | — | — |
| 0x2_4C64 | RFBPTR4*—Last Free RxBD pointer for ring 4 | R/W | All zeros | 19.5.3.9.2/19-109 |
| 0x2_4C68 | Reserved | — | — | — |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|-------------------------------------|---|---------------------|-------------|---------------------|
| 0x2_4C6C | RFBPTR5*—Last Free RxBD pointer for ring 5 | R/W | All zeros | 19.5.3.9.2/19-109 |
| 0x2_4C70 | Reserved | — | — | — |
| 0x2_4C74 | RFBPTR6*—Last Free RxBD pointer for ring 6 | R/W | All zeros | 19.5.3.9.2/19-109 |
| 0x2_4C78 | Reserved | — | — | — |
| 0x2_4C7C | RFBPTR7*—Last Free RxBD pointer for ring 7 | R/W | All zeros | 19.5.3.9.2/19-109 |
| eTSEC Future Expansion Space | | | | |
| 0x2_4CC0 – 0x2_4D94 | Reserved | — | — | — |
| eTSEC IEEE 1588 Registers | | | | |
| 0x2_4E00 | TMR_CTRL* - Timer control register | R/W | 0x0001_0000 | 19.5.3.10.1/19-110 |
| 0x2_4E04 | TMR_TEVENT* - time stamp event register | w1c | All zeros | 19.5.3.10.2/19-111 |
| 0x2_4E08 | TMR_TEMASK* - Timer event mask register | R/W | All zeros | 19.5.3.10.3/19-113 |
| 0x2_4E0C | TMR_PEVENT* - time stamp event register | R/W | All zeros | 19.5.3.10.4/19-113 |
| 0x2_4E10 | TMR_PEMASK* - Timer event mask register | R/W | All zeros | 19.5.3.10.5/19-114 |
| 0x2_4E14 | TMR_STAT* - time stamp status register | R/W | All zeros | 19.5.3.10.6/19-115 |
| 0x2_4E18 | TMR_CNT_H* - timer counter high register | R/W | All zeros | 19.5.3.10.7/19-115 |
| 0x2_4E1C | TMR_CNT_L* - timer counter low register | R/W | All zeros | 19.5.3.10.7/19-115 |
| 0x2_4E20 | TMR_ADD* - Timer drift compensation addend register | R/W | All zeros | 19.5.3.10.8/19-116 |
| 0x2_4E24 | TMR_ACC* - Timer accumulator register | R/W | All zeros | 19.5.3.10.9/19-117 |
| 0x2_4E28 | TMR_PRSC* -Timer prescale | R/W | 0x0000_0002 | 19.5.3.10.10/19-117 |
| 0x2_4E2C | Reserved | — | — | — |
| 0x2_4E30 | TMROFF_H* - Timer offset high | R/W | All zeros | 19.5.3.10.11/19-118 |
| 0x2_4E34 | TMROFF_L* - Timer offset low | R/W | All zeros | 19.5.3.10.11/19-118 |
| 0x2_4E40 | TMR_ALARM1_H* - Timer alarm 1 high register | R/W | 0xFFFF_FFFF | 19.5.3.10.12/19-118 |
| 0x2_4E44 | TMR_ALARM1_L* - Timer alarm 1 high register | R/W | 0xFFFF_FFFF | |
| 0x2_4E48 | TMR_ALARM2_H* - Timer alarm 2 high register | R/W | 0xFFFF_FFFF | |
| 0x2_4E4C | TMR_ALARM2_L* - Timer alarm 2 high register | R/W | 0xFFFF_FFFF | |
| 0x2_4E50– 0x2_4E7C | Reserved | — | — | — |
| 0x2_4E80 | TMR_FIPER1* - Timer fixed period interval | R/W | 0xFFFF_FFFF | 19.5.3.10.13/19-119 |
| 0x2_4E84 | TMR_FIPER2* - Timer fixed period interval | R/W | 0xFFFF_FFFF | |
| 0x2_4E88 | TMR_FIPER*3 - Timer fixed period interval | R/W | 0xFFFF_FFFF | |

Table 19-4. Module Memory Map (continued)

| eTSEC1 Offset | Name ¹ | Access ² | Reset | Section/Page |
|---------------------------|---|---------------------|-----------|---------------------|
| 0x2_4EA0 | TMR_ETTS1_H* - Time stamp of general purpose external trigger | R/W | All zeros | 19.5.3.10.14/19-120 |
| 0x2_4EA4 | TMR_ETTS1_L* - Time stamp of general purpose external trigger | R/W | All zeros | |
| 0x2_4EA8 | TMR_ETTS2_H* - Time stamp of general purpose external trigger | R/W | All zeros | |
| 0x2_4EAC | TMR_ETTS2_L* - Time stamp of general purpose external trigger | R/W | All zeros | |
| 0x2_4EB0 – 0x2_4FFF | Reserved | — | — | |
| Other eTSECs | | | | |
| 0x2_5000– 0x2_5FFF | eTSEC2 REGISTERS ⁴ | | | |

¹ Registers denoted * are new to the enhanced TSEC and not supported by PowerQUICC II Pro TSECs.

² Key: R = read only, WO = write only, R/W = read and write, LH = latches high, SC = self-clearing.

³ Cleared on read.

⁴ eTSEC2 has the same memory-mapped registers that are described for eTSEC1 from 0x 2_4000 to 0x2_4FFF, except the offsets are from 0x 2_5000 to 0x2_5FFF.

19.5.3 Memory-Mapped Register Descriptions

This section provides a detailed description of all the eTSEC registers. Because all of the eTSEC registers are 32 bits wide, only 32-bit register accesses are supported.

19.5.3.1 eTSEC General Control and Status Registers

This section describes general control and status registers used for both transmitting and receiving Ethernet frames. All of the registers are 32 bits wide.

19.5.3.1.1 Controller ID Register (TSEC_ID)

The controller ID register (TSEC_ID) is a read-only register. The TSEC_ID register is used to identify the eTSEC block and revision.




Figure 19-2. TSEC_ID Register

Table 19-5 describes the fields of the TSEC_ID register.

Table 19-5. TSEC_ID Field Descriptions

| Bits | Name | Description |
|-------|-------------|--|
| 0–15 | TSEC_ID | Value identifying the eTSEC (10/100/1000 Ethernet MAC). 0124 Unique identifier for eTSEC with 8 Rx and 8 Tx BD rings. |
| 16–23 | TSEC_REV_MJ | Value identifies the major revision of the eTSEC. 00 Initial revision |
| 24–31 | TSEC_REV_MN | Value identifies the minor revision of the eTSEC. 05 Initial revision |

19.5.3.1.2 Controller ID Register (TSEC_ID2)

The controller ID register (TSEC_ID2) is a read-only register. The TSEC_ID2 register is used to identify the eTSEC block configuration.




Figure 19-3. TSEC_ID2 Register

Table 19-6 describes the fields of the TSEC_ID2 register.

Table 19-6. TSEC_ID2 Field Descriptions

| Bits | Name | Description |
|-------|----------|--|
| 0–7 | — | Reserved |
| 8–15 | TSEC_INT | Interface mode support. See Table 19-7 for settings. |
| 16–23 | — | Reserved |
| 24–31 | TSEC_CFG | Value identifies configuration options of the eTSEC. 00 eTSEC multiple ring, Rx TOE, Filer and Tx TOE supports are off F0 eTSEC multiple ring, Rx TOE, Filer and Tx TOE supports are on 30 eTSEC multiple ring support is OFF and Rx TOE, Filer and Tx TOE supports are on 50 eTSEC multiple ring and filer supports are OFF and Rx TOE and Tx TOE supports are on |

Table 19-7 describes the field settings for TSEC_ID2[TSEC_INT].

Table 19-7. TSEC_ID2[TSEC_INT] Field Settings

| Bit | Mode |
|-----|-------------------------------|
| 8 | 0 1588 protocol not supported |
| | 1 1588 protocol supported |
| 9 | 0 SGMII not supported |
| | 1 SGMII supported |

Table 19-7. TSEC_ID2[TSEC_INT] Field Settings (continued)

| | |
|-------|--|
| 10 | 0 Ethernet mode not supported 1 Ethernet mode supported |
| 11–13 | Reserved |
| 14 | 0 Can be configured to run in Ethernet normal/full mode 1 Ethernet normal/full mode off |
| 15 | 0 Can be configured to run in Ethernet reduced mode 1 Ethernet reduced mode off |

19.5.3.1.3 Interrupt Event Register (IEVENT)

Interrupt events cause bits in the IEVENT register to be set. Software may poll this register at any time to check for pending interrupts. If an event occurs and its corresponding enable bit is set in the interrupt mask register (IMASK), the event also causes a hardware interrupt at the PIC. A bit in the interrupt event register is cleared by writing a 1 to that bit position. A write of 0 has no effect.

Each eTSEC can issue three kinds of hardware interrupt to the PIC:

1. Transmit data frame interrupts—Issued whenever bits TXB or TXF of IEVENT are set to 1 and either transmit interrupt coalescing is disabled or the interrupt coalescing thresholds have been met for TXF. To negate this hardware interrupt, software must clear both TXB and TXF bits.
2. Receive data frame interrupts—Issued whenever bits RXB or RXF of IEVENT are set to 1 and either receive interrupt coalescing is disabled or the interrupt coalescing thresholds have been met for RXF. To negate this hardware interrupt, software must clear both RXB and RXF bits.
3. Error, diagnostic, and special interrupts—Issued whenever bits MAG, GTSC, GRSC, TXC, RXC, BABR, BAPT, LC, CRL, FGPI, FIR, FIQ, DPE, PERR, EBERR, TXE, XFUN, BSY, MSRO, MMRD, or MMRW of IEVENT are set to 1. Software must clear all of these bits to negate an error/diagnostic/special hardware interrupt.
 - Magic Packet reception event is: MAG
 - Operational diagnostics are events on: GTSC, GRSC, TXC, and RXC
 - Interrupts resulting from errors/problems detected in the network or transceiver are: BABR, BAPT, LC, and CRL
 - Interrupts resulting from internal or combination errors are: FIR, FIQ, DPE, PERR, EBERR, TXE, XFUN, and BSY
 - Special function interrupts are: FGPI, MSRO, MMRD, and MMRW

Some of the error interrupts are independently counted in the MIB block counters. Software may choose to mask off these interrupts because these errors are visible to network management through the MIB counters.

Figure 19-4 describes the definition for the IEVENT register.

Offset eTSEC1:0x2_4010; eTSEC2: 0x2_5010

Access: w1c

| | | | | | | | | | | | | | | | | | |
|-------|-----------|-----|-----|-------|------|------|------|------|-----|-----|-----|------|------|-----|-----|------|------|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | |
| R | BABR | RXC | BSY | EBERR | — | MSRO | GTSC | BABT | TXC | TXE | TXB | TXF | — | LC | CRL | XFUN | |
| W | w1c | w1c | w1c | w1c | — | w1c | w1c | w1c | w1c | w1c | w1c | w1c | — | w1c | w1c | w1c | |
| Reset | All zeros | | | | | | | | | | | | | | | | |
| | 16 | 17 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | | |
| R | RXB | — | | | MMRD | MMWR | GRSC | RXF | — | | | FGPI | FIR | FIQ | DPE | PERR | |
| W | w1c | — | | | w1c | w1c | w1c | w1c | — | | | w1c | w1c | w1c | w1c | w1c | |
| R | RXB | — | | | MAG | MMRD | MMWR | GRSC | RXF | — | | | FGPI | FIR | FIQ | DPE | PERR |
| W | w1c | — | | | w1c | w1c | w1c | w1c | w1c | — | | | w1c | w1c | w1c | w1c | w1c |
| R | RXB | — | | | MAG | MMRD | MMWR | GRSC | RXF | — | | | FIR | FIQ | DPE | PERR | |
| W | w1c | — | | | w1c | w1c | w1c | w1c | w1c | — | | | w1c | w1c | w1c | w1c | w1c |
| Reset | All zeros | | | | | | | | | | | | | | | | |

Figure 19-4. IEVENT Register Definition

Table 19-8 describes the fields of the IEVENT register.

Table 19-8. IEVENT Field Descriptions

| Bits | Name | Description |
|------|-------|---|
| 0 | BABR | Babbling receive error. This bit indicates that a frame was received with length in excess of the MAC's maximum frame length register while MACCFG2[Huge Frame] is set. 0 Excessive frame not received. 1 Excessive frame received. |
| 1 | RXC | Receive control interrupt. A control frame was received while MACCFG1[Rx_Flow] is set. As soon as the transmitter finishes sending the current frame, a pause operation is performed. 0 Control frame not received. 1 Control frame received. |
| 2 | BSY | Busy condition interrupt. Indicates that a frame was received and discarded due to a lack of buffers. 0 No frame received and discarded. 1 Frame received and discarded. |
| 3 | EBERR | Internal bus error. This bit indicates that a system bus error occurred while a DMA transaction was underway. As a result, transferred data is expected to be partially or completely invalid. 0 No system bus error occurred. 1 System bus error occurred. |
| 4 | — | Reserved |
| 5 | MSRO | MIB counter overflow. This interrupt is asserted if the count for one of the MIB counters has exceeded the size of its register. 0 MIB count not exceeding its register size. 1 MIB count exceeds its register size. |

Table 19-8. IEVENT Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 6 | GTSC | Graceful transmit stop complete. This interrupt is asserted for one of two reasons. Graceful stop means that the transmitter is put into a pause state after completion of the frame currently being transmitted. <ul style="list-style-type: none"> • A graceful stop, which was initiated by setting DMACTRL[GTS], is now complete. • A transmission of a flow control PAUSE frame, which was initiated by setting TCTRL[TFC_PAUSE], is now complete. 0 No graceful stop interrupt. 1 Graceful stop requested. |
| 7 | BABT | Babbling transmit error. This bit indicates that the transmitted frame length has exceeded the value in the MAC's maximum frame length register and MACCFG2[Huge Frame] is cleared. Frame truncation occurs when this condition occurs. 0 Transmitted frame length not exceeding maximum frame length. 1 Transmitted frame length exceeding maximum frame length when MACCFG2[Huge Frame] = 0. |
| 8 | TXC | Transmit control interrupt. This bit indicates that a control frame was transmitted. 0 Control frame not transmitted. 1 Control frame transmitted. |
| 9 | TXE | Transmit error. This bit indicates that an error occurred on the transmitted channel that has caused TSTAT[THLT] to be set by the eTSEC. This bit is set whenever any transmit error occurs that causes the transmitter to halt (EBERR, LC, CRL, XFUN). 0 No transmit channel error occurred. 1 Transmit channel error occurred. |
| 10 | TXB | Transmit buffer. This bit indicates that a transmit buffer descriptor was updated whose I (interrupt) bit was set in its status word and was not the last buffer descriptor of the frame. 0 No transmit buffer descriptor updated. 1 Transmit buffer descriptor updated. |
| 11 | TXF | Transmit frame interrupt. This bit indicates that a frame was transmitted and that the last corresponding transmit buffer descriptor (TxBD) was updated. This only occurs if the I (interrupt) bit in the status word of the buffer descriptor is set. The specific transmit queue that was updated has its TXF bit set in TSTAT. 0 No frame transmitted/TxBD not updated. 1 Frame transmitted/TxBD updated. |
| 12 | — | Reserved |
| 13 | LC | Late collision. This bit indicates that a collision occurred beyond the collision window (slot time) in half-duplex mode. The frame is truncated with a bad CRC and the remainder of the frame is discarded. 0 No late collision occurred. 1 Late collision occurred. |
| 14 | CRL | Collision retry limit. This bit indicates that the number of successive transmission collisions has exceeded the MAC's half-duplex register's retransmission maximum count (HAFDUP[Retransmission Maximum]). The frame is discarded without being transmitted and the queue halts (TSTAT[THLT _n] set to 1). This only occurs while in half-duplex mode. 0 Successive transmission collisions do not exceed maximum. 1 Successive transmission collisions exceed maximum. |
| 15 | XFUN | Transmit FIFO underrun. This bit indicates that the transmit FIFO became empty before the complete frame was transmitted. 0 Transmit FIFO not underrun. 1 Transmit FIFO underrun. |
| 16 | RXB | Receive buffer. This bit indicates that a receive buffer descriptor was updated which had the I (Interrupt) bit set in its status word and was not the last buffer descriptor of the frame. 0 Receive buffer descriptor not updated. 1 Receiver buffer descriptor updated. |

Table 19-8. IEVENT Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 17–19 | — | Reserved |
| 20 | MAG | Magic Packet detected when the eTSEC is in Magic Packet detection mode (MACCFG2[MPEN] = 1). 0 No Magic Packet received, or Magic Packet mode was not enabled. 1 A Magic Packet was received while in Magic Packet mode. MACCFG2[MPEN] is also cleared upon receiving the Magic Packet. |
| 21 | MMRD | MII management read completion 0 MII management read not issued or in process. 1 MII management read completed that was initiated by a user through the MII Scan or Read cycle command. |
| 22 | MMWR | MII management write completion 0 MII management write not issued or in process. 1 MII management write completed that was initiated by a user write to the MIIMCON register. |
| 23 | GRSC | Graceful receive stop complete. This interrupt is asserted if a graceful receive stop is completed. It allows the user to know if the system has completed the stop and it is safe to write to receive registers (status, control or configuration registers) that are used by the system during normal operation. 0 Graceful stop not completed. 1 Graceful stop completed. |
| 24 | RXF | Receive frame interrupt. This bit indicates that a frame was received and the last receive buffer descriptor (RxBD) in that frame was updated. This occurs either if the I (interrupt) bit in the buffer descriptor status word is set, or an overrun error occurs. The specific receive queue that was updated has its RXF bit set in RSTAT. 0 Frame not received. 1 Frame received. |
| 25–26 | — | Reserved |
| 27 | FGPI | Filer generated general purpose interrupt on a set of filer rule match. This bit will be set upon reception of a frame that matches a GPI rule sequence that is specified in the filer. It is synchronized with the setting of RXF. 0 No filer generated interrupt has occurred. 1 The filer has accepted a frame via a matching rule that the RQFCR[GPI] bit set. |
| 28 | FIR | The receive queue filer result is invalid, either because not enough time between frames was available to find a matching rule, or no entry in the filer table could be matched. 0 Receive queue filer reached a definite result; however, bit FIQ may still be set if a frame was filed to a disabled RxBD ring. 1 Receive queue filer was unable to reach a definite result. In this case, bit FIQ is also set if no entry in the filer table could provide a rule match. |
| 29 | FIQ | Filed frame to invalid receive queue. This bit indicates that either the receive queue filer chose to DMA a received frame to a disabled RxBD ring, or that no rule in the filer table could be matched. 0 Received frames filed to valid queues or rejected. Note that a frame may be rejected if the filer has insufficient time to reach a conclusive result between frames, in which case bit FIR is set. 1 Received frames filed to RxBD rings that are not enabled. The frame is discarded. If bit FIR is also set this indicates that the filer exhausted all of its table entries without a rule match. |

Table 19-8. IEVENT Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 30 | DPE | Internal data parity error. This bit indicates that the eTSEC has detected a parity error on its stored data, which is likely to compromise the validity of recently transferred frames. 0 No parity errors detected. 1 Data held in the FIFO or filter arrays is expected to be corrupted due to a parity error. |
| 31 | PERR | Receive frame parse error for TCP/IP off-load. This bit indicates that a received frame could not be parsed unambiguously, due to encapsulated header type fields contradicting each other. 0 Received frame parsed successfully. 1 Received frame parse revealed header inconsistencies. |

19.5.3.1.4 Interrupt Mask Register (IMASK)

The interrupt mask register provides control over which possible interrupt events in the IEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this register are R/W and cleared upon a hardware reset. If the corresponding bits in both the IEVENT and IMASK registers are set, the PIC receives an interrupt (for each eTSEC these are grouped into transmit, receive, and error/diagnostic interrupts). The interrupt signal remains asserted until either the IEVENT bit is cleared, by writing a 1 to it, or by writing a 0 to the corresponding IMASK bit.

Figure 19-5 describes the IMASK register.

Offset eTSEC1:0x2_4014; eTSEC2:0x2_5014

Access: Read/Write




Figure 19-5. IMASK Register Definition

Table 19-9 describes the fields of the IMASK register.

Table 19-9. IMASK Field Descriptions

| Bits | Name | Description |
|-------|---------|--|
| 0 | BREN | Babbling receiver interrupt enable |
| 1 | RXCEN | Receive control interrupt enable |
| 2 | BSYEN | Busy interrupt enable |
| 3 | EBERREN | Ethernet controller bus error enable |
| 4 | — | Reserved |
| 5 | MSROEN | MIB counter overflow interrupt enable |
| 6 | GTSCEN | Graceful transmit stop complete interrupt enable |
| 7 | BTEN | Babbling transmitter interrupt enable |
| 8 | TXCEN | Transmit control interrupt enable |
| 9 | TXEEN | Transmit error interrupt enable |
| 10 | TXBEN | Transmit buffer interrupt enable |
| 11 | TXFEN | Transmit frame interrupt enable |
| 12 | — | Reserved |
| 13 | LCEN | Late collision enable |
| 14 | CRLLEN | Collision retry limit enable |
| 15 | XFUNEN | Transmit FIFO underrun enable |
| 16 | RXBEN | Receive buffer interrupt enable |
| 17–19 | — | Reserved |
| 20 | MAGEN | Magic packet received interrupt enable |
| 21 | MMRDEN | MII management read completion interrupt enable |
| 22 | MMWREN | MII management write completion interrupt enable |
| 23 | GRSCEN | Graceful receive stop complete interrupt enable |
| 24 | RXFEN | Receive frame interrupt enable |
| 25–26 | — | Reserved |
| 27 | FGPIEN | Filer general purpose interrupt enable |
| 28 | FIREN | Filer invalid result interrupt enable |
| 29 | FIQEN | Filed frame to invalid queue interrupt enable |
| 30 | DPEEN | Data parity error interrupt enable |
| 31 | PERREN | Receive frame parse error enable |

19.5.3.1.5 Error Disabled Register (EDIS)

Figure 19-6 describes the definition for the EDIS register. The error disabled register allows the user to disable an error interruption, possibly to avoid spurious error indications external to the eTSECs.




Figure 19-6. EDIS Register Definition

Table 19-10 describes the fields of the EDIS register.

Table 19-10. EDIS Field Descriptions

| Bits | Name | Description |
|-------|----------|---|
| 0–1 | — | Reserved |
| 2 | BSYDIS | Busy disable. 0 Allow eTSEC to report IEVENT[BSY] status and halt buffer descriptor queue if BSY condition occurs. 1 Do not set IEVENT[BSY] and do not halt buffer descriptor queue if BSY condition occurs. |
| 3 | EBERRDIS | Ethernet controller bus error disable. 0 Allow eTSEC to report IEVENT[EBERR] status and halt buffer descriptor queue if EBERR condition occurs. 1 Do not set IEVENT[EBERR] and do not halt buffer descriptor queue if EBERR condition occurs. |
| 4–6 | — | Reserved |
| 7 | BABTDIS | Babbling transmit error disable. 0 Allow eTSEC to report IEVENT[BABT] status and set the buffer descriptor TR field. 1 Do not set IEVENT[BABT] nor the buffer descriptor TR field. |
| 8 | — | Reserved |
| 9 | TXEDIS | Transmit error disable. 0 Allow eTSEC to report IEVENT[TXE] status. 1 Do not set IEVENT[TXE] if TXE condition occurs. |
| 10–12 | — | Reserved |
| 13 | LCDIS | Late collision disable. 0 Allow eTSEC to report IEVENT[LC] status, set the buffer descriptor LC field, and halt buffer descriptor queue if LC condition occurs. 1 Do not set IEVENT[LC] nor the buffer descriptor LC field, and do not halt buffer descriptor queue if LC condition occurs. |

Table 19-10. EDIS Field Descriptions (continued)

| Bits | Name | Description |
|-------|---------|---|
| 14 | CRLDIS | Collision retry limit disable. 0 Allow eTSEC to report IEVENT[CRL] status, set the buffer descriptor RL field, and halt buffer descriptor queue if CRL condition occurs. 1 Do not set IEVENT[CRL] nor the buffer descriptor RL field, and do not halt buffer descriptor queue if CRL condition occurs. |
| 15 | XFUNDIS | Transmit FIFO underrun disable. 0 Allow eTSEC to report IEVENT[XFUN] status, set the buffer descriptor UN field, and halt buffer descriptor queue if XFUN condition occurs. 1 Do not set IEVENT[XFUN] nor the buffer descriptor UN field, and do not halt buffer descriptor queue if XFUN condition occurs. |
| 16–27 | — | Reserved |
| 28 | FIRDIS | Filer invalid result error disable. 0 Allow eTSEC to report IEVENT[FIR] status. 1 Do not set IEVENT[FIR] if eTSEC fails to reach a definite filer result when attempting to file a received frame, but discard the frame silently. |
| 29 | FIQDIS | Filed frame to invalid queue error disable. 0 Allow eTSEC to report IEVENT[FIQ] status. 1 Do not set IEVENT[FIQ] if eTSEC attempts to file a received frame to an invalid (disabled) RxBD ring, but discard the frame silently. |
| 30 | DPEDIS | Data parity error disable. 0 Allow eTSEC to report IEVENT[DPE] status. 1 Do not set IEVENT[DPE] if a parity error occurs in eTSEC's FIFO or filer arrays. |
| 31 | PERRDIS | Receive frame parse error disable. 0 Allow eTSEC to report IEVENT[PERR] status. 1 Do not set IEVENT[PERR] if a parse error occurs on a received frame. |

19.5.3.1.6 Ethernet Control Register (ECNTRL)

ECNTRL is a register writable by the user to reset, configure, and initialize the eTSEC. Note that the FIFM, GMIIM, TBIM, RPM, and RMM fields are read-only, having been set after sampling signals at power-on-reset. (Refer to the TSEC mode in [Section 4.3.2.2, “Reset Configuration Word High Register \(RCWHR\).”](#)).

Figure 19-7 describes the definition for the ECNTRL register.



Figure 19-7. ECNTRL Register Definition

Table 19-11 describes the fields of the ECNTRL register.

Table 19-11. ECNTRL Field Descriptions

| Bits | Name | Description |
|-------|--------|---|
| 0–16 | — | Reserved |
| 17 | CLRCNT | Clear all statistics counters and carry registers. 0 Allow MIB counters to continue to increment and keep any overflow indicators. 1 Reset all MIB counters and CAR1 and CAR2. This bit is self-resetting. |
| 18 | AUTOZ | Automatically zero MIB counter values and carry registers. 0 The user must write the addressed counter zero after a host read. 1 The addressed counter value is automatically cleared to zero after a host read. This is a steady state signal and must be set prior to enabling the Ethernet controller and must not be changed without proper care. |
| 19 | STEN | MIB counter statistics enabled. 0 Statistics not enabled 1 Enables internal counters to update This is a steady state signal and must be set prior to enabling the Ethernet controller and must not be changed without proper care. |
| 20–24 | — | Reserved |
| 25 | GMIIM | GMII interface mode. Not supported. |
| 26 | TBIM | (Reduced) ten-bit interface mode. If this bit is set, reduced ten-bit interface (RTBI) mode is enabled. This bit can be pin-configured at reset to set or clear. See Section 4.3.2, “Reset Configuration Words.” 0 MII or RMII mode interface 1 RTBI mode interface |
| 27 | RPM | Reduced-pin mode for Gigabit interfaces. If this bit is set, a reduced-pin interface is expected on Ethernet interfaces. RPM and RMM are never set together. This register can be pin-configured at reset to 0 or 1. 0 SGMII or MII in non-reduced-pin mode configuration 1 RGMII or RTBI reduced-pin mode |
| 28 | R100M | RGMII/RMII 100 mode. This bit is ignored unless SGMIIIM, RPM or RMM are set and MACCFG2[I/F Mode] is assigned to 10/100 (01). 0 RGMII is in 10 Mbps mode RMII is in 10 Mbps mode, and every 10th RMII Reference clock is used to transfer data SGMII is in 10 Mbps mode, and every 100th SGMII Reference clock is used to transfer data 1 RGMII is in 100 Mbps mode RMII is in 100 Mbps mode, and data is transferred on every Reference clock SGMII is in 100 Mbps mode, and every 10th SGMII Reference clock is used to transfer data This bit must be cleared for 1-Gbps SGMII operation. |
| 29 | RMM | Reduced-pin mode for 10/100 interfaces. If this bit is set, an RMII pin interface is expected. RMM must be 0 if RPM = 1. This register can be pin-configured at reset to 0 or 1. 0 Non-RMII interface mode 1 RMII interface mode |

Table 19-11. ECNTRL Field Descriptions (continued)

| Bits | Name | Description |
|------|--------|--|
| 30 | SGMIIM | Serial GMII mode. If this bit is set, a SGMII pin interface is expected to be connected via an on chip SerDes. This register can be pin-configured at reset to 0 or 1. 0 SGMII mode disabled. eTSEC connected via a parallel interface. 1 SGMII mode enabled. |
| 31 | — | Reserved |

The different interface configurations indicated by registers ECNTRL and MACCFG2 are summarized in [Table 19-12](#).

Table 19-12. eTSEC Interface Configurations

| Interface Mode | ECNTRL Field | | | | | | | MACCFG2 Field |
|-----------------|-------------------|--------------------|------|-----|-------|-----|--------|---------------|
| | FIFM ¹ | GMIIM ² | TBIM | RPM | R100M | RMM | SGMIIM | I/F Mode |
| MII 10/100 Mbps | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 01 |
| RMII 100 Mbps | 0 | 0 | 0 | 0 | 1 | 1 | 0 | 01 |
| RMII 10 Mbps | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 01 |
| RGMII 1Gbps | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 10 |
| RGMII 100 Mbps | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 01 |
| RGMII 10 Mbps | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 01 |
| RTBI 1Gbps | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 10 |
| SGMII 1 Gbps | 0 | 0 | 1 | 0 | 0 | 0 | 1 | 10 |
| SGMII 100 Mbps | 0 | 0 | 1 | 0 | 1 | 0 | 1 | 01 |
| SGMII 10 Mbps | 0 | 0 | 1 | 0 | 0 | 0 | 1 | 01 |

¹ FIFM bit is not supported.

² GMIIM bit is not supported.

19.5.3.1.7 Pause Time Value Register (PTV)

PTV is a 32-bit register written by the user to store the pause duration used when the eTSEC initiates an IEEE 802.3 PAUSE control frame through TCTRL[TFC_PAUSE]. The low-order 16 bits (PT) represent the pause time and the high-order 16 bits (PTE) represent the extended pause control parameter. The pause time is measured in units of *pause_quanta*, equal to 512 bit times. The pause time can range from 0 to 65,535 *pause_quanta*, or 0 to 33,553,920 bit times. See [Section 19.6.2.9, “Flow Control,”](#) for additional details. [Figure 19-8](#) describes the definition for the PTV register.

Offset eTSEC1:0x2_4028; eTSEC2:0x2_5028

Access: Read/Write




Figure 19-8. PTV Register Definition

Table 19-13 describes the fields of the PTV register.

Table 19-13. PTV Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | PTE | Extended pause control. This field allows software to add a 16-bit additional control parameter into the PAUSE frame to be sent when TCTRL[TFC_PAUSE] is set. Note that current IEEE 802.3 PAUSE frame format requires this parameter to be cleared. |
| 16–31 | PT | Pause time value. Represents the 16-bit pause quanta (that is, 512 bit times). This pause value is used as part of the PAUSE frame to be sent when TCTRL[TFC_PAUSE] is set. See Section 19.6.2.9, “Flow Control,” on page 19-151 for more information. |

19.5.3.1.8 DMA Control Register (DMACTRL)

DMACTRL is writable by the user to configure the DMA block. Figure 19-9 describes the definition for the DMACTRL register.

Offset eTSEC1:0x2_402C; eTSEC2:0x2_502C

Access: Read/Write




Figure 19-9. DMACTRL Register

Table 19-14 describes the fields of the DMACTRL register.

Table 19-14. DMACTRL Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0–15 | — | Reserved |
| 16 | LE | Little-endian descriptor mode enable. This bit controls both the reading and writing of descriptors; data buffers are always transferred in network byte order. 0 RxBs and TxBs are interpreted with big-endian byte ordering, as shown in Section 19.6.7.1, “Data Buffer Descriptors.” 1 RxBs and TxBs are interpreted with little-endian byte ordering. That is, the 16 bits of flags are considered a complete half-word unit, the buffer length is considered another complete half-word unit, and the buffer pointer is considered a complete word unit. |
| 17–23 | — | Reserved |
| 24 | TDSEN | Tx Data snoop enable. 0 Disables snooping of all transmit frames from memory. 1 Enables snooping of all transmit frames from memory. |

Table 19-14. DMACTRL Field Descriptions (continued)

| Bits | Name | Description |
|------|---------|--|
| 25 | TBDSSEN | TxBD snoop enable. 0 Disables snooping of all transmit BD memory accesses. 1 Enables snooping of all transmit BD memory accesses. |
| 26 | — | Reserved |
| 27 | GRS | Graceful receive stop. If this bit is set, the Ethernet controller stops receiving frames following completion of the frame currently being received. (That is, after a valid end of frame was received). The contents of the Rx FIFO are then written to memory, and the IEVENT[GRSC] is set to indicate that all current receive buffers have been closed. Because the receive enable bit of the MAC may still be set, the MAC may continue to receive but the eTSEC ignores the receive data until GRS is cleared. If this bit is cleared, the eTSEC scans the input data stream for the start of a new frame (preamble sequence and start of frame delimiter) and the first valid frame received uses the next RxBD. If GRS is set, the user must monitor the graceful receive stop complete (GRSC) bit in the IEVENT register to insure that the graceful receive stop was completed. The user can then clear IEVENT[GRSC] and can write to receive registers that are accessible to both user and the eTSEC hardware without fear of conflict. 0 eTSEC scans input data stream for valid frame. 1 eTSEC stops receiving frames following completion of current frame. |
| 28 | GTS | Graceful transmit stop. If this bit is set, the Ethernet controller stops transmission after all frames that are currently in the Tx FIFO or scheduled have been transmitted, and the GTSC interrupt in the IEVENT register is asserted. A frame that has started reading buffer descriptors or data from memory is read to completion and transmitted before the GTSC interrupt occurs. However, if no frame has been scheduled for transmission and the Tx FIFO is empty, the GTSC interrupt is asserted immediately. Once transmission has completed, clearing GTS “restart” transmit.0Controller continues. 1 Controller stops transmission after completion of current frame. |
| 29 | TOD | Transmit on demand for TxBD ring 0. This bit is applicable only to the transmitter, and requires both TCTRL[TXSCHED] = 00 and DMACTRL[WOP] = 0. If 1 is written to this bit, the eTSEC immediately begins fetching the next TxBD from ring 0, avoiding waiting the normal polling time to check the TxBD's R bit. This bit is always read as 0. 0 eTSEC continues waiting for the TxBD ring 0 poll timer to expire. 1 eTSEC immediately fetches a new TxBD from ring 0. |
| 30 | WWR | Write with response. This bit gives the user the assurance that a BD was updated in memory before it receives an interrupt concerning a transmit or receive frame. 0 Do not wait for acknowledgement from system for BD writes before setting IEVENT bits. 1 Before setting IEVENT bits TXB, TXF, TXE, XFUN, LC, CRL, RXB, RXF, the eTSEC waits for acknowledgement from system that the transmit or receive BD being updated was stored in memory. |
| 31 | WOP | Wait or poll for TxBD ring 0. This bit, which is applicable only to the transmitter and when TCTRL[TXSCHED] = 00, provides the user the option for the eTSEC to periodically poll TxBDs or to wait for software to tell eTSEC to fetch a buffer descriptor. While operating in the “Wait” mode, the eTSEC allows two additional reads of a descriptor which is not ready before entering a halt state. No interrupt is driven. To resume transmission, software must clear TSTAT[THLT]. 0 Poll TxBD on ring 0 every 512 serial clocks. 1 Do not poll, but wait for TSTAT[THLT] to be cleared by the user. |

19.5.3.1.9 TBI Physical Address Register (TBIPA)

The TBIPA, shown in [Figure 19-10](#), is writable by the user to assign a physical address to the TBI (or RTBI) for MII management configuration. The TBI registers are accessed at the offset of TBIPA. For detailed descriptions of the TBI registers (the MII register set for the ten-bit interface) refer to [Section 19.5.4, “Ten-Bit Interface \(TBI\).”](#)




Figure 19-10. TBIPA Register Definition

Table 19-15 describes the fields of the TBIPA register.

Table 19-15. TBIPA Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0–26 | — | Reserved |
| 27–31 | TBIPA | This field is used to program the PHY address of the ten-bit interface’s MII management bus. To access the TBI register the user must write the TBIPA value to the MIIMADD [PHY Address] register located in the MAC register section. PHY Address 0 is reserved. Refer to Section 19.5.3.5.8, “MII Management Address Register (MIIMADD).” |

19.5.3.2 eTSEC Transmit Control and Status Registers

This section describes the control and status registers that are used specifically for transmitting Ethernet frames. All of the registers are 32 bits wide.

19.5.3.2.1 Transmit Control Register (TCTRL)

This register is writable by the user to configure the transmit block. [Figure 19-11](#) describes the TCTRL register.




Figure 19-11. TCTRL Register Definition

Table 19-16 describes the fields of the TCTRL register.

Table 19-16. TCTRL Field Descriptions

| Bits | Name | Description |
|------|--------|---|
| 0–16 | — | Reserved |
| 17 | IPCSEN | IP header checksum generation enable. When set, the eTSEC offloads IPv4 header checksum generation. See Section 19.6.3.2, “Transmit Path Off-Load and Tx PTP Packet Parsing,” on page 19-159. 0 IP header checksum generation is disabled even if enabled in a transmit frame control block. 1 IP header checksum generation is performed for IPv4 headers as determined by the settings in the current transmit frame control block. |

Table 19-16. TCTRL Field Descriptions (continued)

| Bits | Name | Description |
|-------|-----------|---|
| 18 | TUCSEN | TCP/UDP header checksum generation enable. When set, the eTSEC offloads TCP or UDP header checksum generation. See Section 19.6.3.2, “Transmit Path Off-Load and Tx PTP Packet Parsing,” on page 19-159 . 0 TCP or UDP header checksum generation is disabled even if enabled in a transmit frame control block. 1 TCP or UDP header checksum generation is performed as determined by the settings in the current transmit frame control block. |
| 19 | VLINS | VLAN (IEEE Std. 802.1Q) tag insertion enable. Applicable only for transmission through the Ethernet MAC. 0 Do not insert a VLAN tag into the frame. 1 Insert a VLAN tag into the frame. If the frame FCB has a valid VLAN field, use the FCB to source the VLAN control word, otherwise take the default VLAN control word from register DFVLAN. |
| 20 | THDF | Transmit half-duplex flow control under software control for 10-/100-Mbps half-duplex media. This bit is not self-resetting. 0 Disable back pressure 1 Back pressure is applied to media by raising carrier |
| 21–26 | — | Reserved |
| 27 | RFC_PAUSE | Receive flow control pause frame (written by the eTSEC). This read-only status bit is set if a flow control pause frame was received and the transmitter is paused for the duration defined in the received pause frame. This bit automatically clears after the pause duration is complete. 0 Pause duration complete. 1 Flow control pause frame received. |
| 28 | TFC_PAUSE | Transmit flow control pause frame. Set this bit to transmit a PAUSE frame. If this bit is set, the MAC stops transmission of data frames after the currently transmitting frame completes. Next, the MAC transmits a pause control frame with the duration value obtained from the PTV register. The TXC event occurs after sending the pause control frame. Finally, the controller clears TFC_PAUSE and resumes transmitting data frames as before. Note that pause control frames can still be transmitted if the Tx controller is stopped due to user assertion of DMACTRL[GTS] or reception of a PAUSE frame. 0 No request for Tx PAUSE frame pending or transmission complete. 1 Software request for Tx PAUSE frame pending. |

Table 19-16. TCTRL Field Descriptions (continued)

| Bits | Name | Description |
|-------|---------|--|
| 29–30 | TXSCHED | <p>Transmit ring scheduling algorithm. This field determines which scheme the transmit scheduler uses to arbitrate between the enabled TxBD rings. The scheme chosen also controls how the DMACTRL and TQUEUE bits are interpreted. Ring polling is supported only by mode 00; the other modes require software to restart rings with the TSTAT register. TCP/IP offload can be enabled with any scheduling mode.</p> <p>00 Single polled ring mode. TxBD ring 0 is the only ring serviced, even if other rings are enabled and ready. In this scheduler mode, the DMACTRL[WOP] and DMACTRL[TOD] bits control polling and retry behavior. This mode supports ring polling, and allows fetching of a non-ready TxBD to be retried twice.</p> <p>01 Priority scheduling mode. Frames from enabled TxBD rings are serviced in ascending ring index order.</p> <p>10 Modified weighted round-robin scheduling mode. Each TxBD ring is polled in sequence for frames that are ready for transmission. If a non-ready TxBD is fetched from a ring, that ring is removed from the scheduling pool until software re-enables it. Ready frames are repeatedly transmitted from a chosen ring until its transmission quota is exhausted. The transmission quota for TxBD ring n is set to $WT_n \times 64$ bytes, where WT_n is a weight from the TR03WT/TR47WT registers. If a ring transmits more data than its quota allows, the excess is deducted from its quota on the next transmission opportunity, thereby preventing large frames from monopolizing the eTSEC bandwidth.</p> <p>11 Reserved</p> |
| 31 | — | Reserved |

19.5.3.2.2 Transmit Status Register (TSTAT)

This register is read/write-one-to-clear and is written by the eTSEC to convey DMA status information for each TxBD ring. The halt bit only has meaning for enabled rings. After processing transmit-related interrupts, software should use TSTAT to restart transmission from rings that may have been affected by the interrupt condition. In particular, an error condition that prevents eTSEC from continuing transmission halts DMA from all rings, including the ring that gave rise to the error. [Figure 19-12](#) describes the TSTAT register.

Offset eTSEC1:0x2_4104; eTSEC2:0x2_5104

Access: w1c




Figure 19-12. TSTAT Register Definition

Table 19-17 describes the fields of the TSTAT register.

Table 19-17. TSTAT Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 0 | THLT0 | <p>Transmit halt of ring 0. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN0], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set. Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0 |
| 1 | THLT1 | <p>Transmit halt of ring 1. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN1], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0 |
| 2 | THLT2 | <p>Transmit halt of ring 2. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN2], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0 |

Table 19-17. TSTAT Field Descriptions (continued)

| Bits | Name | Description |
|------|-------|--|
| 3 | THLT3 | <p>Transmit halt of ring 3. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN3], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error: <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read TxBD programming errors: <ul style="list-style-type: none"> • Ready=1 and length=0 </p> |
| 4 | THLT4 | <p>Transmit halt of ring 4. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN4], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error: <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read TxBD programming errors: <ul style="list-style-type: none"> • Ready=1 and length=0 </p> |
| 5 | THLT5 | <p>Transmit halt of ring 5. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN5], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error: <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read TxBD programming errors: <ul style="list-style-type: none"> • Ready=1 and length=0 </p> |

Table 19-17. TSTAT Field Descriptions (continued)

| Bits | Name | Description |
|------|-------|--|
| 6 | THLT6 | <p>Transmit halt of ring 6. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN6], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error: <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read TxBD programming errors: <ul style="list-style-type: none"> • Ready=1 and length=0 </p> |
| 7 | THLT7 | <p>Transmit halt of ring 7. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN7], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include: Bus error: <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read TxBD programming errors: <ul style="list-style-type: none"> • Ready=1 and length=0 </p> |
| 8–15 | — | Reserved |
| 16 | TXF0 | Transmit frame event occurred on ring 0. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring. |
| 17 | TXF1 | Transmit frame event occurred on ring 1. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring. |
| 18 | TXF2 | Transmit frame event occurred on ring 2. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring. |
| 19 | TXF3 | Transmit frame event occurred on ring 3. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring. |
| 20 | TXF4 | Transmit frame event occurred on ring 4. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring. |
| 21 | TXF5 | Transmit frame event occurred on ring 5. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring. |
| 22 | TXF6 | Transmit frame event occurred on ring 6. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring. |

Table 19-17. TSTAT Field Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 23 | TXF7 | Transmit frame event occurred on ring 7. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring. |
| 24–31 | — | Reserved |

19.5.3.2.3 Default VLAN Control Word Register (DFVLAN)

This register defines the default value for the VLAN Ethertype and control word when VLAN tags are automatically inserted by the eTSEC, and no per-frame VLAN data is supplied by software. On receive, this register defines a customizable VLAN Ethertype for automatic deletion. Note that an Ethertype of 0x8808 (Control Word) is not permitted as a custom VLAN tag. Frames with an Ethertype of 0x8808 are dropped by the receiver. In the case of frames containing stacked VLAN tags, this register defines the tag associated with the outer or metropolitan area VLAN. [Figure 19-13](#) describes the DFVLAN register.




Figure 19-13. DFVLAN Register Definition

[Table 19-18](#) describes the fields of the DFVLAN register.

Table 19-18. DFVLAN Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | TAG | This is the default Ethertype used to tag VLAN frames. On transmit, this tag is inserted ahead of the VLAN control word; TAG should be set to 0x8100 for IEEE 802.1Q VLAN. On receive, an Ethertype matching TAG or an Ethertype of 0x8100 marks a VLAN-tagged frame. Note that if using DFVLAN to set a custom ethertype (that is, using a value other than 0x8100), packets received with a custom tag are not counted by any of the RMON counters. Affected counters include TRMGV, RMCA, RBCA, RXCF, RXPF, RXUO, RALN, RFLR, ROVR, RJBR, TMCA, TBCA, TXPF, TXCF. |
| 16–18 | PRI | This is the default value used for the IEEE Std. 802.1p frame priority. |
| 19 | CFI | This is the default value used for the IEEE Std. 802.1Q canonical format indicator. |
| 20–31 | VID | This is the default value used for the virtual-LAN identifier in VLAN-tagged frames. A value of zero is defined as the null VLAN, however field PRI may be still set independently. |

19.5.3.2.4 Transmit Interrupt Coalescing Register (TXIC)

The TXIC register enables and configures the operational parameters for interrupt coalescing associated with transmitted frames. Figure 19-14 describes the definition for the TXIC register.

Offset eTSEC1:0x2_4110; eTSEC2:0x2_5110

Access: Read/Write




Figure 19-14. TXIC Register Definition

Table 19-19 describes the fields of the TXIC register.

Table 19-19. TXIC Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0 | ICEN | Interrupt coalescing enable 0 Interrupt coalescing is disabled. Interrupts are raised as they are received. 1 Interrupt coalescing is enabled. If the eTSEC transmit frame interrupt is enabled (IMASK[TXFEN] is set), an interrupt is raised when the threshold number of frames is reached (defined by TXIC[ICFT]) or when the threshold timer expires (determined by TXIC[ICTT]). |
| 1 | ICCS | Interrupt coalescing timer clock source. 0 The coalescing timer advances count every 64 eTSEC Tx interface clocks (TSEC _n _GTX_CLK). 1 The coalescing timer advances count every 64 system clocks. This mode is recommended for FIFO operation. |
| 2 | — | Reserved |
| 3–10 | ICFT | Interrupt coalescing frame count threshold. While interrupt coalescing is enabled (TXIC[ICEN] is set), this value determines how many frames are transmitted before raising an interrupt. The eTSEC threshold counter is reset to ICFT following an interrupt. The value of ICFT must be greater than zero to avoid unpredictable behavior. |
| 11–15 | — | Reserved |
| 16–31 | ICTT | Interrupt coalescing timer threshold. While interrupt coalescing is enabled (TXIC[ICEN] is set), this value determines the maximum amount of time after transmitting a frame before raising an interrupt. If frames have been transmitted but the frame count threshold has not been met, an interrupt is raised when the threshold timer reaches zero. The threshold timer is reset to the value in this field and begins counting down upon transmission of the first frame having its TxBD[<i>i</i>] bit set. The threshold value is represented in units of 64 clock periods as specified by the timer clock source (TXIC[ICCS]). The value of ICTT must be greater than zero to avoid unpredictable behavior. |

19.5.3.2.5 Transmit Queue Control Register (TQUEUE)

The TQUEUE register, shown in Figure 19-15, selectively enables each of the TxBD rings 0–7. By default, TxBD ring 0 is enabled.




Figure 19-15. TQUEUE Register Definition

Table 19-20 describes the TQUEUE register.

Table 19-20. TQUEUE Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16 | EN0 | Transmit queue 0 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission. |
| 17 | EN1 | Transmit queue 1 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission. |
| 18 | EN2 | Transmit queue 2 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission. |
| 19 | EN3 | Transmit queue 3 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission. |
| 20 | EN4 | Transmit queue 4 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission. |
| 21 | EN5 | Transmit queue 5 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission. |
| 22 | EN6 | Transmit queue 6 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission. |
| 23 | EN7 | Transmit queue 7 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission. |
| 24–31 | — | Reserved |

19.5.3.2.6 TxBD Ring 0–3 Weighting Register (TR03WT)

When modified weighted round-robin Tx scheduling is enabled (TCTRL[TXSCHEd] = 10), this register determines the weighting applied to each transmit queue for queues 0 to 3. For priority-based scheduling,

TR03WT has no effect. A description of how queue weights affect eTSEC's round-robin algorithm appears in [Section 19.6.4.3.2, "Modified Weighted Round-Robin Queuing \(MWRR\)."](#) Figure 19-16 describes the TR03WT register.




Figure 19-16. TR03WT Register Definition

Table 19-21 describes the fields of the TR03WT register.

Table 19-21. TR03WT Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–7 | WT0 | Weighting value for TxBd ring 0 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of $WT0 \times 64$ bytes of data are scheduled for transmission from TxBd ring 0. Clearing this field prevents transmission. |
| 8–15 | WT1 | Weighting value for TxBd ring 1 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of $WT1 \times 64$ bytes of data are scheduled for transmission from TxBd ring 1. Clearing this field prevents transmission. |
| 16–23 | WT2 | Weighting value for TxBd ring 2 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of $WT2 \times 64$ bytes of data are scheduled for transmission from TxBd ring 2. Clearing this field prevents transmission. |
| 24–31 | WT3 | Weighting value for TxBd ring 3 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of $WT3 \times 64$ bytes of data are scheduled for transmission from TxBd ring 3. Clearing this field prevents transmission. |

19.5.3.2.7 TxBd Ring 4–7 Weighting Register (TR47WT)

When modified weighted round-robin Tx scheduling is enabled (TCTRL[TXSCHED] = 10), this register determines the weighting applied to each enabled transmit queue for queues 4 to 7. For priority-based scheduling, TR47WT has no effect. A description of how queue weights affect eTSEC's modified weighted round-robin algorithm appears in [Section 19.6.4.3.2, "Modified Weighted Round-Robin Queuing \(MWRR\)."](#) Figure 19-17 describes the definition for the TR47WT register.



Figure 19-17. TR47WT Register Definition

Table 19-22 describes the fields of the TR47WT register.

Table 19-22. TR47WT Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–7 | WT4 | Weighting value for TxBd ring 4 when TCTRL[TXSCHEd] = 10. On each round of the Tx scheduler, a minimum of WT4 × 64 bytes of data are scheduled for transmission from TxBd ring 4. Clearing this field prevents transmission. |
| 8–15 | WT5 | Weighting value for TxBd ring 5 when TCTRL[TXSCHEd] = 10. On each round of the Tx scheduler, a minimum of WT5 × 64 bytes of data are scheduled for transmission from TxBd ring 5. Clearing this field prevents transmission. |
| 16–23 | WT6 | Weighting value for TxBd ring 6 when TCTRL[TXSCHEd] = 10. On each round of the Tx scheduler, a minimum of WT6 × 64 bytes of data are scheduled for transmission from TxBd ring 6. Clearing this field prevents transmission. |
| 24–31 | WT7 | Weighting value for TxBd ring 7 when TCTRL[TXSCHEd] = 10. On each round of the Tx scheduler, a minimum of WT7 × 64 bytes of data are scheduled for transmission from TxBd ring 7. Clearing this field prevents transmission. |

19.5.3.2.8 Transmit Buffer Descriptor Pointers 0–7 (TBPTR0–TBPTR7)

TBPTR0–TBPTR7 each contains the low-order 32 bits of the next transmit buffer descriptor address for their respective TxBd ring. Figure 19-18 describes the TBPTR registers. These registers takes on the value of their ring’s associated TBASE when the TBASE register is written by software. Software must not write TBPTR0–TBPTR7 while eTSEC is actively transmitting frames. However, TBPTR0– TBPTR7 can be modified when the transmitter is disabled or when no Tx buffer is in use (after a GRACEFUL STOP TRANSMIT command is issued and the frame completes its transmission) in order to change the next TxBd eTSEC transmits.



Figure 19-18. TBPTR0–TBPTR7 Register Definition

Table 19-23 describes the fields of the TBPTR_n register.

Table 19-23. TBPTR_n Field Descriptions

| Bits | Name | Description |
|-------|--------------------|--|
| 0–28 | TBPTR _n | Current TxBd pointer for TxBd ring <i>n</i> . Points to the current BD being processed or to the next BD the transmitter uses when it is idling. When the end of the TxBd ring is reached, eTSEC initializes TBPTR _n to the value in the corresponding TBASE _n . The TBPTR register is internally written by the eTSEC’s DMA controller during transmission. The pointer increments by eight (bytes) each time a descriptor is closed successfully by the eTSEC. Note that the three least significant bits of this register are read-only and zero. After an error condition, the eTSEC returns TBPTR _n to point to the first BD of the frame partially transmitted. |
| 29–31 | — | Reserved |

19.5.3.2.9 Transmit Descriptor Base Address Registers (TBASE0–TBASE7)

The TBASE n registers are written by the user with the base address of each TxBD ring n . Each such value must be divisible by eight, since the three least significant bits always write as 000. Figure 19-19 describes the definition for the TBASE n registers.




Figure 19-19. TBASE Register Definition

Table 19-24 describes the fields of the TBASE n registers.

Table 19-24. TBASE0–TBASE7 Field Descriptions

| Bits | Name | Description |
|-------|-----------|---|
| 0–28 | TBASE n | Transmit base for ring n . TBASE defines the starting location in the memory map for the eTSEC TxBDs. This field must be 8-byte aligned. Together with setting the W (wrap) bit in the last BD, the user can select how many BDs to allocate for the transmit packets. The user must initialize TBASE before enabling the eTSEC transmit function on the associated ring. |
| 29–31 | — | Reserved |

19.5.3.2.10 Transmit Time Stamp Identification Register (TMR_TXTS1–2_ID)

Transmit time stamp identification register (TMR_TXTS n _ID). This register holds the identification number of the transmitted frame corresponding to the timestamp captured in TMR_TXTS n _H/L. Each time the eTSEC is instructed to capture the timestamp of an outgoing frame via TxFCB[PTP] the associated field in TxFCB[PTP_ID] is stored in this register, overwriting the previous value.

This register is read only in normal operation. Figure 19-20 describes the definition for the TMR_TXTS n _ID register.




Figure 19-20. TMR_TXTS n _ID Register Definition

Table 19-25 describes the fields of the TMR_TXTS n _ID register.

Table 19-25. TMR_TXTS n _ID Register Field Descriptions

| Bits | Name | Description |
|-------|---------|------------------------------------|
| 0–15 | — | Reserved |
| 16–31 | TXTS_ID | Tx time stamp identification field |

19.5.3.2.11 Transmit Time Stamp Register (TMR_TXTS1–2_H/L)

Transmit stamp register (TMR_TXTS_n_H/L). This register holds the value of the TMR_CNT_H/L when a frame tagged for timestamp capture (via Tx FCB[PTP]) is transmitted. Upon transmission of the start of frame symbol of such a frame, the value in TMR_CNT_H/L is copied into TMR_TXTS_n_H/L.

This register is read only in normal operation. Figure 19-21 depicts TMR_TXTS_n_H/L.




Figure 19-21. TMR_TXTS_n_H/L Register Definition

Table 19-26 describes the fields of the TMR_TXTS_n_H/L register.

Table 19-26. TMR_TXTS_n_H/L Register Field Descriptions

| Bits | Name | Description |
|------|----------|--|
| 0–63 | TXTS_H/L | Time stamp field of the transmitted PTP packet’s start of frame detection. |

19.5.3.3 eTSEC Receive Control and Status Registers

This section describes the control and status registers that are used specifically for receiving Ethernet frames. All of the registers are 32 bits wide.

19.5.3.3.1 Receive Control Register (RCTRL)

The RCTRL register is programmed by the user and controls the operational mode of the receiver. It must be written only after a system reset (at initialization) or after a graceful receive stop has completed.

Figure 19-22 describes the RCTRL register.




Figure 19-22. RCTRL Register Definition

Table 19-27 describes the fields of the RCTRL register.

Table 19-27. RCTRL Field Descriptions

| Bits | Name | Description |
|-------|--------|---|
| 0–6 | — | Reserved |
| 7 | TS | Time stamp incoming packets as padding bytes. PAL field is set to 8 if the PAL field is programmed to less than 8. Must be set to zero if TMR_CTRL[TE]=0. |
| 8–10 | — | Reserved |
| 11–15 | PAL | Packet alignment padding length. If not zero, PAL (1–31) bytes of zero padding are inserted before the start of each received frame, but following the RxFCB if TOE is enabled. For Ethernet where optional preamble extraction is enabled, the padding appears before the preamble, otherwise the padding precedes the layer 2 header. The value of PAL can be set so that the start of the IP header in the receive data buffer is aligned to a 32-bit boundary. Normally, setting PAL = 2 provides minimal padding to ensure such alignment of the IP header. Note that the minimum zero padding value for this field should be PAL–8 if the TS field is set and 0 when PAL is < 8. |
| 16 | — | Reserved |
| 17 | LFC | Lossless flow control. When set, the eTSEC determines the number of free BDs (through RQPARM n [LEN] and RBTPTR n) in each active ring. Should the free BD count in an active ring drop below its setting for RQPARM n [FBTHR], the eTSEC asserts link layer flow control. For full-duplex Ethernet connections, the eTSEC emits a pause frame as if TCTRL[TFC_PAUSE] was set. For FIFO packet interface connections, the RFC signal is asserted. 0 Disabled. This is the default 1 Enabled, calculate the free BDs in each active ring and assert link layer flow control if required. |
| 18 | VLEX | Enable automatic VLAN tag extraction and deletion from Ethernet frames. 0 Do not delete VLAN tags from received Ethernet frames. 1 If a VLAN tag is seen after the Ethernet source address, and PRSDEP is non-zero, delete the VLAN tag and return the VLAN control word in the frame control block returned with this frame. Note that if PRSDEP is cleared, VLEX must be cleared as well. (VLAN tag extraction is only supported when the parser is enabled.) |
| 19 | FILREN | Filer enable. When set, the receive frame filer is enabled. This file accepted frames to a particular RxBD ring according to rules defined in the filer table. In this case, PRSDEP must not be cleared. 0 Do not search the receive queue filer table for received frames. All received frames are sent to RxBD ring 0 by default. 1 Search the receive queue filer table for received frames, and let the filer determine the index of the RxBD ring for each frame. Note that if PRSDEP is cleared, FILREN must be cleared as well. |
| 20 | FSQEN | Enable single-queue mode for the receive frame filer. This bit is ignored unless FILREN is also set. 0 The filer chooses the RxBD ring using the least significant bits of the virtual queue ID as a ring index. 1 The filer always attempts to file received frames to ring 0, regardless of virtual queue ID. This mode is intended for operating the filer as a packet classification engine. |
| 21 | GHTX | Group address hash table extend. By default, the group address hash table is 256 entries (as defined by registers GADDR0–GADDR7); registers IGADDR0–IGADDR7 are then used to define the individual address hash table. When this bit is set, the hash table is extended to a total of 512 entries (IGADDR0–IGADDR7 are then the first 256 entries of the extended 512-entry group address hash table). 0 Both the individual and group hash functions are the 8 MSBs of the CRC-32 of the Ethernet destination address. 1 The group hash function is the 9 MSBs of the CRC-32 of the Ethernet destination address. The individual address hash function is unavailable. |

Table 19-27. RCTRL Field Descriptions (continued)

| Bits | Name | Description |
|-------|--------|--|
| 22 | IPCSEN | IP Checksum verification enable. See Section 19.6.3.3, “Receive Path Off-Load.” 0 IPv4 header checksums are not verified by the eTSEC—even if layer 3 parsing is enabled. 1 Perform IPv4 header checksum verification if PRSDEP > 01. |
| 23 | TUCSEN | TCP or UDP Checksum verification enable. See Section 19.6.3.3, “Receive Path Off-Load.” 0 TCP or UDP checksums are not verified by the eTSEC—even if layer 4 parsing is enabled. 1 Perform TCP or UDP checksum verification if PRSDEP = 11. |
| 24–25 | PRSDEP | Parser control. The level of parser layer recognition is determined as follows: 00 Parser disabled. Receive frame filter must also be disabled by clearing RCTRL[FILREN]. 01 Only L2 (Ethernet) protocols are recognized. 10 L2 and L3 (IP) protocols are recognized. 11 L2, L3, and L4 (TCP/UDP) protocols are recognized. If this field is non-zero, a TOE frame control block is prepended to the received frame, and the first RxB D points to the FCB. Note that if PRSDEP is cleared, VLEX must be cleared as well. (VLAN tag extraction is only supported when the parser is enabled.) Also, if PRSDEP is cleared, FILREN must also be cleared. |
| 26 | — | Reserved |
| 27 | BC_REJ | Broadcast frame reject. If this bit is set, frames with DA (destination address) = FFFF_FFFF_FFFF are rejected unless RCTRL[PROM] is set. If both BC_REJ and RCTRL[PROM] are set, then frames with broadcast DA are accepted and the M (MISS) bit is set in the receive BD. |
| 28 | PROM | Promiscuous mode. All Ethernet frames, regardless of destination address, are accepted. |
| 29 | RSF | Receive short frame mode. When set, enables the reception of frames shorter than 64 bytes. 0 Ethernet frames less than 64B in length are silently dropped. 1) Frames more than 16B and less than 64B in length are accepted upon a DA match. Note that frames less than or equal to 16B in length are always silently dropped. |
| 30 | EMEN | Exact match MAC address enable. If this bit is set, the MAC01ADDR1–MAC15ADDR1 and MAC01ADDR2–MAC15ADDR2 registers are recognized as containing MAC addresses aliasing the MAC’s station address. Setting this bit therefore allows eTSEC to receive Ethernet frames having a destination address matching one of these 15 addresses. |
| 31 | — | Reserved |

19.5.3.3.2 Receive Status Register (RSTAT)

The eTSEC writes to this register under the following conditions:

- A frame interrupt event occurred on one or more RxB D rings
- The receiver runs out of descriptors due to a busy condition on a RxB D ring
- The receiver was halted because an error condition was encountered while receiving a frame

Writing 1 to any bit of this register clears it. Software should clear the QHLT bit to take eTSEC’s receiver function out of halt state for the associated queue. [Figure 19-23](#) describes the definition for the RSTAT register.




Figure 19-23. RSTAT Register Definition

[Table 19-28](#) describes the fields of the RSTAT register.

Table 19-28. RSTAT Field Descriptions

| Bits | Name | Description |
|------|-------|---|
| 0–7 | — | Reserved |
| 8 | QHLT0 | RxBD queue 0 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT0 to be set.) The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted. |
| 9 | QHLT1 | RxBD queue 1 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT1 to be set.) The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted. |
| 10 | QHLT2 | RxBD queue 2 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT2 to be set.) The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted. |
| 11 | QHLT3 | RxBD queue 3 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT3 to be set.) The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted. |
| 12 | QHLT4 | RxBD queue 4 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT4 to be set.) The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted. |

Table 19-28. RSTAT Field Descriptions (continued)

| Bits | Name | Description |
|-------|-------|--|
| 13 | QHLT5 | RxBD queue 5 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT5 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted. |
| 14 | QHLT6 | RxBD queue 6 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT6 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted. |
| 15 | QHLT7 | RxBD queue 7 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT7 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted. |
| 16–23 | — | Reserved |
| 24 | RXF0 | Receive frame event occurred on ring 0. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring. |
| 25 | RXF1 | Receive frame event occurred on ring 1. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring. |
| 26 | RXF2 | Receive frame event occurred on ring 2. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring. |
| 27 | RXF3 | Receive frame event occurred on ring 3. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring. |
| 28 | RXF4 | Receive frame event occurred on ring 4. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring. |
| 29 | RXF5 | Receive frame event occurred on ring 5. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring. |
| 30 | RXF6 | Receive frame event occurred on ring 6. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring. |
| 31 | RXF7 | Receive frame event occurred on ring 7. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring. |

19.5.3.3.3 Receive Interrupt Coalescing Register (RXIC)

The RXIC register enables and configures the operational parameters for interrupt coalescing associated with received frames. [Figure 19-24](#) describes the RXIC register.

Offset eTSEC1:0x2_4310; eTSEC2:0x2_5310

Access: Read/Write




Figure 19-24. RXIC Register Definition

Table 19-29 describes the fields of the RXIC register.

Table 19-29. RXIC Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0 | ICEN | Interrupt coalescing enable 0 Interrupt coalescing is disabled. Interrupts are raised as they are received. 1 Interrupt coalescing is enabled. If the eTSEC receive frame interrupt is enabled (IMASK[RXFEN] is set), an interrupt is raised when the threshold number of frames is reached (defined by RXIC[ICFT]) or when the threshold timer expires (determined by RXIC[ICTT]). |
| 1 | ICCS | Interrupt coalescing timer clock source. 0 The coalescing timer advances count every 64 eTSEC Rx interface clocks (TSECn_GTX_CLK). 1 The coalescing timer advances count every 64 system clocks. This mode is recommended for FIFO operation. |
| 2 | — | Reserved |
| 3–10 | ICFT | Interrupt coalescing frame count threshold. While interrupt coalescing is enabled (RXIC[ICE] is set), this value determines how many frames are received before raising an interrupt. The eTSEC threshold counter is reset to ICFT following an interrupt. The value of ICFT must be greater than zero avoid unpredictable behavior. |
| 11–15 | — | Reserved |
| 16–31 | ICTT | Interrupt coalescing timer threshold. While interrupt coalescing is enabled (RXIC[ICE] is set), this value determines the maximum amount of time after receiving a frame before raising an interrupt. If frames have been received but the frame count threshold has not been met, an interrupt is raised when the threshold timer reaches zero. The threshold timer is reset to the value in this field and begins counting down upon receiving the first frame having its RxBD[I] bit set. The threshold value is represented in units equal to 64 periods of the clock specified by RXIC[ICCS]. ICTT must be greater than zero to avoid unpredictable behavior. |

19.5.3.3.4 Receive Queue Control Register (RQUEUE)

The RQUEUE register enables each of the RxBD rings 0–7. By default, RxBD ring 0 is enabled. Figure 19-25 describes the definition for the RQUEUE register.




Figure 19-25. RQUEUE Register Definition

Table 19-30 describes the RQUEUE register.

Table 19-30. RQUEUE Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–23 | — | Reserved |
| 24 | EN0 | Receive queue 0 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception. |

Table 19-30. RQUEUE Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 25 | EN1 | Receive queue 1 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 26 | EN2 | Receive queue 2 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 27 | EN3 | Receive queue 3 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 28 | EN4 | Receive queue 4 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 29 | EN5 | Receive queue 5 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 30 | EN6 | Receive queue 6 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 31 | EN7 | Receive queue 7 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |

19.5.3.3.5 Receive Bit Field Extract Control Register (RBIFX)

The RBIFX register provides a set of four 6-bit offsets for locating up to four octets in a received frame and passing them to the receive queue filer as the user-defined ARB property. Through RBIFX a custom ARB filer property can be constructed from arbitrary bytes, which allows frame filing on the basis of bitfields not ordinarily provided to the filer, such as bits from the Ethernet preamble or TCP flags. The value of property ARB is the concatenation of {B0, B1, B2, B3} to 32-bits, where B0–B3 are the bytes as defined by RBIFX.

Figure 19-26 describes the definition for the RBIFX register.




Figure 19-26. RBIFX Register Definition

Table 19-31 describes the RBIFX register.

Table 19-31. RBIFX Field Descriptions

| Bits | Name | Description |
|-------|----------|---|
| 0–1 | B0CTL | Location of byte 0 of property ARB. 00 Byte 0 is not extracted, and appears as zero in property ARB. 01 Byte 0 is located in the received frame at offset (B0OFFSET – 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B0OFFSET less than 8 are reserved in FIFO modes. 10 Byte 0 is located in the received frame at offset B0OFFSET bytes from the byte after the last byte of the layer 2 header. 11 Byte 0 is located in the received frame at offset B0OFFSET bytes from the byte after the last byte of the layer 3 header. |
| 2–7 | B0OFFSET | Offset relative to the header defined by B0CTL that locates byte 0 of property ARB. An effective offset of zero points to the first byte of the specified header. |
| 8–9 | B1CTL | Location of byte 1 of property ARB. 00 Byte 1 is not extracted, and appears as zero in property ARB. 01 Byte 1 is located in the received frame at offset (B1OFFSET – 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B1OFFSET less than 8 are reserved in FIFO modes. 10 Byte 0 is located in the received frame at offset B1OFFSET bytes from the byte after the last byte of the layer 2 header. 11 Byte 0 is located in the received frame at offset B1OFFSET bytes from the byte after the last byte of the layer 3 header. |
| 10–15 | B1OFFSET | Offset relative to the header defined by B1CTL that locates byte 1 of property ARB. An effective offset of zero points to the first byte of the specified header. |
| 16–17 | B2CTL | Location of byte 2 of property ARB. 00 Byte 2 is not extracted, and appears as zero in property ARB. 01 Byte 2 is located in the received frame at offset (B2OFFSET – 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B2OFFSET less than 8 are reserved in FIFO modes. 10 Byte 0 is located in the received frame at offset B2OFFSET bytes from the byte after the last byte of the layer 2 header. 11 Byte 0 is located in the received frame at offset B2OFFSET bytes from the byte after the last byte of the layer 3 header. |
| 18–23 | B2OFFSET | Offset relative to the header defined by B2CTL that locates byte 2 of property ARB. An effective offset of zero points to the first byte of the specified header. |
| 24–25 | B3CTL | Location of byte 3 of property ARB. 00 Byte 3 is not extracted, and appears as zero in property ARB. 01 Byte 3 is located in the received frame at offset (B3OFFSET – 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B3OFFSET less than 8 are reserved in FIFO modes. 10 Byte 0 is located in the received frame at offset B3OFFSET bytes from the byte after the last byte of the layer 2 header. 11 Byte 0 is located in the received frame at offset B3OFFSET bytes from the byte after the last byte of the layer 3 header. |
| 26–31 | B3OFFSET | Offset relative to the header defined by B3CTL that locates byte 3 of property ARB. An effective offset of zero points to the first byte of the specified header. |

19.5.3.3.6 Receive Queue Filer Table Address Register (RQFAR)

RQFAR, shown in Figure 19-27, contains the index of the current, indirectly accessible entry of the received queue filer table. Each table entry occupies a pair of 32-bit words, denoted RQCTRL and RQPROP. To access the RQCTRL and RQPROP words of entry n , write n to RQFAR. Then read or write the indexed RQCTRL and RQPROP words by reading or writing the RQFCR and RQFPR registers, respectively.




Figure 19-27. Receive Queue Filer Table Address Register Definition

Table 19-32 describes the fields of the RQFAR register.

Table 19-32. RQFAR Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0–23 | — | Reserved |
| 24–31 | RQFAR | Current index of receive queue filer table, which spans a total of 256 entries. |

19.5.3.3.7 Receive Queue Filer Table Control Register (RQFCR)

RQFCR is accessed to read or write the RQCTRL words in entries of the receive queue filer table. The table entries are described in greater detail in Section 19.6.4.2, “Receive Queue Filer.” The word accessed through RQFCR is defined by the current value of RQFAR.

Figure 19-28 describes the definition for the RQFCR register.




Figure 19-28. Receive Queue Filer Table Control Register Definition

Table 19-33 describes the fields of the RQFCR register.

Table 19-33. RQFCR Field Descriptions

| Bit | Name | Description |
|------|------|---|
| 0 | GPI | General purpose interrupt. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the filer will instruct the Rx descriptor controller to set IEVENT[FGPI] when the corresponding receive frame is written to memory. If the timer is enabled (TMR_CTRL[TE] = 1), then TMR_PEVENT[RXP] will also be set. |
| 1–15 | — | Reserved, should be written with zero. |

Table 19-33. RQFCR Field Descriptions (continued)

| Bit | Name | Description |
|-------|------|--|
| 16–21 | Q | Receive queue index, from 0 to 63, inclusive, written into the Rx frame control block associated with the received frame. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the frame is sent to either RxBD ring 0 (if RCTRL[FSQEN] = 1) or the RxBD ring with index (Q mod 8) and the filing table search is terminated. In the case where RCTRL[FSQEN] = 0, 8 virtual receive queues are overlaid on every RxBD ring, and software needs to consult the RQ field of the Rx frame control block to determine which virtual receive queue was chosen. |
| 22 | CLE | Cluster entry/exit (used in combination with AND bit). This bit brackets clusters, marking the start and end entries of a cluster. Clusters cannot be nested. 0 Regular RQCTRL entry. 1 If entry matches and AND = 1, treat subsequent entries as belonging to a nested cluster and enter the cluster; otherwise skip all entries up to and including the next cluster exit. If AND = 0, exit current cluster. |
| 23 | REJ | Reject frame. This bit and its specified action are ignored if AND = 1. 0 If entry matches, accept frame and file it to RxBD ring Q. 1 If entry matches, reject frame and discard it, ignoring Q. |
| 24 | AND | Match this entry and the next entry as a pair. 0 Match property[PID] against RQPROP, independent of the next entry. 1 Match property[PID] against RQPROP. If matched and CLE = 0, attempt to match next entry, otherwise, skip all entries up to and including the entry with AND = 0. If matched and CLE = 1, enter cluster of entries, otherwise, skip all entries up to and including the entry with CLE = 1 (cluster exit). |
| 25–26 | CMP | Comparison operation to perform on the RQPROP entry at this index when PID > 0. The property value extracted by the frame parser is masked by the 32-bit <i>mask_register</i> prior to comparison against RQPROP. However, the property value is not permanently altered by the value in <i>mask_register</i> . By default, <i>mask_register</i> is initialized to 0xFFFF_FFFF before each frame is processed. In the case where PID = 0, CMP is interpreted as follows: 00/01 Filer <i>mask_register</i> is set to all 32 bits of RQPROP, and this entry always <i>matches</i> . 10/11 Filer <i>mask_register</i> is set to all 32 bits of RQPROP, and this entry always <i>fails to match</i> . In the case where PID > 0, CMP is interpreted as follows (& is bit-wise AND operator): 00 <i>property</i> [PID] & <i>mask_register</i> = RQPROP 01 <i>property</i> [PID] & <i>mask_register</i> >= RQPROP 10 <i>property</i> [PID] & <i>mask_register</i> != RQPROP 11 <i>property</i> [PID] & <i>mask_register</i> < RQPROP |
| 27 | — | Reserved, should be written with zero. |
| 28–31 | PID | Property identifier. The value in the RQFPR[RQPROP] entry at this index is interpreted according to PID (see Table 19-34). |

19.5.3.3.8 Receive Queue Filer Table Property Register (RQFPR)

RQFPR (see [Figure 19-29](#)) is accessed to read or write the RQPROP words in entries of the receive queue filer table. The table entries are described in greater detail in [Section 19.6.4.2, “Receive Queue Filer.”](#) The word accessed through RQFPR is defined by the current value of RQFAR. [Figure 19-29](#) and [Figure 19-30](#) describe the fields of the RQFPR register according to property ID, designated in RQFCR[PID].




Figure 19-29. Receive Queue Filer Table Property IDs 0, 2–15 Register Definition




Figure 19-30. Receive Queue Filer Table Property ID1 Register Definition

Table 19-34 describes the fields of the RQFPR register.

Table 19-34. RQFPR Field Descriptions

| PID ¹ | Bit | Name | Description |
|------------------|------|------|---|
| 0000 | 0–31 | MASK | Mask bits to be written to Filer <i>mask_register</i> for masking of property values. The rule match/fail status for this PID is determined by RQCTRL[<i>CMP</i>]. Since <i>mask_register</i> is bit-wise ANDed with properties, every bit of MASK that is cleared also results in the corresponding property bit being cleared in comparisons. Therefore setting MASK to 0xFFFF_FFFF ensures that all property bits participate in rule matches. |

Table 19-34. RQFPR Field Descriptions (continued)

| PID ¹ | Bit | Name | Description |
|------------------|-------|--|--|
| 0001 | 0–15 | — | Reserved |
| | 16 | EBC | Set if the destination Ethernet address is to the broadcast address. |
| | 17 | VLN | Set if a VLAN tag (Ethertype DFVLAN[TAG] or 0x8100) was seen in the frame. |
| | 18 | CFI | Set to the value of the Canonical Format Indicator in the VLAN control tag if VLAN is set, zero otherwise. |
| | 19 | JUM | Set if a jumbo Ethernet frame was parsed. |
| | 20 | IPF | Set if a fragmented IPv4 or IPv6 header was encountered. See the descriptions of receive FCB fields IP and PRO in Section 19.6.3.3, “Receive Path Off-Load,” for more information on determining the status of received packets for which IPF is set. |
| | 21 | — | Reserved |
| | 22 | IP4 | Set if an IPv4 header was parsed. |
| | 23 | IP6 | Set if an IPv6 header was parsed. |
| | 24 | ICC | Set if the IPv4 header checksum was checked. |
| | 25 | ICV | Set if the IPv4 header checksum was verified correct. |
| | 26 | TCP | Set if a TCP header was parsed. |
| | 27 | UDP | Set if a UDP header was parsed. |
| | 28–29 | — | Reserved. |
| | 0010 | 0–7 | ARB |
| 8–15 | | User-defined arbitrary bit field property: byte 1 extracted. Defaults to 0x00. | |
| 16–23 | | User-defined arbitrary bit field property: byte 2 extracted. Defaults to 0x00. | |
| 24–31 | | User-defined arbitrary bit field property: byte 3 extracted. Defaults to 0x00. | |
| 0011 | 0–7 | — | Reserved, should be written with zero. |
| | 8–31 | DAH | Destination MAC address, most significant 24 bits. Defaults to 0x000000. |
| 0100 | 0–7 | — | Reserved, should be written with zero. |
| | 8–31 | DAL | Destination MAC address, least significant 24 bits. Defaults to 0x000000. |
| 0101 | 0–7 | — | Reserved, should be written with zero. |
| | 8–31 | SAH | Source MAC address, most significant 24 bits. Defaults to 0x000000. |
| 0110 | 0–7 | — | Reserved, should be written with zero. |
| | 8–31 | SAL | Source MAC address, least significant 24 bits. Defaults to 0x000000. |

Table 19-34. RQFPR Field Descriptions (continued)

| PID ¹ | Bit | Name | Description |
|------------------|-------|------|--|
| 0111 | 0–15 | — | Reserved, should be written with zero. |
| | 16–31 | ETY | <p>Ethertype of next layer protocol, that is, last ethertype if layer 2 headers nest. Defaults to 0xFFFF. Using the filer to match ETY does not work in the case of PPPoE packets, because the PPPoE ethertype in the original packet, 0x8864, is always overwritten with the PPP protocol field. Thus, matches on ETY == 0x8864 always fail.</p> <p>Instead, software should use PID=1 fields IP4 (ETY = 0x0021) and IP6 (ETY = 0x0057) to distinguish PPPoE session packets carrying IPv4 and IPv6 datagrams. Other PPP protocols are encoded in the ETY field, but many of them overlap with real ethertype definitions. Consult IANA and IEEE for possible ambiguities.</p> <p>Packets with a value in the length/type field greater than 1500 and less than 1536 are treated as payload length. If the eTSEC is used in a network where there are packets carrying a type designation between 1500 and 1536 (note there are none currently publicly defined by IANA), then the S/W must confirm the parser and filer results by checking the type/length field after the packet has been written to memory to see if it falls in this range.</p> <p>Note that the eTSEC filer gets multiple packet attributes as a result of parsing the packet. The behavior of the eTSEC is that it pulls the innermost ethertype found in the packet; this means that in many supported protocols that have inner etherypes, in order to file based on the outer ethertype, arbitrary extraction should be used instead of the ETY PID. There are four cases that need to be highlighted.</p> <ol style="list-style-type: none"> 1. The jumbo ethertype (0x8870)—In this case, the eTSEC assumes that the following header is LLC/SNAP. LLC/SNAP has an associated Ethertype, and the ETY field is populated with that ethertype. This makes it impossible to file on jumbo frames. In this case, one can use arbitrary extracted bytes to pull the outermost Ethertype. 2. The PPPoE ethertype described above. 3. The VLAN tag ethertype (0x8100)—In this case, one can use the PID=1 VLN bit to indicate that the packet had a VLAN tag. 4. The MPLS tagged packets. In this case, one can use arbitrary extraction bytes to compare to the actual ethertype if a filer rule is intending to file based on an MPLS label existence. |
| 1000 | 0–19 | — | Reserved, should be written with zero. |
| | 20–31 | VID | VLAN network identifier (as per IEEE Std 802.1Q). This value defaults to 0x000 if no VLAN tag was found, or the VLAN tag contained only priority information. |
| 1001 | 0–28 | — | Reserved, should be written with zero. |
| | 29–31 | PRI | VLAN user priority (as per IEEE Std 802.1p). This value defaults to 000 (best effort priority) if no VLAN tag was found. |
| 1010 | 0–23 | — | Reserved, should be written with zero. |
| | 24–31 | TOS | IPv4 header Type Of Service field or IPv6 Traffic Class field. This value defaults to 0x00 (default RFC 2474 best-effort behavior) if no IP header appeared. Note that for IPv6 the Traffic Class field is extracted using the IP header definition in RFC 2460. IPv6 headers formed using the earlier RFC 1883 have a different format and must be handled with software. |
| 1011 | 0–23 | — | Reserved, should be written with zero. |
| | 24–31 | L4P | Layer 4 protocol identifier as per published IANA specification. This is the last recognized protocol type recognized in the case of IPv6 extension headers. This value defaults to 0xFF to indicate that no layer 4 header was recognized (possibly due to absence of an IP header). |
| 1100 | 0–31 | DIA | Destination IP address. If an IPv4 header was found, this is the entire destination address. If an IPv6 header was found, this is the 32 most significant bits of the 128-bit destination address. This value defaults to all zeros if no IP header appeared. |

Table 19-34. RQFPR Field Descriptions (continued)

| PID ¹ | Bit | Name | Description |
|------------------|-------|------|--|
| 1101 | 0–31 | SIA | Source IP address. If an IPv4 header was found, this is the entire source address. If an IPv6 header was found, this is the 32 most significant bits of the 128-bit source address. This value defaults to all zeros if no IP header appeared. |
| 1110 | 0–15 | — | Reserved, should be written with zero. |
| | 16–31 | DPT | Destination port number for TCP or UDP headers. This value defaults to 0x0000 if no TCP or UDP headers were recognized. |
| 1111 | 0–15 | — | Reserved, should be written with zero. |
| | 16–31 | SPT | Source port number for TCP or UDP headers. This value defaults to 0x0000 if no TCP or UDP headers were recognized. |

¹ PID is the property identifier field of the filer table control entry (see RQFCR[PID]) at the same index.

19.5.3.3.9 Maximum Receive Buffer Length Register (MRBLR)

The MRBLR register is written by the user. It informs the eTSEC how much space is in the receive buffer pointed to by the RxBD. [Figure 19-31](#) describes the definition for the MRBLR.




Figure 19-31. MRBLR Register Definition

Table 19-35. MRBLR Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–25 | MRBL | Maximum receive buffer length. MRBL is the number of bytes that the eTSEC receiver writes to the receive buffer. The MRBL register is written by the user with a multiple of 64 for all modes. The eTSEC can write fewer bytes to the buffer than the value set in MRBL if a condition such as an error or end-of-frame occurs, but it never exceeds the MRBL value; therefore, user-supplied buffers must be at least as large as the MRBL. MRBL must be set, together with the number of buffer descriptors, to ensure adequate space for received frames. See Section 19.5.3.5.5, “Maximum Frame Length Register (MAXFRM),” for further discussion. |
| 26–31 | — | To ensure that MRBL is a multiple of 64, these bits are reserved and should be cleared. |

19.5.3.3.10 Receive Buffer Descriptor Pointers 0–7 (RBPTR0–RBPTR7)

RBPTR0–RBPTR7 each contains the low-order 32 bits of the next receive buffer descriptor address for their respective RxBD ring. [Figure 19-32](#) describes the RBPTR registers. These registers takes on the value of their ring’s associated RBASE when the RBASE register is written by software. Software must not write RBPTR_n while eTSEC is actively receiving frames. However, RBPTR_n can be modified when the receiver is disabled or when no Rx buffer is in use (after a GRACEFUL STOP RECEIVE command is issued and the frame completes its reception) in order to change the next RxBD eTSEC receives.




Figure 19-32. RBPTR0–RBPTR7 Register Definition

Table 19-36 describes the fields of the RBPTR_n register.

Table 19-36. RBPTR_n Field Descriptions

| Bits | Name | Description |
|-------|--------------------|---|
| 0–28 | RBPTR _n | Current RxBD pointer for RxBD ring <i>n</i> . Points to the current BD being processed or to the next BD the receiver uses when it is idling. After reset or when the end of the RxBD ring is reached, eTSEC initializes RBPTR _n to the value in the corresponding RBASE _n . The RBPTR register is internally written by the eTSEC’s DMA controller during reception. The pointer increments by 8 (bytes) each time a descriptor is closed successfully by the eTSEC. Note that the 3 least-significant bits of this register are read only and zero. |
| 29–31 | — | Reserved |

19.5.3.3.11 Receive Descriptor Base Address Registers (RBASE0–RBASE7)

The RBASE_n registers are written by the user with the base address of each RxBD ring *n*. Each such value must be divisible by eight, since the 3 least-significant bits always write as 000. Figure 19-33 describes the RBASE_n registers.




Figure 19-33. RBASE Register Definition

Table 19-37 describes the fields of the RBASE_n registers.

Table 19-37. RBASE0–RBASE7 Field Descriptions

| Bits | Name | Description |
|-------|--------------------|---|
| 0–28 | RBASE _n | Receive base for ring <i>n</i> . RBASE defines the starting location in the memory map for the eTSEC RxBDs. This field must be 8-byte aligned. Together with setting the W (wrap) bit in the last BD, the user can select how many BDs to allocate for the receive packets. The user must initialize RBASE before enabling the eTSEC receive function on the associated ring. |
| 29–31 | — | Reserved |

19.5.3.3.12 Receive Stamp Register (TMR_RXTS_H/L)

Receive time stamp register (RXTS_H/L). This register holds the value present in TMR_CNT_H/L when the eTSEC detects a new incoming Ethernet frame. This register is only updated when the precision time

stamp logic is enable via TMR_CTRL[TE]. This register is read only in normal operation. Figure 19-34 describes the definition for the RXTS_H/L register.




Figure 19-34. TMR_RXTS_H/L Register Definition

Table 19-38 describes the fields of the TMR_RXTS_H/L register.

Table 19-38. TMR_RXTS_H/L Register Field Descriptions

| Bits | Name | Description |
|------|--------------|--|
| 0–63 | TMR_RXTS_H/L | Value of the eTSEC precision timer upon detection of a start of frame symbol for the received frame. |

19.5.3.4 MAC Functionality

This section describes the MAC registers and provides a brief overview of the functionality that can be exercised through the use of these registers, particularly those that provide functionality not explicitly required by the IEEE 802.3 standard. All of the MAC registers are 32 bits wide.

19.5.3.4.1 Configuring the MAC

The MAC configuration registers 1 and 2 provide for configuring the MAC in multiple ways:

- Adjusting the preamble length—The length of the preamble can be adjusted from the nominal seven bytes to some other (non-zero) value. Should custom preamble insertion/extraction be configured, then this register must be left at its default value.
- Varying pad/CRC combinations—Three different pad/CRC combinations are provided to handle a variety of system requirements. Simplest are frames that already have a valid frame check sequence (FCS) field. The other two options include appending a valid CRC or padding and then appending a valid CRC, resulting in a minimum frame of 64 octets. In addition to the programmable register set, the pad/CRC behavior can be dynamically adjusted on a per-packet basis.

19.5.3.4.2 Controlling CSMA/CD

The half-duplex register (HAFDUP) allows control over the carrier-sense multiple access/collision detection (CSMA/CD) logic of the eTSEC. Half-duplex mode is only supported for 10- and 100-Mbps operation. Following the completion of the packet transmission the part begins timing the inter packet gap (IPG) as programmed in the back-to-back IPG configuration register. The system is now free to begin another frame transfer.

In full-duplex mode both the carrier sense (CRS) and collision (COL) indications from the PHY are ignored, but in half-duplex mode the eTSEC defers to CRS, and following a carrier event, times the IPG using the non-back-to-back IPG configuration values that include support for the optional

two-thirds/one-third CRS deferral process. This optional IPG mechanism enhances system robustness and ensures fair access to the medium. During the first two-thirds of the IPG, the IPG timer is cleared if CRS is sensed. During the final one-third of the IPG, CRS is ignored and the transmission begins once IPG is timed. The two-thirds/one-third ratio is the recommended value.

19.5.3.4.3 Handling Packet Collisions

While transmitting a packet in half-duplex mode, the eTSEC is sensitive to COL. If a collision occurs, it aborts the packet and outputs the 32-bit jam sequence. The jam sequence is comprised of several bits of the CRC, inverted to guarantee an invalid CRC upon reception. A signal is sent to the system indicating that a collision occurred and that the start of the frame is needed for retransmission. The eTSEC then backs off of the medium for a time determined by the truncated binary exponential back off (BEB) algorithm. Following this back-off time, the packet is retried. The back-off time can be skipped if configured through the half-duplex register. However, this is non-standard behavior and its use must be carefully applied. Should any one packet experience excessive collisions, the packet is aborted. The system should flush the frame and move to the next one in line. If the system requests to send a packet while the eTSEC is deferring to a carrier, the eTSEC simply waits until the end of the carrier event and the timing of IPG before it honors the request.

If packet transmission attempts experience collisions, the eTSEC outputs the jam sequence and waits some amount of time before retrying the packet. This amount of time is determined by a controlled randomization process called truncated binary exponential back-off. The amount of time is an integer number of slot times. The number of slot times to delay before the n th retransmission attempt is chosen as a uniformly distributed random integer r in the range:

$$0 \leq r \leq 2^k, \text{ where } k = \min(n, 10).$$

So after the first collision, the eTSEC backs off either 0 or 1 slot times. After the fifth collision, the eTSEC backs off between 0 and 32 slot times. After the tenth collision, the maximum number of slot times to back off is 1024. This can be adjusted through the half-duplex register. An alternate truncation point, such as 7 for instance, can be programmed. On average, the MAC is more aggressive after seven collisions than other stations on the network.

19.5.3.4.4 Controlling Packet Flow

Packet flow can be dealt with in a number of ways within eTSEC. A default retransmit attempt limit of 15 can be reduced using the half-duplex register. The slot time or collision window can be used to gate the retry window and possibly reduce the amount of transmit buffering within the system. The slot time for 10/100 Mbps is 512 bit times. Because the slot time begins at the beginning of the packet (including preamble), the end occurs around the 56th byte of the frame data. Slot time in 1000-Mbps mode is not supported.

Full-duplex flow control is provided for in IEEE 802.3x. Currently the standard does not address flow control in half-duplex environments. Common in the industry, however, is the concept of back pressure. The eTSEC implements the optional back pressure mechanism using the raise carrier method. If the system receive logic wishes to stop the reception of packets in a network-friendly way, transmit half-duplex flow control (THDF) is set (TCTRL[THDF]). If the medium is idle, the eTSEC raises carrier by transmitting preamble. Other stations on the half-duplex network then defer to the carrier.

In the event the preamble transmission happens to cause a collision, the eTSEC ensures the minimum 96-bit presence on the wire, then drops preamble and waits a back-off time depending on the value of the back-pressure-no-back-off configuration bit HAFDUP[BP No BackOff]. These transmitting-preamble-for-back pressure collisions are not counted. If HAFDUP[BP No BackOff] is set, the eTSEC waits an inter-packet gap before resuming the transmission of preamble following the collision and does not defer. If HAFDUP[BP No BackOff] is cleared, the eTSEC adheres to the truncated BEB algorithm that allows the possibility of packets being received. This also can be detrimental in that packets can now experience excessive collisions, causing them to be dropped in the stations from which they originate. To reduce the likelihood of lost packets and packets leaking through the back pressure mechanism, HAFDUP[BP No BackOff] must be set.

The eTSEC drops carrier (cease transmitting preamble) periodically to avoid excessive defer conditions in other stations on the shared network. If, while applying back pressure, the eTSEC is requested to send a packet, it stops sending preamble, and waits one IPG before sending the packet. HAFDUP[BP No BackOff] applies for any collision that occurs during the sending of this packet. Collisions for packets while half duplex back pressure is asserted are counted. The eTSEC does not defer while attempting to send packets while in back pressure. Again, back pressure is non-standard, yet it can be effective in reducing the flow of receive packets.

19.5.3.4.5 Controlling PHY Links

Control and status to and from the PHY is provided through the two-wire MII management interface described in IEEE 802.3u. The MII management registers (MII management configuration, command, address, control, status, and indicator registers) are used to exercise this interface between a host processor and one or more PHY devices.

The eTSEC MII's registers provide the ability to perform continuous read cycles (called a scan cycle); although, scan cycles are not explicitly defined in the standard. If requested (by setting MIIMCOM[Scan Cycle]), the part performs repetitive read cycles of the PHY status register, for example. In this way, link characteristics may be monitored more efficiently. The different fields in the MII management indicator register (scan, not valid and busy) are used to indicate availability of each read of the scan cycle to the host from MIIMSTAT[PHY scan].

Yet another parameter that can be modified through the MII registers is the length of the MII management interface preamble. After establishing that a PHY supports preamble suppression, the host may so configure the eTSEC. While enabled, the length of MII management frames are reduced from 64 clocks to 32 clocks. This effectively doubles the efficiency of the interface.

19.5.3.5 MAC Registers

This section describes the MAC registers.

19.5.3.5.1 MAC Configuration 1 Register (MACCFG1)

MACCFG1 is written by the user. [Figure 19-35](#) describes the definition for the MACCFG1 register.

Offset eTSEC1:0x2_4500; eTSEC2:0x2_5500 Access: Mixed




Figure 19-35. MACCFG1 Register Definition

[Table 19-39](#) describes the fields of the MACCFG1 register.

Table 19-39. MACCFG1 Field Descriptions

| Bits | Name | Description |
|-------|--------------|--|
| 0 | Soft_Reset | Soft reset. This bit is cleared by default. See Section 19.6.2.2, “Soft Reset and Reconfiguring Procedure,” for more information on setting this bit. 0 Normal operation. 1 Place the entire MAC in reset except for the host interface. |
| 1–11 | — | Reserved |
| 12 | Reset Rx MC | Reset receive MAC control block. This bit is cleared by default. 0 Normal operation. 1 Place the receive part of the MAC in reset. This block detects control frames and contains the pause timers. |
| 13 | Reset Tx MC | Reset transmit MAC control block. This bit is cleared by default. 0 Normal operation. 1 Place the transmit part of the MAC in reset. This block multiplexes data and control frame transfers. It also responds to XOFF PAUSE control frames. |
| 14 | Reset Rx Fun | Reset receive function block. This bit is cleared by default. 0 Normal operation. 1 Place the receive function in reset. This block performs the receive frame protocol. |
| 15 | Reset Tx Fun | Reset transmit function block. This bit is cleared by default. 0 Normal operation. 1 Place the transmit function in reset. This block performs the frame transmission protocol. |
| 16–22 | — | Reserved |
| 23 | Loop Back | Loop back. This bit is cleared by default. 0 Normal operation. 1 Loop back the MAC transmit outputs to the MAC receive inputs. |

Table 19-39. MACCFG1 Field Descriptions (continued)

| Bits | Name | Description |
|-------|--------------|--|
| 24–25 | — | Reserved |
| 26 | Rx_Flow | Receive flow. This bit is cleared by default. Must be 0 if MACCFG2[Full Duplex] = 0. 0 The receive MAC control ignores PAUSE flow control frames. 1 The receive MAC control detects and acts on PAUSE flow control frames. Note: Should not be set when operating in Half-Duplex mode |
| 27 | Tx_Flow | Transmit flow. This bit is cleared by default. Must be 0 if MACCFG2[Full Duplex] = 0. 0 The transmit MAC control may not send PAUSE flow control frames if requested by the system. Note: 1The transmit MAC control may send PAUSE flow control frames if requested by the system.Should not be set when operating in Half-Duplex mode |
| 28 | Sync'd Rx EN | Receive enable synchronized to the receive stream. (Read-only) 0 Frame reception is not enabled. 1 Frame reception is enabled. |
| 29 | Rx_EN | Receive enable. This bit is cleared by default. If set, prior to clearing this bit, set DMACTRL[GRS] then confirm subsequent occurrence of the graceful receive stop interrupt (IEVENT[GRSC] is set). 0 The MAC may not receive frames from the PHY. 1 The MAC may receive frames from the PHY. |
| 30 | Sync'd Tx EN | Transmit enable synchronized to the transmit stream. (Read-only) 0 Frame transmission is not enabled. 1 Frame transmission is enabled. |
| 31 | Tx_EN | Transmit enable. This bit is cleared by default. If set, prior to clearing this bit, set DMACTRL[GTS] then confirm subsequent occurrence of the graceful transmit stop interrupt (IEVENT[GTSC] is set). 0 The MAC may not transmit frames from the system. 1 The MAC may transmit frames from the system. |

19.5.3.5.2 MAC Configuration 2 Register (MACCFG2)

The MACCFG2 register is written by the user. [Figure 19-36](#) describes the definition for the MACCFG2 register.

Offset eTSEC1:0x2_4504; eTSEC2:0x2_5504

Access: Read/Write




Figure 19-36. MACCFG2 Register Definition

Table 19-40 describes the fields of the MACCFG2 register.

Table 19-40. MACCFG2 Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | | | | | | | | | |
|---------------------|------------------------|--|---------------------------|--------------|-------------------|---------------------------|---------------------|------------------------|-----|-----|---------|------------------------|----|----|----------|------------------------|----|-----|---------------------|------------------------|----|----|
| 0–15 | — | Reserved | | | | | | | | | | | | | | | | | | | | |
| 16–19 | Preamble Length | This field determines the length in bytes of the preamble field preceding each Ethernet start-of-frame delimiter byte. Values from 0x3 to 0xF are supported by the controller. The default value of 0x7 should not be altered in order to guarantee reliable operation with IEEE 802.3 compliant hardware. | | | | | | | | | | | | | | | | | | | | |
| 20–21 | — | Reserved | | | | | | | | | | | | | | | | | | | | |
| 22–23 | I/F Mode | This field determines the type of interface to which the MAC is connected. Its default is 00. 00 Reserved bit mode (not supported) (10 Mbps GENDEC/GPSI) 01 Nibble mode (MII) (10/100 Mbps MII/RMII) 10 Byte mode (1000 Mbps). Reserved if neither GMII or TBI are supported. 11 Reserved | | | | | | | | | | | | | | | | | | | | |
| 24 | PreAM RxEN | User defined preamble enable for received frames. This bit is cleared by default. 0 The MAC skips the Ethernet preamble without returning it. 1 The MAC recovers the received Ethernet preamble and passes it to the driver at the start of each received frame. If the preamble is less than 7 bytes, 0's are prepended to pad it to 7 bytes. Not applicable to or RMII 10/100 modes. | | | | | | | | | | | | | | | | | | | | |
| 25 | PreAM TxEN | User defined preamble enable for transmitted frames. This bit is cleared by default. 0 The MAC generates a standard Ethernet preamble. 1 If a user-defined preamble has been passed to the MAC it is transmitted instead of the standard preamble. Otherwise the standard Ethernet preamble is generated. The Preamble Length field should be left at its default setting if a user-defined preamble is transmitted. Not applicable to or RMII 10/100 modes. | | | | | | | | | | | | | | | | | | | | |
| 26 | Huge Frame | Huge frame enable. This bit is cleared by default. 0 Limit the length of frames received to less than or equal to the maximum frame length value (MAXFRM[Maximum Frame]) and limit the length of frames transmitted to less than the maximum frame length. See Section 19.6.7, "Buffer Descriptors," for further details of buffer descriptor bit updating. <table border="1" data-bbox="451 1213 1446 1495"> <thead> <tr> <th>Frame type</th> <th>Frame length</th> <th>Packet truncation</th> <th>Buffer descriptor updated</th> </tr> </thead> <tbody> <tr> <td>Receive or transmit</td> <td>> maximum frame length</td> <td>yes</td> <td>yes</td> </tr> <tr> <td>Receive</td> <td>= maximum frame length</td> <td>no</td> <td>no</td> </tr> <tr> <td>Transmit</td> <td>= maximum frame length</td> <td>no</td> <td>yes</td> </tr> <tr> <td>Receive or transmit</td> <td>< maximum frame length</td> <td>no</td> <td>no</td> </tr> </tbody> </table> 1 Frames are transmitted and received regardless of their relationship to the maximum frame length. Note that if Huge Frame is cleared, the user must ensure that adequate buffer space is allocated for received frames. See Section 19.5.3.5.5, "Maximum Frame Length Register (MAXFRM)," for further information. | Frame type | Frame length | Packet truncation | Buffer descriptor updated | Receive or transmit | > maximum frame length | yes | yes | Receive | = maximum frame length | no | no | Transmit | = maximum frame length | no | yes | Receive or transmit | < maximum frame length | no | no |
| Frame type | Frame length | Packet truncation | Buffer descriptor updated | | | | | | | | | | | | | | | | | | | |
| Receive or transmit | > maximum frame length | yes | yes | | | | | | | | | | | | | | | | | | | |
| Receive | = maximum frame length | no | no | | | | | | | | | | | | | | | | | | | |
| Transmit | = maximum frame length | no | yes | | | | | | | | | | | | | | | | | | | |
| Receive or transmit | < maximum frame length | no | no | | | | | | | | | | | | | | | | | | | |
| 27 | Length check | Length check. This bit is cleared by default. 0 No length field checking is performed. 1 The MAC checks the frame's length field on receive to ensure it matches the actual data field length. Transmitted frames are not checked. | | | | | | | | | | | | | | | | | | | | |

Table 19-40. MACCFG2 Field Descriptions (continued)

| Bits | Name | Description |
|------|-------------|--|
| 28 | MPEN | Magic packet enable for Ethernet modes. This bit is cleared by default. MPEN should be enabled only after GRACEFUL RECEIVE STOP and GRACEFUL TRANSMIT STOP are completed successfully (in other words, transmission and reception have stopped). 0 Normal receive behavior on receive, or Magic Packet mode has exited with reception of a valid Magic Packet. 1 Commence Magic Packet detection by the MAC provided that frame reception is enabled in MACCFG1. In this mode the MAC ignores all received frames until the specific Magic Packet frame is received, at which point this bit is cleared by the eTSEC, and a maskable interrupt through IEVENT[MAG] occurs. |
| 29 | PAD/CRC | Pad and append CRC. This bit is cleared by default. This bit must be set when in half-duplex mode (MACCFG2[Full Duplex] is cleared). 0 Frames presented to the MAC have a valid length and contain a CRC. 1 The MAC pads all transmitted short frames and appends a CRC to every frame regardless of padding requirement. |
| 30 | CRC EN | CRC enable. If the configuration bit PAD/CRC ENABLE or the per-packet PAD/CRC ENABLE is set, CRC ENABLE is ignored. This bit is cleared by default. 0 Frames presented to the MAC have a valid length and contain a valid CRC. 1 The MAC appends a CRC on all frames. Clear this bit if frames presented to the MAC have a valid length and contain a valid CRC. |
| 31 | Full Duplex | Full duplex configure. This bit is cleared by default. 0 The MAC operates in half-duplex mode only. 1 The MAC operates in full-duplex mode. |

19.5.3.5.3 Inter-Packet Gap/Inter-Frame Gap Register (IPGIFG)

The IPGIFG register is written by the user. [Figure 19-37](#) describes the definition for IPGIFG.




Figure 19-37. IPGIFG Register Definition

[Table 19-41](#) describes the fields of the IPGIFG register.

Table 19-41. IPGIFG Field Descriptions

| Bits | Name | Description |
|------|---|--|
| 0 | — | Reserved |
| 1–7 | Non-Back-to-Back Inter-Packet-Gap, Part 1 | This is a programmable field representing the optional carrier sense window referenced in IEEE 802.3/4.2.3.2.1 ‘carrier deference’. If carrier is detected during the timing of IPGR1, the MAC defers to carrier. If, however, carrier becomes active after IPGR1, the MAC continues timing IPGR2 and transmits, knowingly causing a collision, thus ensuring fair access to medium. Its range of values is 0x00 to IPGR2. Its default is 0x40 (64d) which follows the two-thirds/one-third guideline. |
| 8 | — | Reserved |

Table 19-41. IPGIFG Field Descriptions (continued)

| Bits | Name | Description |
|-------|---|---|
| 9–15 | Non-Back-to-Back Inter-Packet-Gap, Part 2 | This is a programmable field representing the non-back-to-back inter-packet-gap in bits. Its default is 0x60 (96d), which represents the minimum IPG of 96 bits. |
| 16–23 | Minimum IFG Enforcement | This is a programmable field representing the minimum number of bits of IFG to enforce between frames. A frame is dropped whose IFG is less than that programmed. The default setting of 0x50 (80d) represents half of the nominal minimum IFG which is 160 bits. |
| 24 | — | Reserved |
| 25–31 | Back-to-Back Inter-Packet-Gap | This is a programmable field representing the IPG between back-to-back packets. This is the IPG parameter used exclusively in full-duplex mode and in half-duplex mode if two transmit packets are sent back-to-back. Set this field to the number of bits of IPG desired. The default setting of 0x60 (96d) represents the minimum IPG of 96 bits. |

19.5.3.5.4 Half-Duplex Register (HAFDUP)

The HAFDUP register is written by the user. [Figure 19-38](#) describes the HAFDUP register.

Offset eTSEC1:0x2_450C; eTSEC2:0x2_550C

Access: Read/Write




Figure 19-38. Half-Duplex Register Definition

[Table 19-42](#) describes the fields of the HAFDUP register.

Table 19-42. HAFDUP Field Descriptions

| Bits | Name | Description |
|------|--------------------------|--|
| 0–7 | — | Reserved |
| 8–11 | Alternate BEB Truncation | This field is used while ALTERNATE BINARY EXPONENTIAL BACKOFF ENABLE is set. The value programmed is substituted for the Ethernet standard value of ten. Its default is 0xA. |
| 12 | Alt BEB | Alternate binary exponential backoff. This bit is cleared by default. 0 The Tx MAC follows the standard binary exponential back off rule. 1 The Tx MAC uses the ALTERNATE BINARY EXPONENTIAL BACKOFF TRUNCATION setting instead of the 802.3 standard tenth collision. The standard specifies that any collision after the tenth uses one less than 210 as the maximum backoff time. |
| 13 | BP No BackOff | Back pressure no backoff. This bit is cleared by default. 0 The Tx MAC follows the binary exponential back off rule. 1 The Tx MAC immediately re-transmits, following a collision, during back pressure operation. |

MIIM registers for that eTSEC. For example, if a RTBI interface is required on eTSEC2, then the MIIM registers starting at offset 0x2_5520 are used to configure it.

Figure 19-40 describes the definition for the MIIMCFG register.




Figure 19-40. MII Management Configuration Register Definition

Table 19-44 describes the fields of the MIIMCFG register.

Table 19-44. MIIMCFG Field Descriptions

| Bits | Name | Description |
|-------|------------|--|
| 0 | Reset Mgmt | Reset management. This bit is cleared by default. 0 Allow the MII MGMT to perform mgmt read/write cycles if requested through the host interface. 1 Reset the MII MGMT. |
| 1–26 | — | Reserved |
| 27 | No Pre | Preamble suppress. This bit is cleared by default. 0 The MII MGMT performs Mgmt read/write cycles with 32 clocks of preamble. 1 The MII MGMT suppresses preamble generation and reduces the Mgmt cycle from 64 clocks to 32 clocks. This is in accordance with IEEE 802.3/22.2.4.4.2. |
| 28 | — | Reserved |
| 29–31 | MgmtClk | This field determines the clock frequency of the MII management clock (EC_MDC). Its default value is 111. Note: The eTSEC system clock is selected by the SCCR register. (See Chapter 4, “Reset, Clocking, and Initialization.”) 00x 1/4 of the eTSEC system clock divided by 8 010 1/6 of the eTSEC system clock divided by 8 011 1/8 of the eTSEC system clock divided by 8 100 1/10 of the eTSEC system clock divided by 8 101 1/14 of the eTSEC system clock divided by 8 110 1/20 of the eTSEC system clock divided by 8 111 1/28 of the eTSEC system clock divided by 8 |

19.5.3.5.7 MII Management Command Register (MIIMCOM)

The MIIMCOM register is written by the user. Figure 19-41 describes the definition for MIIMCOM.




Figure 19-41. MIIMCOM Register Definition

Table 19-45 describes the fields of the MIIMCOM register.

Table 19-45. MIIMCOM Descriptions

| Bits | Name | Description |
|------|------------|--|
| 0–29 | — | Reserved |
| 30 | Scan Cycle | Scan cycle. This bit is cleared by default. 0 Normal operation. 1 The MII management continuously performs read cycles. This is useful for monitoring link fail, for example. |
| 31 | Read Cycle | Read cycle. This bit is cleared by default but is not self-clearing once set. 0 Normal operation. 1 The MII management performs a single read cycle upon the transition of this bit from 0 to 1 using the PHY address (at MIIMADD[PHY Address]) and the register address (at MIIMADD[Register Address]). The 0-to-1 transition of this bit also causes the MIIMIND[Busy] bit to be set. The read is complete when the MIIMIND[Busy] bit clears. Data is returned in register MIIMSTAT[PHY Status]. |

19.5.3.5.8 MII Management Address Register (MIIMADD)

The MIIMADD register is written by the user. Figure 19-42 shows the MIIMADD register.




Figure 19-42. MIIMADD Register Definition

Table 19-46 describes the fields of the MIIMADD register.

Table 19-46. MIIMADD Field Descriptions

| Bits | Name | Description |
|-------|------------------|--|
| 0–18 | — | Reserved |
| 19–23 | PHY Address | This field represents the 5-bit PHY address field of Mgmt cycles. Up to 31 PHYs can be addressed (0 is reserved). Its default value is 0x00. |
| 24–26 | — | Reserved |
| 27–31 | Register Address | This field represents the 5-bit register address field of Mgmt cycles. Up to 32 registers can be accessed. Its default value is 0x00. |

19.5.3.5.9 MII Management Control Register (MIIMCON)

MIIMCON, shown in [Figure 19-43](#), is written by the user.




Figure 19-43. MII Mgmt Control Register Definition

[Table 19-47](#) describes the fields of the MIIMCON register.

Table 19-47. MIIMCON Field Descriptions

| Bits | Name | Description |
|-------|-------------|---|
| 0–15 | — | Reserved |
| 16–31 | PHY Control | If written, an MII Mgmt write cycle is performed using this 16-bit data, the pre-configured PHY address (at MIIMADD[PHY Address]) and the register address (at MIIMADD[Register Address]). Its default value is 0x0000. |

19.5.3.5.10 MII Management Status Register (MIIMSTAT)

The MIIMSTAT register is read only by the user. [Figure 19-44](#) describes the definition for the MIIMSTAT register.




Figure 19-44. MIIMSTAT Register Definition

[Table 19-48](#) describes the fields of the MIIMSTAT register.

Table 19-48. MIIMSTAT Field Descriptions

| Bits | Name | Description |
|-------|------------|--|
| 0–15 | — | Reserved |
| 16–31 | PHY Status | Following an MII Mgmt read cycle, the 16-bit data can be read from this location. Its default value is 0x0000. |

19.5.3.5.11 MII Management Indicator Register (MIIMIND)

The MIIMIND register is read-only by the user. [Figure 19-45](#) describes the definition for the MIIMIND register.



Figure 19-45. MII Mgmt Indicator Register Definition

Table 19-49. MIIMIND Field Descriptions

| Bits | Name | Description |
|------|-----------|---|
| 0–28 | — | Reserved |
| 29 | Not Valid | Not valid. 0 MII Mgmt read cycle has completed and the read data is valid. 1 MII Mgmt read cycle has not completed and the read data is not yet valid. |
| 30 | Scan | Scan in progress. 0 A scan operation (continuous MII Mgmt read cycles) is not in progress. 1 A scan operation (continuous MII Mgmt read cycles) is in progress. |
| 31 | Busy | Busy. 0 MII Mgmt block is not currently performing an MII Mgmt read or write cycle. 1 MII Mgmt block is currently performing an MII Mgmt read or write cycle. |

19.5.3.5.12 Interface Status Register (IFSTAT)

[Figure 19-46](#) shows the IFSTAT register.




Figure 19-46. Interface Status Register Definition

[Table 19-50](#) describes the fields of the FSTAT register.

Table 19-50. IFSTAT Field Descriptions

| Bits | Name | Description |
|------|------|-------------|
| 0–21 | — | Reserved |

Table 19-50. IFSTAT Field Descriptions (continued)

| Bits | Name | Description |
|-------|--------------|--|
| 22 | Excess Defer | Excessive transmission defer. This bit latches high and is cleared when read. This bit is cleared by default. 0 Normal operation. 1 The MAC excessively defers a transmission. |
| 23–31 | — | Reserved |

19.5.3.5.13 MAC Station Address Part 1 Register (MACSTNADDR1)

The MACSTNADDR1 register is written by the user. The value of the station address written into MACSTNADDR1 and MACSTNADDR2 is byte reversed from how it would appear in the DA field of a frame in memory. For example, for a station address of 0x12345678ABCD, MACSTNADDR1 is set to 0xCDAB7856 and MACSTNADDR2 is set to 0x34120000.

Figure 19-47 shows the MACSTNADDR1 register.




Figure 19-47. MAC Station Address Part 1 Register Definition

Table 19-51 describes the fields of the MACSTNADDR1 register.

Table 19-51. MACSTNADDR1 Field Descriptions

| Bit | Name | Description |
|-------|----------------------------|---|
| 0–7 | Station Address, 6th Octet | This field holds the sixth octet of the station address. The sixth octet (station address bits 40–47) defaults to a value of 0x0. |
| 8–15 | Station Address, 5th Octet | This field holds the fifth octet of the station address. The fifth octet (station address bits 32–39) defaults to a value of 0x0. |
| 16–23 | Station Address, 4th Octet | This field holds the fourth octet of the station address. The fourth octet (station address bits 24–31) defaults to a value of 0x0. |
| 24–31 | Station Address, 3rd Octet | This field holds the third octet of the station address. The third octet (station address bits 16–23) defaults to a value of 0x0. |

19.5.3.5.14 MAC Station Address Part 2 Register (MACSTNADDR2)

The MACSTNADDR2 register is written by the user. Figure 19-48 describes the definition for the MACSTNADDR2 register.




Figure 19-48. MAC Station Address Part 2 Register Definition

Table 19-52 describes the fields of the MACSTNADDR2 register.

Table 19-52. MACSTNADDR2 Field Descriptions

| Bit | Name | Description |
|-------|----------------------------|--|
| 0–7 | Station Address, 2nd Octet | This field holds the second octet of the station address. The second octet (station address bits 8–15) defaults to a value of 0x0. |
| 8–15 | Station Address, 1st Octet | This field holds the first octet of the station address. The first octet (station address bits 0–7) defaults to a value of 0x0. |
| 16–31 | — | Reserved |

19.5.3.5.15 MAC Exact Match Address 1–15 Part 1 Registers (MAC01ADDR1–MAC15ADDR1)

The MAC01ADDR1–MAC15ADDR1 registers are written by the user with the unicast or multicast addresses aliasing the MAC. Figure 19-49 describes the definition for all of the fifteen MAC_nADDR1 registers. The value of the address written into MAC_xADDR1 and MAC_nADDR2 is byte reversed from how it would appear in the DA field of a frame in memory. For example, for a MAC address of 0x12345678ABCD, MAC_nADDR1 is set to 0xCDAB7856 and MAC_nADDR2 is set to 0x34120000. For any valid, non-zero MAC address received, exact match registers can be excluded individually by clearing them to all zero bytes.




Figure 19-49. MAC Exact Match Address *n* Part 1 Register Definition

Table 19-53 describes the fields of a MAC_nADDR1 register.

Table 19-53. MAC_nADDR1 Field Descriptions

| Bit | Name | Description |
|-------|--------------------------------|--|
| 0–7 | Exact Match Address, 6th Octet | Holds the sixth octet of the exact match address. The sixth octet (destination address bits 40–47) defaults to a value of 0x0. |
| 8–15 | Exact Match Address, 5th Octet | Holds the fifth octet of the exact match address. The fifth octet (destination address bits 32–39) defaults to a value of 0x0. |
| 16–23 | Exact Match Address, 4th Octet | Holds the fourth octet of the exact match address. The fourth octet (destination address bits 24–31) defaults to a value of 0x0. |
| 24–31 | Exact Match Address, 3rd Octet | Holds the third octet of the exact match address. The third octet (destination address bits 16–23) defaults to a value of 0x0. |

19.5.3.5.16 MAC Exact Match Address 1–15 Part 2 Registers (MAC01ADDR2–MAC15ADDR2)

The MAC01ADDR2–MAC15ADDR2 registers are written by the user with the unicast or multicast addresses aliasing the MAC. Figure 19-50 describes the definition for all of the fifteen MAC_xADDR2 registers.

Offset eTSEC1:0x2_454C+8×*n*; eTSEC2:0x2_554C+8×*n*

Access: Read/Write




Figure 19-50. MAC Exact Match Address x Part 2 Register Definition

Table 19-54 describes the fields of a MAC_xADDR2 register.

Table 19-54. MAC01ADDR2–MAC15ADDR2 Field Descriptions

| Bit | Name | Description |
|-------|--------------------------------|--|
| 0–7 | Exact Match Address, 2nd Octet | This field holds the second octet of the exact match address. The second octet (destination address bits 8–15) defaults to a value of 0x0. |
| 8–15 | Exact Match Address, 1st Octet | This field holds the first octet of the exact match address. The first octet (destination address bits 0–7) defaults to a value of 0x0. |
| 16–31 | — | Reserved |

19.5.3.6 MIB Registers

This section describes the MIB registers. The eTSEC RMON module has 37 separate statistics counters, which simply count or accumulate statistical events that occur as packets transmitted and received. These counters support RMON MIB group 1, RMON MIB group 2 if table counters, RMON MIB group 3, RMON MIB group 9, RMON MIB 2, and the IEEE 802.3 Ethernet MIB.

An interrupt can be generated upon any one counter’s rollover condition through a carry interrupt output from the RMON. Each counter’s rollover condition can be discretely masked from causing an interrupt by internal masking registers. In addition, each individual counter value may be reset on read access, or all counters may be simultaneously reset by setting ECNTL[CLRCNT].

The majority of MIB counters are Ethernet-specific.

NOTE

RMON counters do not comprehend custom VLAN tagged frames. Affected counters include TRMGV, RMCA, RBCA, RXCF, RXPF, RXUO, RALN, RFLR, ROVR, RJBR, TMCA, TBCA, TXPF, TXCF. Specifically, custom VLAN tagged frames are not afforded the ability to be greater than 1518, as compared to the IEEE standard tagged frames.

NOTE

The transmit and receive frame counters (TR64, TR127, TR 255, TR511, TR1K, TRMAX, and TRMGV) do not increment for aborted frames (collision retry limit exceeded, late collision, underrun, EBERR, TxFIFO data error, frame truncated due to exceeding MAXFRM, or excessive deferral).

19.5.3.6.1 Transmit and Receive 64-Byte Frame Counter (TR64)

Figure 19-51 describes the definition for the TR64 register.




Figure 19-51. Transmit and Receive 64-Byte Frame Register Definition

Table 19-55 describes the fields of the TR64 register.

Table 19-55. TR64 Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–9 | — | Reserved |
| 10–31 | TR64 | Transmit and receive 64-byte frame counter—Increment for each good or bad frame transmitted and received which is 64 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes). |

19.5.3.6.2 Transmit and Receive 65- to 127-Byte Frame Counter (TR127)

Figure 19-52 describes the definition for the TR127 register.




Figure 19-52. Transmit and Receive 65- to 127-Byte Frame Register Definition

Table 19-56 describes the fields of the TR127 register.

Table 19-56. TR127 Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0–9 | — | Reserved |
| 10–31 | TR127 | Transmit and receive 65- to 127-byte frame counter—Increments for each good or bad frame transmitted and received which is 65–127 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes). |

19.5.3.6.3 Transmit and Receive 128- to 255-Byte Frame Counter (TR255)

Figure 19-53 describes the definition for the TR255 register.




Figure 19-53. Transmit and Received 128- to 255-Byte Frame Register Definition

Table 19-57 describes the fields of the TR255 register.

Table 19-57. TR255 Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0–9 | — | Reserved |
| 10–31 | TR255 | Transmit and receive 128- to 255-byte frame counter—Increments for each good or bad frame transmitted and received which is 128–255 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes). |

19.5.3.6.4 Transmit and Receive 256- to 511-Byte Frame Counter (TR511)

Figure 19-54 describes the definition for the TR511 register.




Figure 19-54. Transmit and Received 256- to 511-Byte Frame Register Definition

Table 19-58 describes the fields of the TR511 register.

Table 19-58. TR511 Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0–9 | — | Reserved |
| 10–31 | TR511 | Increments for each good or bad frame transmitted and received which is 256–511 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes). |

19.5.3.6.5 Transmit and Receive 512- to 1023-Byte Frame Counter (TR1K)

Figure 19-55 shows the TR1K register.



Figure 19-55. Transmit and Received 512- to 1023-Byte Frame Register Definition

Table 19-59 describes the fields of the TR1K register.

Table 19-59. TR1K Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–9 | — | Reserved |
| 10–31 | TR1K | Increments for each good or bad frame transmitted and received which is 512–1023 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes). |

19.5.3.6.6 Transmit and Receive 1024- to 1518-Byte Frame Counter (TRMAX)

Figure 19-56 describes the definition for the TRMAX register.




Figure 19-56. Transmit and Received 1024- to 1518-Byte Frame Register Definition

Table 19-60 describes the fields of the TRMAX register.

Table 19-60. TRMAX Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0–9 | — | Reserved |
| 10–31 | TRMAX | Increments for each good or bad frame transmitted and received which is 1024–1518 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes). |

19.5.3.6.7 Transmit and Receive 1519- to 1522-Byte VLAN Frame Counter (TRMGV)

Figure 19-57 describes the definition for the TRMGV register.



Figure 19-57. Transmit and Received 1519- to 1522-Byte VLAN Frame Register Definition

Table 19-61 describes the fields of the TRMGV register.

Table 19-61. TRMGV Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0–9 | — | Reserved |
| 10–31 | TRMGV | Increments for each good or bad frame transmitted and received which is 1519–1522 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes). |

19.5.3.6.8 Receive Byte Counter (RBYT)

Figure 19-58 shows the RBYT register.




Figure 19-58. Receive Byte Counter Register Definition

Table 19-62 describes the fields of the RBYT register.

Table 19-62. RBYT Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–31 | RBYT | Receive byte counter. The statistic counter register increments by the byte count of frames received, including those in bad packets, excluding preamble and SFD but including FCS bytes. |

19.5.3.6.9 Receive Packet Counter (RPKT)

Figure 19-59 describes the definition for the RPKT register.



Figure 19-59. Receive Packet Counter Register Definition

Table 19-63 describes the fields of the RPKT register.

Table 19-63. RPKT Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–9 | — | Reserved |
| 10-31 | RPKT | Receive packet counter. Increments for each frame received packet (including bad packets, all unicast, broadcast, and multicast packets). |

19.5.3.6.10 Receive FCS Error Counter (RFCS)

Figure 19-60 describes the definition for the RFCS register.



Figure 19-60. Receive FCS Error Counter Register Definition

Table 19-64 describes the fields of the RFCS register.

Table 19-64. RFCS Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–31 | RFCS | Receive FCS error counter. In Ethernet mode, increments for each frame received that has an integral 64–1518 length and contains a frame check sequence error. |

19.5.3.6.11 Receive Multicast Packet Counter (RMCA)

Figure 19-61 describes the definition for the RMCA register.




Figure 19-61. Receive Multicast Packet Counter Register Definition

Table 19-65 describes the fields of the RMCA register.

Table 19-65. RMCA Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–9 | — | Reserved |
| 10–31 | RMCA | Receive multicast packet counter. Increments for each multicast frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN), excluding broadcast frames. This count does not include range/length errors. |

19.5.3.6.12 Receive Broadcast Packet Counter (RBCA)

Figure 19-62 describes the definition for the RBCA register.




Figure 19-62. Receive Broadcast Packet Counter Register Definition

Table 19-66 describes the fields of the RBCA register.

Table 19-66. RBCA Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–9 | — | Reserved |
| 10–31 | RBCA | Receive broadcast packet counter. Increments for each broadcast frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN), excluding multicast frames. Does not include range/length errors. |

19.5.3.6.13 Receive Control Frame Packet Counter (RXCF)

Figure 19-63 describes the definition for the RXCF register.




Figure 19-63. Receive Control Frame Packet Counter Register Definition

Table 19-67 describes the fields of the RXCF register.

Table 19-67. RXCF Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–31 | RXCF | Receive control frame packet counter. Increments for each MAC control frame received (PAUSE and unsupported) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN). |

19.5.3.6.14 Receive Pause Frame Packet Counter (RXPF)

Figure 19-64 describes the definition for the RXPF register.




Figure 19-64. Receive Pause Frame Packet Counter Register Definition

Table 19-68 describes the fields of the RXPF register.

Table 19-68. RXPF Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–31 | RXPF | Receive PAUSE frame packet counter. Increments each time a PAUSE MAC control frame is received with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN). |

19.5.3.6.15 Receive Unknown Opcode Packet Counter (RXUO)

Figure 19-65 describes the definition for the RXUO register.



Figure 19-65. Receive Unknown OPCode Packet Counter Register Definition

Table 19-69 describes the fields of the RXUO register.

Table 19-69. RXUO Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–31 | RXUO | Receive unknown opcode counter. Increments each time a MAC control frame is received which contains an opcode other than PAUSE, but the frame has valid CRC and length 64 to 1518 (non VLAN) or 1522 (VLAN). |

19.5.3.6.16 Receive Alignment Error Counter (RALN)

Figure 19-66 describes the definition for the RALN register.



Figure 19-66. Receive Alignment Error Counter Register Definition

Table 19-70 describes the fields of the RALN register.

Table 19-70. RALN Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | — | Reserved |
| 16–31 | RALN | Receive alignment error counter. Increments for each received frame from 64 to 1518 (non VLAN) or 1522 (VLAN) which contains an invalid FCS and is not an integral number of bytes. |

19.5.3.6.17 Receive Frame Length Error Counter (RFLR)

Figure 19-67 describes the definition for the RFLR register.



Figure 19-67. Receive Frame Length Error Counter Register Definition

Table 19-71 describes the fields of the RFLR register.

Table 19-71. RFLR Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–31 | RFLR | Receive frame length error counter. Increments for each frame received in which the 802.3 length field did not match the number of data bytes actually received (46–1500 bytes). The counter does not increment if the length field is not a valid 802.3 length, such as an EtherType value. Frames tagged with a single VLAN tag are checked for valid length based on bytes 17-18 (rather than 13-14). Frames tagged (stacked) with multiple VLAN tags are not checked for valid length. |

19.5.3.6.18 Receive Code Error Counter (RCDE)

Figure 19-68 describes the definition for the RCDE register.



Figure 19-68. Receive Code Error Counter Register Definition

Table 19-72 describes the fields of the RCDE register.

Table 19-72. RCDE Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | — | Reserved |
| 16–31 | RCDE | Receive code error counter. Increments each time a valid carrier is present and at least one invalid data symbol is detected. |

19.5.3.6.19 Receive Carrier Sense Error Counter (RCSE)

Figure 19-69 describes the definition for the RCSE register.



Figure 19-69. Receive Carrier Sense Error Counter Register Definition

Table 19-73 describes the fields of the RCSE register.

Table 19-73. RCSE Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | — | Reserved |
| 16–31 | RCSE | Receive false carrier counter. Counts the number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface. The count represented by an instance of this object is incremented at most once per transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt. The event is reported along with the statistics generated on the next received frame, as defined by a 1 on TSECn_RX_ER and an 0xE on TSECn_RXD. Only one false carrier condition can be detected and logged between frames. |

19.5.3.6.20 Receive Undersize Packet Counter (RUND)

Figure 19-70 describes the definition for the RUND register.




Figure 19-70. Receive Undersize Packet Counter Register Definition

Table 19-74 describes the fields of the RUND register.

Table 19-74. RUND Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–31 | RUND | Receive undersize packet counter. Increments each time a frame is received which is less than 64 bytes in length and contains a valid FCS and were otherwise well formed. This count does not include range length errors. |

19.5.3.6.21 Receive Oversize Packet Counter (ROVR)

Figure 19-71 describes the definition for the ROVR register.



Figure 19-71. Receive Oversize Packet Counter Register Definition

Table 19-75 describes the fields of the ROVR register.

Table 19-75. ROVR Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | — | Reserved |
| 16–31 | ROVR | Receive oversize packet counter. Increments each time a frame is received which exceeded 1518 (non VLAN) or 1522 (VLAN) and contains a valid FCS and was otherwise well formed. |

19.5.3.6.22 Receive Fragments Counter (RFRG)

Figure 19-72 describes the definition for the RFRG register.




Figure 19-72. Receive Fragments Counter Register Definition

Table 19-76 describes the fields of the RFRG register.

Table 19-76. RFRG Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | — | Reserved |
| 16–31 | RFRG | Receive fragments counter. Increments for each frame received which is less than 64 bytes in length and contains an invalid FCS. This includes integral and non-integral lengths. |

19.5.3.6.23 Receive Jabber Counter (RJBR)

Figure 19-73 describes the definition for the RJBR register.



Figure 19-73. Receive Jabber Counter Register Definition

Table 19-77 describes the fields of the RJBR register.

Table 19-77. RJBR Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–31 | RJBR | Receive jabber counter. Increments for frames received which exceed 1518 (non VLAN) or 1522 (VLAN) bytes and contain an invalid FCS. This includes alignment errors. |

19.5.3.6.24 Receive Dropped Packet Counter (RDRP)

Figure 19-74 describes the definition for the RDRP register.




Figure 19-74. Receive Dropped Packet Counter Register Definition

Table 19-78 describes the fields of the RDRP register.

Table 19-78. RDRP Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | — | Reserved |
| 16–31 | RDRP | Receive dropped packets counter. Increments for frames received which are streamed to system but are later dropped due to lack of system resources. |

19.5.3.6.25 Transmit Byte Counter (TBYT)

Figure 19-75 depicts the TBYT register.




Figure 19-75. Transmit Byte Counter Register Definition

Table 19-79 describes the fields of the TBYT register.

Table 19-79. TBYT Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–31 | TBYT | Transmit byte counter. Increments by the number of bytes that were put on the wire including fragments of frames that were involved with collisions. This count does not include preamble/SFD or jam bytes, except for half-duplex flow control (back-pressure triggered by TCTRL[THDF]=1). For THDF, the sum total of ‘phantom’ preamble bytes transmitted for flow control purposes is included in the TBYT increment value of the next frame to be transmitted, up to 65,535 bytes of frame and phantom preamble. Note that the value of TBYT may be greater than the actual number of bytes transmitted if the frame is truncated because it exceeds MAXFRM. |

19.5.3.6.26 Transmit Packet Counter (TPKT)

Figure 19-76 describes the definition for the TPKT register.




Figure 19-76. Transmit Packet Counter Register Definition

Table 19-80 describes the fields of the TPKT register.

Table 19-80. TPKT Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–9 | — | Reserved |
| 10–31 | TPKT | Transmit packet counter. Increments for each transmitted packet (including bad packets, excessive deferred packets, excessive collision packets, late collision packets, all unicast, broadcast, and multicast packets). |

19.5.3.6.27 Transmit Multicast Packet Counter (TMCA)

Figure 19-77 describes the definition for the TMCA register.




Figure 19-77. Transmit Multicast Packet Counter Register Definition

Table 19-81 describes the fields of the TMCA register.

Table 19-81. TMCA Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–9 | — | Reserved |
| 10–31 | TMCA | Transmit multicast packet counter. Increments for each multicast valid frame transmitted (excluding broadcast frames) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN). |

19.5.3.6.28 Transmit Broadcast Packet Counter (TBCA)

Figure 19-78 describes the definition for the TBCA register.



Figure 19-78. Transmit Broadcast Packet Counter Register Definition

Table 19-82 describes the fields of the TBCA register.

Table 19-82. TBCA Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–9 | — | Reserved |
| 10–31 | TBCA | Transmit broadcast packet counter. Increments for each broadcast frame transmitted (excluding multicast frames) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN). |

19.5.3.6.29 Transmit Pause Control Frame Counter (TXPF)

Figure 19-79 describes the definition for the TXPF register.



Figure 19-79. Transmit Pause Control Frame Counter Register Definition

Table 19-83 describes the fields of the TXPF register.

Table 19-83. TXPF Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–15 | — | Reserved |
| 16–31 | TXPF | Transmit PAUSE frame packet counter. Increments each time a valid PAUSE MAC control frame is transmitted with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN). |

19.5.3.6.30 Transmit Deferral Packet Counter (TDFR)

Figure 19-80 describes the definition for the TDFR register.



Figure 19-80. Transmit Deferral Packet Counter Register Definition

Table 19-84 describes the fields of the TDFR register.

Table 19-84. TDFR Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–19 | — | Reserved |
| 20–31 | TDFR | Transmit deferral packet counter. Increments for each frame, which was deferred on its first transmission attempt. This count does not include frames involved in collisions. |

19.5.3.6.31 Transmit Excessive Deferral Packet Counter (TEDF)

Figure 19-81 describes the definition for the TEDF register.




Figure 19-81. Transmit Excessive Deferral Packet Counter Register Definition

Table 19-85 describes the fields of the TEDF register.

Table 19-85. TEDF Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–19 | — | Reserved |
| 20–31 | TEDF | Transmit excessive deferral packet counter. Increments for frames aborted which were deferred for an excessive period of time (3036 byte times). |

19.5.3.6.32 Transmit Single Collision Packet Counter (TSCL)

Figure 19-82 describes the definition for the TSCL register.



Figure 19-82. Transmit Single Collision Packet Counter Register Definition

Table 19-86 describes the fields of the TSCL register.

Table 19-86. TSCL Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–19 | — | Reserved |
| 20–31 | TSCL | Transmit single collision packet counter. Increments for each frame transmitted which experienced exactly one collision during transmission. |

19.5.3.6.33 Transmit Multiple Collision Packet Counter (TMCL)

Figure 19-83 describes the definition for the TMCL register.




Figure 19-83. Transmit Multiple Collision Packet Counter Register Definition

Table 19-87 describes the fields of the TMCL register.

Table 19-87. TMCL Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–19 | — | Reserved |
| 20–31 | TMCL | Transmit multiple collision packet counter. Increments for each frame transmitted which experienced 2–15 collisions (including any late collisions) during transmission as defined using the Half_Duplex[RETRANSMISSION MAXIMUM] field. |

19.5.3.6.34 Transmit Late Collision Packet Counter (TLCL)

Figure 19-84 describes the definition for the TLCL register.




Figure 19-84. Transmit Late Collision Packet Counter Register Definition

Table 19-88 describes the fields of the TLCL register.

Table 19-88. TLCL Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–19 | — | Reserved |
| 20–31 | TLCL | Transmit late collision packet counter. Increments for each frame transmitted which experienced a late collision during a transmission attempt. Late collisions are defined using the collision window field of the half-duplex [26:31] register. |

19.5.3.6.35 Transmit Excessive Collision Packet Counter (TXCL)

Figure 19-85 describes the definition for the TXCL register.




Figure 19-85. Transmit Excessive Collision Packet Counter Register Definition

Table 19-89 describes the fields of the TXCL register.

Table 19-89. TXCL Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–19 | — | Reserved |
| 20–31 | TXCL | Transmit excessive collision packet counter. Increments for each frame that experienced 16 collisions during transmission and was aborted. |

19.5.3.6.36 Transmit Total Collision Counter (TNCL)

Figure 19-86 describes the definition for the TNCL register.



Figure 19-86. Transmit Total Collision Counter Register Definition

Table 19-90 describes the fields of the TNCL register.

Table 19-90. TNCL Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–19 | — | Reserved |
| 20–31 | TNCL | Transmit total collision counter. Increments by the number of collisions experienced during the transmission of a frame as defined as the simultaneous presence of signals on the DO and RD circuits (That is, transmitting and receiving at the same time). Note: This count does not include collisions that result in an excessive collision condition. |

19.5.3.6.37 Transmit Drop Frame Counter (TDRP)

Figure 19-87 describes the definition for the TDRP register.



Figure 19-87. Transmit Drop Frame Counter Register Definition

Table 19-91 describes the fields of the TDRP register.

Table 19-91. TDRP Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–15 | — | Reserved |
| 16–31 | TDRP | Transmit drop frame counter. Increments each time a memory error or an underrun has occurred. |

19.5.3.6.38 Transmit Jabber Frame Counter (TJBR)

Figure 19-88 describes the definition for the TJBR register.




Figure 19-88. Transmit Jabber Frame Counter Register Definition

Table 19-92 describes the fields of the TJBR register.

Table 19-92. TJBR Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–19 | — | Reserved |
| 20–31 | TJBR | Transmit jabber frame counter. Increments for each oversized transmitted frame with an incorrect FCS value. |

19.5.3.6.39 Transmit FCS Error Counter (TFCS)

Figure 19-89 describes the definition for the TFCS register.




Figure 19-89. Transmit FCS Error Counter Register Definition

Table 19-93 describes the fields of the TFCS register.

Table 19-93. TFCS Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–19 | — | Reserved |
| 20–31 | TFCS | Transmit FCS error counter. Increments for every valid sized packet with an incorrect FCS value. |

19.5.3.6.40 Transmit Control Frame Counter (TXCF)

Figure 19-90 describes the definition for the TXCF register.




Figure 19-90. Transmit Control Frame Counter Register Definition

Table 19-94 describes the fields of the TXCF register.

Table 19-94. TXCF Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–19 | — | Reserved |
| 20–31 | TXCF | Transmit control frame counter. Increments for every control frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN). |

19.5.3.6.41 Transmit Oversize Frame Counter (TOVR)

Figure 19-91 describes the definition for the TOVR register.




Figure 19-91. Transmit Oversized Frame Counter Register Definition

Table 19-95 describes the fields of the TOVR register.

Table 19-95. TOVR Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–19 | — | Reserved |
| 20–31 | TOVR | Transmit oversize frame counter. Increments each time a frame is transmitted which exceeds 1518 bytes (non VLAN) or 1522 bytes (VLAN) with a correct FCS value. |

19.5.3.6.42 Transmit Undersize Frame Counter (TUND)

Figure 19-92 describes the definition for the TUND register.



Figure 19-92. Transmit Undersize Frame Counter Register Definition

Table 19-96 describes the fields of the TUND register.

Table 19-96. TUND Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–19 | — | Reserved |
| 20–31 | TUND | Transmit undersize frame counter. Increments for every frame less than 64 bytes, with a correct FCS value. |

19.5.3.6.43 Transmit Fragment Counter (TFRG)

Figure 19-93 describes the definition for the TFRG register.




Figure 19-93. Transmit Fragment Counter Register Definition

Table 19-97 describes the fields of the TFRG register.

Table 19-97. TFRG Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–19 | — | Reserved |
| 20–31 | TFRG | Transmit fragment counter. Increments for every frame less than 64 bytes, with an incorrect FCS value. |

19.5.3.6.44 Carry Register 1 (CAR1)

Carry register bits are cleared on carry register writes when the respective bits are set. Figure 19-94 describes the definition for the CAR1 register.




Figure 19-94. Carry Register 1 (CAR1) Register Definition

Table 19-98 describes the fields of the CAR1 register.

Table 19-98. CAR1 Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 0 | C164 | Carry register 1 TR64 counter carry bit |
| 1 | C1127 | Carry register 1 TR127 counter carry bit |
| 2 | C1255 | Carry register 1 TR255 counter carry bit |
| 3 | C1511 | Carry register 1 TR511 counter carry bit |

Table 19-98. CAR1 Field Descriptions (continued)


| Bits | Name | Description |
|------|-------|--|
| 4 | C11K | Carry register 1 TR1K counter carry bit |
| 5 | C1MAX | Carry register 1 TRMAX counter carry bit |
| 6 | C1MGV | Carry register 1 TRMGV counter carry bit |
| 7–13 | — | Reserved |
| 14 | C1REJ | Carry register 1 RREJ counter carry bit |
| 15 | C1RBY | Carry register 1 RBYT counter carry bit |
| 16 | C1RPK | Carry register 1 RPKT counter carry bit |
| 17 | C1RFC | Carry register 1 RFCS counter carry bit |
| 18 | C1RMC | Carry register 1 RMCA counter carry bit |
| 19 | C1RBC | Carry register 1 RBCA counter carry bit |
| 20 | C1RXC | Carry register 1 RXCF counter carry bit |
| 21 | C1RXP | Carry register 1 RXPF counter carry bit |
| 22 | C1RXU | Carry register 1 RXUO counter carry bit |
| 23 | C1RAL | Carry register 1 RALN counter carry bit |
| 24 | C1RFL | Carry register 1 RFLR counter carry bit |
| 25 | C1RCD | Carry register 1 RCDE counter carry bit |
| 26 | C1RCS | Carry register 1 RCSE counter carry bit |
| 27 | C1RUN | Carry register 1 RUND counter carry bit |
| 28 | C1ROV | Carry register 1 ROVR counter carry bit |
| 29 | C1RFR | Carry register 1 RFRG counter carry bit |
| 30 | C1RJB | Carry register 1 RJBR counter carry bit |
| 31 | C1RDR | Carry register 1 RDRP counter carry bit |

Table 19-99. CAR2 Field Descriptions (continued)

| Bits | Name | Description |
|------|-------|---|
| 29 | C2TNC | Carry register 2 TNCL counter carry bit |
| 30 | — | Reserved, should be cleared |
| 31 | C2TDP | Carry register 2 TDRP counter carry bit |

19.5.3.6.46 Carry Mask Register 1 (CAM1)

While one of the below mask bits are cleared, the corresponding carry bit in CAR1 is allowed to cause interrupt indications in register IEVENT[MSR0]. These bits all default to a set state. [Figure 19-96](#) describes the definition for the CAM1 register.


Figure 19-96. Carry Mask Register 1 (CAM1) Register Definition

[Table 19-100](#) describes the fields of the CAM1 register.

Table 19-100. CAM1 Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 0 | M164 | Mask register 1 TR64 counter carry bit mask |
| 1 | M1127 | Mask register 1 TR127 counter carry bit mask |
| 2 | M1255 | Mask register 1 TR255 counter carry bit mask |
| 3 | M1511 | Mask register 1 TR511 counter carry bit mask |
| 4 | M11k | Mask register 1 TR1K counter carry bit mask |
| 5 | M1MAX | Mask register 1 TRMAX counter carry bit mask |
| 6 | M1MGV | Mask register 1 TRMGV counter carry bit mask |
| 7–13 | — | Reserved |
| 14 | M1REJ | Mask register 1 RREJ counter carry bit mask |
| 15 | M1RBY | Mask register 1 RBYT counter carry bit mask |
| 16 | M1RPK | Mask register 1 RPKT counter carry bit mask |
| 17 | M1RFC | Mask register 1 RFCS counter carry bit mask |
| 18 | M1RMC | Mask register 1 RMCA counter carry bit mask |

Table 19-101. CAM2 Field Descriptions (continued)

| Bits | Name | Description |
|------|-------|---|
| 14 | M2TCF | Mask register 2 TXCF counter carry bit mask |
| 15 | M2TOV | Mask register 2 TOVR counter carry bit mask |
| 16 | M2TUN | Mask register 2 TUND counter carry bit mask |
| 17 | M2TFG | Mask register 2 TFRG counter carry bit mask |
| 18 | M2TBY | Mask register 2 TBYT counter carry bit mask |
| 19 | M2TPK | Mask register 2 TPKT counter carry bit mask |
| 20 | M2TMC | Mask register 2 TMCA counter carry bit mask |
| 21 | M2TBC | Mask register 2 TBCA counter carry bit mask |
| 22 | M2TPF | Mask register 2 TXPF counter carry bit mask |
| 23 | M2TDF | Mask register 2 TDFR counter carry bit mask |
| 24 | M2TED | Mask register 2 TEDF counter carry bit mask |
| 25 | M2TSC | Mask register 2 TSCL counter carry bit mask |
| 26 | M2TMA | Mask register 2 TMCL counter carry bit mask |
| 27 | M2TLC | Mask register 2 TLCL counter carry bit mask |
| 28 | M2TXC | Mask register 2 TXCL counter carry bit mask |
| 29 | M2TNC | Mask register 2 TNCL counter carry bit mask |
| 30 | — | Reserved |
| 31 | M2TDP | Mask register 2 TDRP counter carry bit mask |

19.5.3.6.48 Receive Filer Rejected Packet Counter (RREJ)

Figure 19-98 describes the definition for the RREJ register.




Figure 19-98. Receive Filer Rejected Packet Counter Register Definition

Table 19-102 describes the fields of the RREJ register.

Table 19-102. RREJ Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–9 | — | Reserved |
| 10–31 | RREJ | Receive filer rejected packet counter. Increments for each frame with valid CRC received, but rejected by the receive queue filer—either due to a matching rule that asserted the REJ flag or due to filing to a RxBD ring that was not enabled (see IEVENT[FIQ] error). |

19.5.3.7 Hash Function Registers

This section provides detailed descriptions of the registers used for hash functions. All of the registers are 32 bits wide. The DA field of every received frame is processed through a 32-bit CRC generator (CRC-32 polynomial), and the 8 or 9 most significant bits of the CRC are mapped to a hash table entry. The user can enable a hash entry by setting its bit. A hash entry usually represents a set of addresses. A hash table hit occurs if the DA CRC result points to an enabled hash entry. Software may need to further filter the address in order to eliminate false-positive hits in the hash table.

If $RCTRL[GHTX] = 0$, the 8 most significant bits of the CRC are used as the hash table index. In this case, registers IGADDR0–IGADDR7 comprise a 256-entry hash table exclusively for individual (unicast) address matching, while registers GADDR0–GADDR7 comprise a 256-entry hash table for group (multicast) address matching. If $RCTRL[GHTX] = 1$, the group hash table is extended to all 512 entries, and the 9 most significant bits of the CRC are used as the hash table index. In this case, registers IGADDR0–IGADDR7 hold hash table entries 0–255 for group addresses, while registers GADDR0–GADDR7 hold entries 256–511 of the extended group hash table.

See [Section 19.6.2.7.2, “Hash Table Algorithm,”](#) for more information on the hash algorithm.

19.5.3.7.1 Individual/Group Address Registers 0–7 (IGADDR n)

The IGADDR n registers are written by the user. Together these registers represent, depending on $RCTRL[GHTX]$, either the 256 entries of the individual address hash table, or the first 256 entries of the extended group address hash table used in the address recognition process. The user can enable a hash entry by setting the appropriate bit. A hash table hit occurs if the DA CRC-32 result points to an enabled hash entry.

[Figure 19-99](#) describes the definition for the IGADDR n register.




Figure 19-99. IGADDR n Register Definition

[Table 19-103](#) describes the fields of the IGADDR n register.

Table 19-103. IGADDR n Field Descriptions

| Bits | Name | Description |
|------|------------|---|
| 0–31 | IGADDR n | Represents the 32-bit value associated with the corresponding register. When $RCTRL[GHTX] = 0$, IGADDR0 contains entries 0–31 of the 256-entry individual hash table and IGADDR7 represents entries 224–255. When $RCTRL[GHTX] = 1$, IGADDR0 contains entries 0–31 of the 512-entry extended group hash table and IGADDR7 represents entries 224–255. |

19.5.3.7.2 Group Address Registers 0–7 (GADDR n)

The GADDR n registers are written by the user. Together these registers represent, depending on $RCTRL[GHTX]$, either the 256 entries of the group address hash table, or the last 256 entries of the

extended group address hash table used in the address recognition process. The user can enable a hash entry by setting the appropriate bit. A hash table hit occurs if the DA CRC result points to an enabled hash entry. [Figure 19-100](#) describes the definition for the $GADDR_n$ register.




Figure 19-100. $GADDR_n$ Register Definition

[Table 19-104](#) describes the fields of the $GADDR_n$ register.

Table 19-104. $GADDR_n$ Field Descriptions

| Bits | Name | Description |
|------|-----------|---|
| 0–31 | $GADDR_n$ | Represents the 32-bit value associated with the corresponding register. When $RCTRL[GHTX] = 0$, $GADDR_0$ contains entries 0–31 of the 256-entry group hash table and $GADDR_7$ represents entries 224–255. When $RCTRL[GHTX] = 1$, $GADDR_0$ contains entries 256–287 of the 512-entry extended group hash table and $GADDR_7$ represents entries 480–511. |

19.5.3.8 DMA Attribute Registers

This section describes the two eTSEC DMA attribute registers.

19.5.3.8.1 Attribute Register (ATTR)

The attribute register defines memory access attributes and transaction types used to access buffer descriptors, to write receive data, and to read transmit data. Snoop enable attributes may be set for reading buffer descriptors and for reading transmit data.

[Figure 19-101](#) describes the definition for the ATTR register.




Figure 19-101. ATTR Register Definition

Table 19-105 describes the fields of the ATTR register.

Table 19-105. ATTR Field Descriptions

| Bits | Name | Description |
|-------|---------|---|
| 0–23 | — | Reserved |
| 24 | RDSSEN | Rx data snoop enable 0 Disables snooping of all receive frames data to memory. 1 Enables snooping of all receive frames data to memory. |
| 25 | RBDSSEN | RxBD snoop enable. 0 Disables snooping of all receive BD memory accesses. 1 Enables snooping of all receive BD memory accesses. |
| 26–31 | — | Reserved |

19.5.3.9 Lossless Flow Control Configuration Registers

When enabled through RCTRL[LFC], the eTSEC tracks location of the last free BD in each Rx BD ring through the value of RFBPTR n . Using this pointer and the ring length stored in RQPRM n [LEN], the eTSEC continuously calculates the number of free BDs in the ring. Whenever the calculated number of free BDs in the ring drops below the pause threshold specified in RQPRM n [FBTHR], the eTSEC issues link layer flow control. It continues to assert flow control until the free BD count for each active ring reaches or exceeds RQPRM n [FBTHR]. See section 19.6.5.1/19-174 for the theory of operation of these registers.

19.5.3.9.1 Receive Queue Parameters 0–7 (RQPRM0–PQPRM7)

The RQPRM n registers specify the minimum number of BDs required to prevent flow control being asserted and the total number of Rx BDs in their respective ring. Whenever the free BD count calculated by the eTSEC for any active ring drops below the value of RQPRM n [FBTHR] for that ring, link level flow control is asserted. Software must not write to RQPRM n while LFC is enabled and the eTSEC is actively receiving frames. However, software may modify these registers after disabling LFC by clearing RCTRL[LFC]. Note that packets may be lost due to lack of RxBDs while RCTRL[LFC] is clear. Software can prevent packet loss by manually generating pause frames (through TCTRL[TFC_PAUSE]) to cover the time when RCTRL[LFC] is clear. Figure 19-102 describes the definition for the RQPRM n register.




Figure 19-102. RQPRM Register Definition

Table 19-106 describes the fields of the RQPRM register.

Table 19-106. RQPRM Field Descriptions

| Bits | Name | Description |
|------|-------|---|
| 0–7 | FBTHR | Free BD threshold. Minimum number of BDs required for normal operation. If the eTSEC calculated number of free BDs drops below this threshold, link layer flow control is asserted. |
| 8–31 | LEN | Ring length. Total number of Rx BDs in this ring. |

19.5.3.9.2 Receive Free Buffer Descriptor Pointer Registers 0–7 (RFBPTR0–RFBPTR7)

The RFBPTR n registers specify the location of the last free buffer descriptor in their respective ring. These registers live in the same 32b address space – and must share the same 4 most significant bits – as RBPTR n . That is, RFBPTR n and its associated RBPTR n must remain in the same 256MB page. Like RBPTR n , whenever RBASE n is updated, RFBPTR n is initialized to the value of RBASE n . This indicates that the ring is completely empty. As buffers are freed and their respective BDs are returned (by setting the EMPTY bit) to the ring, software is expected to update this register. The eTSEC then performs modulo arithmetic involving RBASE n , RBPTR n and RFBPTR n to determine the number of free BDs remaining in the ring. If, at any stage, the value written to RFBPTR n matches that of the respective RBPTR n the eTSEC free BD calculation assumes that the ring is now completely empty. For more information on the recommended use of these registers, see [Section 19.6.5.1, “Back Pressure Determination through Free Buffers.”](#)

Figure 19-103 describes the definition for the RFBPTR n register.




Figure 19-103. RFBPTR0–RFBPTR7 Register Definition

Table 19-107 describes the fields of the RFBPTR n registers.

Table 19-107. RFBPTR0–RFBPTR7 Field Descriptions

| Bits | Name | Description |
|-------|--------|--|
| 0–28 | RFBPTR | Pointer to the last free BD in RxBD Ring n . When RBASE n is updated, eTSEC initializes RFBPTR n to the value in the corresponding RBASE n . Software may update this register at any time to inform the eTSEC the location of the last free BD in the ring. Note that the 3 least-significant bits of this register are read only and zero. |
| 29–31 | — | Reserved. |

19.5.3.10 IEEE 1588-Compatible Timestamping Registers

IEEE 1588 compliant timestamping on this device is accomplished using the per-port transmit timestamping registers within each Ethernet controller memory space (See [Section 19.5.3.2.10, “Transmit Time Stamp Identification Register \(TMR_TXTS1–2_ID\),”](#) and [Section 19.5.3.2.11, “Transmit Time Stamp Register \(TMR_TXTS1–2_H/L\)”](#)) in conjunction with the following common registers, which are

located within the memory space for eTSEC1. Because the common 1588 timestamping registers exist within the eTSEC1 memory space, the eTSEC1 controller must remain enabled in order to use 1588 timestamping for any Ethernet port.

19.5.3.10.1 Timer Control Register (TMR_CTRL)

This register is used to reset, configure, and initialize the eTSEC precision timer clock. The control of all timer function is performed via programming eTSEC1. The register in eTSEC1 is shared for all eTSECs. Figure 19-7 describes the definition for the TMR_CTRL register.

Register fields not described below are reserved.




Figure 19-104. TMR_CTRL Register Definition

Table 19-108 describes the fields of the TMR_CTRL register. Register fields not described below are reserved.

Table 19-108. TMR_CTRL Register Field Descriptions

| Bits | Name | Description |
|-------|-------------|--|
| 0 | ALM1P | Alarm1 output polarity 0 active high output 1 active low output |
| 1 | ALM2P | Alarm2 output polarity 0 active high output 1 active low output |
| 3 | FS | FIPER start indication 0 Fiper is enabled through timer enable 1 Fiper is enabled through timer enable and alarm indication. |
| 6–15 | TCLK_PERIOD | 1588 timer reference clock period. The timer clock counter will increment by TCLK_PERIOD every time the accumulator register overflows. This clock period must be larger than the clock period of the timer reference clock. For applications where user does not want the clock period to be added, they can program this field to 1 to count the clock ticks. This field defaulted to 1 to count overflow ticks. For nanosecond granularity on 1588 timer counter rate, the TCLK_PERIOD should be calculated using the following equation: $TCLK_PERIOD = 10^9 / Nominal_Frequency$ |
| 18–21 | — | Reserved |

Table 19-108. TMR_CTRL Register Field Descriptions (continued)

| Bits | Name | Description |
|-------|-------|---|
| 22 | ETEP2 | External trigger 2 edge polarity 0 Time stamp on the rising edge of the external trigger 1 Time stamp on the falling edge of the external trigger |
| 23 | ETEP1 | External trigger 1 edge polarity 0 time stamp on the rising edge of the external trigger 1 time stamp on the falling edge of the external trigger |
| 24 | COPH | Generated clock (TSEC_TMR_GCLK) output phase. 0 non-inverted divided clock is output 1 inverted divided clock is output |
| 25 | CIPH | Oscillator input clock phase. 0 non-inverted timer input clock 1 inverted timer input clock (NOTE: this setting is reserved if CKSEL=01.) |
| 26 | TMSR | Timer soft reset. When enabled, it resets all the timer registers and state machines. 0 normal operation 1 place entire timer in reset except control and config registers NOTE: Prior to initiating timer reset (setting TMSR), must gracefully stop receiver (See MACCFG1[RX_EN] description). User programmable registers are not reset by the soft reset, for example, TMR_CTRL, TMR_TEMASK, TMR_PEMASK, TMR_ADD, TMR_PRSC, TMROFF_H/L, TMR_ALARMn, and TMR_FIPERn. |
| 27 | — | Reserved |
| 28 | BYP | Bypass drift compensated clock 0 64-bit clock counter is incremented on the accumulator overflow 1 64-bit clock counter is directly driven from the external oscillator ignoring accumulator overflow |
| 29 | TE | 1588 timer enable. If not enabled, all the timer registers and state machines are disabled. 0 timer not enabled 1 timer enabled and resume normal operation |
| 30–31 | CKSEL | 1588 Timer reference clock source select. 00 External high precision timer reference clock (TSEC_TMR_CLK) 01 eTSEC system clock 10 eTSEC1 transmit clock 11 RTC clock input Note that the 1588 reference clock must be no slower than 1/5 the Rx_clk frequency. |

19.5.3.10.2 Timer Event Register (TMR_TEVENT)

The eTSEC precision timer implementation can generate additional interrupts that are independent of the frame based events that controlled via IEVENT. The timer interrupts are not affected by any interrupt coalescing that may be specified in TXIC/RXIC. Software may poll this register at any time to check for pending interrupts. If an event occurs and its corresponding enable bit is set in the event mask register (TEMASK), the event also causes a hardware interrupt at the PIC. A bit in the timer event register is cleared by writing a 1 to that bit position. Figure 19-4 describes the definition for the TMR_TEVENT register.

Offset eTSEC1:0x2_4E04

Access: w1c




Figure 19-105. TMR_TEVENT Register Definition

Table 19-109 describes the fields of the TMR_TEVENT register fields for the timer.

Table 19-109. TMR_TEVENT Register Field Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–6 | — | Reserved |
| 6 | ETS2 | External trigger 2 timestamp sampled 0 external trigger timestamp not sampled 1 external trigger timestamp sampled |
| 7 | ETS1 | External trigger 1 timestamp sampled 0 external trigger timestamp not sampled 1 external trigger timestamp sampled |
| 8–13 | — | Reserved |
| 14 | ALM2 | Current time equaled alarm time register 2 0 alarm time has not be reached yet 1 alarm time has been reached |
| 15 | ALM1 | Current time equaled alarm time register 1 0 alarm time has not be reached yet 1 alarm time has been reached |
| 16–23 | — | Reserved |
| 24 | PP1 | Indicates that a periodic pulse has been generated based on FIPER1 register. 0 periodic pulse not generated 1 periodic pulse generated |
| 25 | PP2 | Indicates that a periodic pulse has been generated based on FIPER2 register. 0 periodic pulse not generated 1 periodic pulse generated |
| 26 | PP3 | Indicates that a periodic pulse has been generated based on FIPER3 register. 0 periodic pulse not generated 1 periodic pulse generated |
| 27–31 | — | Reserved |

19.5.3.10.3 Timer Event Mask Register (TMR_TEMASK)

Timer event mask register. The event mask register provides control over which possible interrupt events in the TMR_TEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this register are R/W and cleared upon a hardware reset. Figure 19-106 describes the definition for the TMR_TEMASK register.




Figure 19-106. TMR_TEMASK Register Definition

Table 19-110 describes the fields of the TMR_TEMASK register fields for the timer.

Table 19-110. TMR_TEMASK Register Field Descriptions

| Bits | Name | Description |
|-------|--------|--|
| 0–5 | — | Reserved |
| 6 | ETS2EN | External trigger 2 timestamp sample event enable |
| 7 | ETS1EN | External trigger 1 timestamp sample event enable |
| 8–13 | — | Reserved |
| 14 | ALM2EN | Timer ALM1 event enable |
| 15 | ALM1EN | Timer ALM2 event enable |
| 16–23 | — | Reserved |
| 24 | PP1EN | Periodic pulse event 1 enable |
| 25 | PP2EN | Periodic pulse event 2 enable |
| 26–31 | — | Reserved |

19.5.3.10.4 Timer PTP Packet Event Register (TMR_PEVENT)

The eTSEC precision timer logic can generate interrupts upon the capture of a timestamp due to either transmission or reception of a frame. If an event occurs and its corresponding enable bit is set in the event mask register (PEMASK), the event also causes a hardware interrupt at the PIC. A bit in the timer event register is cleared by writing a 1 to that bit position. Figure 19-107 describes the definition for the TMR_PEVENT register.




Figure 19-107. TMR_PEVENT Register Definition

Table 19-111 describes the fields of the TMR_PEVENT register fields for the timer.

Table 19-111. TMR_PEVENT Register Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–21 | — | Reserved |
| 22 | TXP2 | Indicates that a PTP frame has been transmitted and its timestamp is stored in TXTS2 register. 0 PTP packet not transmitted 1 PTP packet has been transmitted |
| 23 | TXP1 | Indicates that a PTP frame has been transmitted and its timestamp is stored in TXTS1 register. 0 PTP packet not transmitted 1 PTP packet has been transmitted |
| 24–30 | — | Reserved |
| 31 | RXP | Indicates that a PTP frame has been received 0 PTP packet not received 1 PTP packet has been received |

19.5.3.10.5 Timer Event Mask Register (TMR_PEMASK)

Timer event mask register. The event mask register provides control over which possible interrupt events in the TMR_PEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this register are R/W and cleared upon a hardware reset. Figure 19-108 describes the definition for the TMR_PEMASK register.




Figure 19-108. TMR_PEMASK Register Definition

Table 19-112 describes the fields of the TMR_PEMASK register fields for the timer.

Table 19-112. TMR_PEMASK Register Field Descriptions

| Bits | Name | Description |
|-------|--------|------------------------------------|
| 0–21 | — | Reserved |
| 22 | TXP2EN | Transmit PTP packet event 2 enable |
| 23 | TXP1EN | Transmit PTP packet event 1 enable |
| 24–30 | — | Reserved |
| 31 | RXPEN | Receive PTP packet event enable |

19.5.3.10.6 Timer Status Register (TMR_STAT)

This register requires the eTSEC filer to be enabled (via RCTRL[FILREN]). When eTSEC generates an interrupt based on the timestamp event for a received packet, the queue ID which the incoming packet will be sent to is captured in this register. This register update is synchronized with the RXF interrupt of the corresponding received packet. Writing 1 to any bit of this register clears it. Figure 19-109 describes the definition for the TMR_STAT register.




Figure 19-109. TMR_STAT Register Definition

Table 19-113 describes the fields of the TMR_STAT register.

Table 19-113. TMR_STAT Register Field Descriptions

| Bits | Name | Description |
|-------|----------|--|
| 0–25 | — | Reserved |
| 26–31 | STAT_VEC | Timer general purpose status vector. It will store the 6-bit queue number generated by the filer. User to decode this status vector. For example, user can encode received PTP packet message types (Sync, Delay_req, Follow_up, Delay_resp, Management) in the filer virtual queue field. |

19.5.3.10.7 Timer Counter Register (TMR_CNT_H/L)

The timer register (TMR_CNT_H/L) represents accurate time in terms clock ticks or in nano-seconds. Writes to these registers will override the previous time. The register in eTSEC1 is shared for all eTSECs. This is a read/write register. Figure 19-110 describes the definition for the TMR_CNT_H/L register.




Figure 19-110. TMR_CNT_H Register Definition

Table 19-114 describes the fields of the TMR_CNT_H/L register.

Table 19-114. TMR_CNT_H/L Register Field Descriptions

| Bits | Name | Description |
|------|-------------|--|
| 0–63 | TMR_CNT_H/L | <p>Value of the current time counter. Current time is calculated by adding TMROFF_H/L with the TMR_CNT_H/L counter. This register can be written through the register writes. Writes to the TMR_CNT_L register copies the written value into the shadow TMR_CNT_L register. Writes to the TMR_CNT_H register copies the values written into the shadow TMR_CNT_H register. Contents of the shadow registers are copied into the TMR_CNT_L and TMR_CNT_H registers following a write into the TMR_CNT_H register. Writes to these registers have precedence over the timer increment. The user must write to TMR_CNT_L register first.</p> <p>Reads from the TMR_CNT_L register copies the entire 64-bit clock time of the read enable into the TMR_CNT_H/L shadow registers. Read instruction from the TMR_CNT_H register reads the value stored in the TMR_CNT_H shadow register. The user must read the TMR_CNT_L register first to get correct 64-bit TMR_CNT_H/L counter values.</p> |

19.5.3.10.8 Timer Drift Compensation Addend Register (TMR_ADD)

Timer drift compensation addend register (TMR_ADD) is used to hold timer frequency compensation value (FreqCompensationValue). The nominal frequency of the clock counter is determined by the FreqDivRatio and the clock frequency (FreqClock). This register is programmed with $2^{32}/\text{FreqDivRatio}$. Frequency division ratio (FreqDivRatio) is the ratio between the frequency of the oscillator (TimerOsc) and the desired clock frequency (NominalFreq). FreqDivRatio is a design constant chosen to be greater than 1.0001. The ADDEND value is added to the 32-bit accumulator register at every rising edge of the oscillator clock (TimerOsc). The clock counter is incremented at every carry pulse of the accumulator. Only one of this register is required for the entire group of eTSECs. Figure 19-111 describes the definition of the TMR_ADD register.



Figure 19-111. TMR_ADD Register Definition

Table 19-115 describes the fields of the TMR_ADD register fields for the timer.

Table 19-115. TMR_ADD Register Field Descriptions

| Bits | Name | Description |
|------|--------|---|
| 0–31 | ADDEND | Timer drift compensation addend register value. It is programmed with a value of $2^{32}/\text{FreqDivRatio}$. For example, TimerOsc = 50 MHz NominalFreq = 40 MHz FreqDivRatio = 1.25 $\text{ADDEND} = \text{ceil}(2^{32}/1.25) = 0xCCCC_CCCD$ |

19.5.3.10.9 Timer Accumulator Register (TMR_ACC)

Timer accumulator register accumulates the value of the addend register into it. An overflow pulse of the accumulator is used to increment the timer clock by TMR_CTRL[TCLK_PERIOD]. This register is read only in normal operation. The register in eTSEC1 is shared for all eTSECs. Figure 19-112 describes the definition of the TMR_ACC register.




Figure 19-112. TMR_ACC Register Definition

Table 19-116 describes the fields of the TMR_ACC register.

Table 19-116. TMR_ACC Register Field Descriptions

| Bits | Name | Description |
|------|---------|-----------------------------------|
| 0–31 | TMR_ACC | 32-bit timer accumulator register |

19.5.3.10.10 Timer Prescale Register (TMR_PRSC)

Timer generated output clock prescale register. It is used to adjust output clock frequency that is put onto the 1588 clock output signal. The register in eTSEC1 is shared for all eTSECs. Figure 19-113 describes the definition for the TMR_PRSC register.




Figure 19-113. TMR_PRSC Register Definition

Table 19-119 describes the fields of the TMR_ALARM n _H/L register.

Table 19-119. TMR_ALARM n _H/L Register Field Descriptions

| Bits | Name | Description |
|------|-----------|--|
| 0–63 | ALARM_H/L | Alarm time comparator register. The corresponding alarm event in TMR_TEVENT is set when the current time counter becomes equal to or greater than the alarm time compare value in TMR_ALARM n _L/H. Writing the TMR_ALARM n _L register deactivates the alarm event after it has fired. Writing the TMR_ALARM n _L followed by the TMR_ALARM n _H register rearms the alarm function with the new compare value. The value programmed in this register must be an integer multiple of TMR_CTRL[TCLK_PERIOD] in order to get correct result. This register is reset to all ones to avoid false alarm after reset. In FS mode the alarm trigger is used as an indication to the fiber start down counting. Only alarm 1 supports this mode. In FS mode, alarm polarity bit should be configured to 0 (rising edge). |

19.5.3.10.13 Timer Fixed Interval Period Register (TMR_FIPER1–3)

Timer fixed interval period pulse generator register. It is used to generate periodic pulses. This register is reset with 0xFFFF_FFFF to prevent any false pulse upon initialization. The down count register loads the value programmed in the fixed period interval (FIPER). FIPER register must be programmed before the timer is enabled. At every tick of the timer accumulator overflow, the counter decrements by the value of TMR_CTRL[TCLK_PERIOD]. It generates a pulse when the down counter value reaches zero. It reloads the down counter in the cycle following a pulse.

Should a user wish to use the TMR_FIPER1 register to generate a 1 PPS event, the following setup should be used:

- Program TMR_FIPER1 to a value that will generate a pulse every second,
- Program TMR_ALARM1 to the correct time for the first PPS event
- Enable the timer

The eTSEC will then wait for TMR_ALARM1 to expire before enabling the count down of TMR_FIPER1. The end result will be that TMR_FIPER1 will pulse every second after the original timer ALARM1 expired.

Note:

In the case where the PPS signals are required to be phased aligned to the prescale output clock, the alarm value should be configured to **1 clock period less** than the wanted value.

In order to keep tracking the prescale output clock, each time before enabling the FIPER, the user must reset the FIPER by writing a new value to the register. The ratio between the prescale register value and the FIPER value should be devisable by the clk period.

$$\text{FIPER_VALUE} = (\text{prescale_value} \times \text{tclk_per} \times N) - \text{tclk_per}$$

For example:

$$\text{prescale} = 9$$

$$\text{clock period} = 10$$

The FIPER can get the following values: 80, 170, 260

The three registers in eTSEC1 are shared for all eTSECs. [Figure 19-116](#) describes the definition for the TMR_FIPER register.




Figure 19-116. TMR_FIPER_n Register Definition

[Table 19-120](#) describes the fields of the TMR_FIPER register.

Table 19-120. TMR_FIPER Register Field Descriptions

| Bits | Name | Description |
|------|-------|---|
| 0–31 | FIPER | Fixed interval pulse period register. This field must be programmed to an integer multiple of TMR_CTRL[TCLK_PERIOD] value to ensure a period pulse being generated correctly. |

19.5.3.10.14 External Trigger Stamp Register (TMR_ETTS1–2_H/L)

General purpose external trigger -stamp register (TMR_ETTS_n_H/L). This register holds time at the programmable edge of the external trigger. The registers in eTSEC1 are shared for all eTSECs. This register is read only in normal operation. [Figure 19-117](#) describes the definition for the TMR_ETTS_n_H/L register.




Figure 19-117. TMR_ETTS1-2_H/L Register Definition

[Table 19-121](#) describes the fields of the TMR_ETTS_n_H/L register.

Table 19-121. TMR_ETTS1-2_H Register Field Descriptions

| Bits | Name | Description |
|------|----------|--|
| 0–63 | ETTS_H/L | Time stamp field at the programmable edge of the external trigger. |

19.5.4 Ten-Bit Interface (TBI)

This section describes the ten-bit interface (TBI), reduced ten-bit interface (RTBI), and the TBI/RTBI MII set of registers. TBI and RTBI operate in the same manner (the only difference is that RTBI has reduced I/O signalling).

19.5.4.1 TBI Transmit Process

The eTSEC's TBI implements the transmit portion of the physical coding sublayer as found in Clause 36 of IEEE 802.3z. In SerDes mode, packets conveyed across the GMII are encapsulated and encoded into 10-bit symbols and output to the SerDes. In GMII mode, the GMII signals are passed through to the attached GMII PHY.

19.5.4.1.1 Packet Encapsulation

If TX_EN is de-asserted the eTSEC outputs an idle stream. If TX_EN is asserted, a Start_of_Packet symbol is output. This symbol replaces the first byte of the preamble field. All other bytes of the packet pass through an 8B10B encoding module. After the last byte of the FCS field is signaled through the GMII, the MAC de-asserts TX_EN. The eTSEC then outputs an End_of_Packet symbol. Then, depending on the position of the End_of_Packet symbols (being in either an odd or even position) the eTSEC outputs one or two Carrier_Extend symbols. Following the last Carrier_Extend symbol, the eTSEC resumes sending idle codes. If, during a packet, the eTSEC wishes to mark a byte invalid, TX_ER is asserted. The eTSEC, upon detection of TX_ER, substitutes the data symbol for an Error_Propagation symbol.

19.5.4.1.2 8B10B Encoding

Every eight-bit data octet has two (not necessarily different) ten-bit symbols associated with it. Depending on the running disparity (the cumulative difference of ones and zeros) the eTSEC module chooses the appropriate symbol.

Special encapsulation symbols are called ordered_sets. Ordered_sets are comprised of one to four ten-bit symbols. Ordered_sets can be found in Clause 36 of the IEEE 802.3z specification.

19.5.4.1.3 Preamble Shortening

Because the idle ordered_set comprises two symbols and begins on an even symbols boundary, packets can only begin on an even boundary. However, the GMII has no such restriction and may signal TX_EN on an odd boundary. If this happens, the eTSEC delays the Start_of_Packet symbol, effectively ignoring the first byte of preamble; thus, a seven octet preamble becomes six octets on the Ten-Bit Interface.

19.5.4.2 TBI Receive Process

The eTSEC's TBI Implements the receive portion of the physical coding sublayer as found in Clause 36 of IEEE 802.3z. The Receive portion includes the Synchronization state machine. In SerDes mode, the eTSEC first attempts to acquire synchronization on the link by examining received symbols. Once synchronization is acquired, received packets are decoded and sent across the Receive GMII interface. In GMII mode, the GMII signals are passed through to the MAC.

19.5.4.2.1 Synchronization

The eTSEC examines received symbols looking for the seven bit 'comma' string embedded in some special symbols. Both the idle ordered_set and the Configuration ordered_set contain a symbol which has the comma. Once a certain number of codes with comma are detected, the eTSEC is considered to have acquired synchronization.

19.5.4.2.2 Auto-Negotiation for 1000BASE-X

Once synchronization is acquired, ordered_sets are decoded. If Configuration ordered_sets are received, the eTSEC decodes the two octet data field and the sixteen-bit Configuration data is stored and used to Auto-Negotiate with the link partner. in the Receive Configuration Register (RXCR[15:0]) an internal register used to receive all the link partners informations and used to compare to local ability during negotiation. Not visible to user. If, during Auto-Negotiation an invalid symbol is detected, Auto-Negotiation re-starts. After Auto-Negotiation is completed the TBI MII Status Register SR[AN done] in set. In this mode, packets may be received from the link partner.

19.5.4.3 TBI MII Set Register Descriptions

This section describes the TBI MII registers. All of the TBI registers are 16 bits wide. The TBI registers are accessed at the offset of the TBI physical address. The eTSEC's TBI physical address is stored in the TBIPA register. Writing to the TBI registers is performed in a way similar to writing to an external PHY, using the MII management interface. Using TBIPA in place of the PHY address, in the MIIMADD[PHY Address] field, and setting the MIIMADD[Register Address] to the appropriate address offset that corresponds to the register that one wants to read or write (see Table 19-122), the user can read (set MIIMCOM[read cycle]) or write (writing to MIIMCON[PHY control]) to the TBI block. Refer to the TBI physical address register in Section 19.5.3.1, “eTSEC General Control and Status Registers,” and the TBI MII register set in Table 19-122. Notice that jitter diagnostics and TBI control are not IEEE 802.3 required registers and are only used for test and control of the eTSEC TBI block. The TBI's TBI control register (TBI) is for configuring the eTSEC ten-bit interface block. However, because this TBI block has an MII management interface (just like any other PHY), it has an IEEE 802.3 register called the control register (CR).

Table 19-122. TBI MII Register Set

| Offset Address | Name | Access | Size | Section/page |
|--|---|------------------|---------|------------------------------------|
| TEN-BIT INTERFACE (TBI) REGISTERS | | | | 19.5.4/19-120 |
| 0x00 | Control (CR) | R/W ¹ | 16 bits | 19.5.4.3.1/19-123 |
| 0x01 | Status (SR) | R, LH, LL | 16 bits | 19.5.4.3.2/19-124 |
| 0x02–0x03 | Reserved | R | 2 bytes | — |
| 0x04 | AN advertisement (ANA) | RW, R | 16 bits | 19.5.4.3.2/19-124 |
| 0x05 | AN link partner base page ability (ANLPBPA) | R | 16 bits | 19.5.4.3.4/19-127 |
| 0x06 | AN expansion (ANEX) | R, LH | 16 bits | 19.5.4.3.5/19-128 |
| 0x07 | AN next page transmit (ANNPT) | R/W, R | 16 bits | 19.5.4.3.6/19-128 |
| 0x08 | AN link partner ability next page (ANLPANP) | R | 16 bits | 19.5.4.3.7/19-129 |
| 0x0F | Extended status (EXST) | R | 16 bits | 19.5.4.3.8/19-130 |
| 0x10 | Jitter diagnostics (JD) | R/W | 16 bits | 19.5.4.3.9/19-131 |
| 0x11 | TBI control (TBICON) | R/W | 16 bits | 19.5.4.3.10/19-132 |

¹ R = Read-only, WO = Write Only, R/W = Read and Write, LH = Latches High, LL = Latches Low, SC = Self-clearing,

19.5.4.3.1 Control Register (CR)

Figure 19-118 describes the definition for the CR register.




Figure 19-118. Control Register Definition

Table 19-123 describes the fields of the CR register.

Table 19-123. CR Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | | | | |
|-------------------------|-------------|--|-------------------------|-------|-------|----------|---|---|----------|---|---|-----------|---|---|----------|---|---|
| 0 | PHY Reset | PHY reset. This bit is cleared by default. This bit is self-clearing. 0 Normal operation. 1 The internal state of the TBI is reset. This in turn may change the state of the TBI link partner. | | | | | | | | | | | | | | | |
| 1 | — | Reserved | | | | | | | | | | | | | | | |
| 2 | Speed[0] | Speed selection. This bit defaults to a cleared state and should always be cleared, which corresponds to 1000 Mbps speed. Setting this field controls the speed at which the TBI operates. The table for Speed[1] provides the appropriate encoding. Its default is bit[2] = '0'; bit[9] = '1'. | | | | | | | | | | | | | | | |
| 3 | AN Enable | Auto-negotiation enable. This bit is set by default. 0 The values programmed in bits 2, 7 and 9 determine the operating condition of the link. 1 Auto-negotiation process enabled. | | | | | | | | | | | | | | | |
| 4–5 | — | Reserved | | | | | | | | | | | | | | | |
| 6 | Reset AN | Reset auto-negotiation. This bit is cleared by default and is self-clearing. 0 Normal operation. 1 The auto-negotiation process restarts. This action is only available if auto-negotiation is enabled. | | | | | | | | | | | | | | | |
| 7 | Full Duplex | Duplex mode. This bit is set by default. 0 Reserved. 1 Full-duplex operation. | | | | | | | | | | | | | | | |
| 8 | — | Reserved, should be cleared. | | | | | | | | | | | | | | | |
| 9 | Speed[1] | Speed selection. This bit defaults to a set state and should always be set, which corresponds to 1000 Mbps speed. Setting this field controls the speed at which the TBI operates. The following table provides the appropriate encoding. Its default is bit[2] = '0'; bit[9] = '1'. | | | | | | | | | | | | | | | |
| | | <table border="1" style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="width: 40%;">Maximum Operating Speed</th> <th style="width: 10%;">Bit 2</th> <th style="width: 10%;">Bit 9</th> </tr> </thead> <tbody> <tr> <td>Reserved</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> </tr> <tr> <td>Reserved</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> </tr> <tr> <td>1000 Mbps</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> </tr> <tr> <td>Reserved</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> </tr> </tbody> </table> | Maximum Operating Speed | Bit 2 | Bit 9 | Reserved | 0 | 0 | Reserved | 1 | 0 | 1000 Mbps | 0 | 1 | Reserved | 1 | 1 |
| Maximum Operating Speed | Bit 2 | Bit 9 | | | | | | | | | | | | | | | |
| Reserved | 0 | 0 | | | | | | | | | | | | | | | |
| Reserved | 1 | 0 | | | | | | | | | | | | | | | |
| 1000 Mbps | 0 | 1 | | | | | | | | | | | | | | | |
| Reserved | 1 | 1 | | | | | | | | | | | | | | | |
| 10–15 | — | Reserved | | | | | | | | | | | | | | | |

19.5.4.3.3 AN Advertisement Register (ANA)

Figure 19-120 describes the definition for the ANA register.




Figure 19-120. AN Advertisement Register Definition

Table 19-125 describes the fields of the ANA register.

Table 19-125. ANA Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | | | | |
|--------------|----------------|---|--------------|----------------|-------------|---|---|-------------------|---|---|--------------------------------------|---|---|-----------------|---|---|---|
| 0 | Next Page | <p>Next page configuration. The local device sets this bit to either request next page transmission or advertise next page exchange capability.</p> <p>0 The local device wishes not to engage in next page exchange. 1 The local device has no next pages but wishes to allow reception of next pages. If the local device has no next pages and the link partner wishes to send next pages, the local device shall send null message codes and have the message page set to 0b000_0000_0001, as defined in annex 28C.</p> | | | | | | | | | | | | | | | |
| 1 | — | Reserved. (Ignore on read) | | | | | | | | | | | | | | | |
| 2–3 | Remote Fault | <p>The local device's remote fault condition is encoded in bits 2 and 3 of the base page. Values are shown in the following table. The default value is 00. Indicate a fault by setting a non-zero remote fault encoding and re-negotiating.</p> <table border="1" style="margin-left: auto; margin-right: auto; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 15%;">RF1 bit[3]</th> <th style="width: 15%;">RF2 bit[2]</th> <th style="width: 70%;">Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No error, link OK</td> </tr> <tr> <td>0</td> <td>1</td> <td>Offline</td> </tr> <tr> <td>1</td> <td>0</td> <td>Link_Failure</td> </tr> <tr> <td>1</td> <td>1</td> <td>Auto-Negotiation_Error</td> </tr> </tbody> </table> | RF1 bit[3] | RF2 bit[2] | Description | 0 | 0 | No error, link OK | 0 | 1 | Offline | 1 | 0 | Link_Failure | 1 | 1 | Auto-Negotiation_Error |
| RF1 bit[3] | RF2 bit[2] | Description | | | | | | | | | | | | | | | |
| 0 | 0 | No error, link OK | | | | | | | | | | | | | | | |
| 0 | 1 | Offline | | | | | | | | | | | | | | | |
| 1 | 0 | Link_Failure | | | | | | | | | | | | | | | |
| 1 | 1 | Auto-Negotiation_Error | | | | | | | | | | | | | | | |
| 4–6 | — | Reserved, should be cleared. | | | | | | | | | | | | | | | |
| 7–8 | Pause | <p>The local device's PAUSE capability is encoded in bits 7 and 8, and the decodes are shown in the following table. For priority resolution information consult Table 19-126.</p> <table border="1" style="margin-left: auto; margin-right: auto; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 15%;">PAUSE bit[8]</th> <th style="width: 15%;">ASM_DIR bit[7]</th> <th style="width: 70%;">Capability</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No PAUSE</td> </tr> <tr> <td>0</td> <td>1</td> <td>Asymmetric PAUSE toward link partner</td> </tr> <tr> <td>1</td> <td>0</td> <td>Symmetric PAUSE</td> </tr> <tr> <td>1</td> <td>1</td> <td>Both symmetric PAUSE and Asymmetric PAUSE toward local device</td> </tr> </tbody> </table> | PAUSE bit[8] | ASM_DIR bit[7] | Capability | 0 | 0 | No PAUSE | 0 | 1 | Asymmetric PAUSE toward link partner | 1 | 0 | Symmetric PAUSE | 1 | 1 | Both symmetric PAUSE and Asymmetric PAUSE toward local device |
| PAUSE bit[8] | ASM_DIR bit[7] | Capability | | | | | | | | | | | | | | | |
| 0 | 0 | No PAUSE | | | | | | | | | | | | | | | |
| 0 | 1 | Asymmetric PAUSE toward link partner | | | | | | | | | | | | | | | |
| 1 | 0 | Symmetric PAUSE | | | | | | | | | | | | | | | |
| 1 | 1 | Both symmetric PAUSE and Asymmetric PAUSE toward local device | | | | | | | | | | | | | | | |

Table 19-125. ANA Field Descriptions (continued)

| Bits | Name | Description |
|-------|-------------|---|
| 9 | Half Duplex | Half-duplex capability. 0 Designates local device as not capable of half-duplex operation. 1 Designates local device as capable of half-duplex operation. |
| 10 | Full Duplex | Full-duplex capability. 0 Designates the local device as not capable of full-duplex operation. 1 Designates the local device as capable of full-duplex operation. |
| 11–15 | — | Reserved, should be cleared. |

Table 19-126 describes the resolution of pause priority.

Table 19-126. PAUSE Priority Resolution

| Local Device | | Link Partner | | Local Resolution | Link Partner Resolution |
|--------------|---------|--------------|---------|---|---|
| PAUSE | ASM_DIR | PAUSE | ASM_DIR | | |
| 0 | 0 | x | x | Disable PAUSE transmit Disable PAUSE receive | Disable PAUSE transmit Disable PAUSE receive |
| 0 | 1 | 0 | x | Disable PAUSE transmit Disable PAUSE receive | Disable PAUSE transmit Disable PAUSE receive |
| 0 | 1 | 1 | 0 | Disable PAUSE transmit Disable PAUSE receive | Disable PAUSE transmit Disable PAUSE receive |
| 0 | 1 | 1 | 1 | Enable PAUSE transmit Disable PAUSE receive | Disable PAUSE transmit Enable PAUSE receive |
| 1 | 0 | 0 | x | Disable PAUSE transmit Disable PAUSE receive | Disable PAUSE transmit Disable PAUSE receive |
| 1 | 0 | 1 | x | Enable PAUSE transmit Enable PAUSE receive | Enable PAUSE transmit Enable PAUSE receive |
| 1 | 1 | 0 | 0 | Disable PAUSE transmit Disable PAUSE receive | Disable PAUSE transmit Disable PAUSE receive |
| 1 | 1 | 0 | 1 | Disable PAUSE transmit Enable PAUSE receive | Enable PAUSE transmit Disable PAUSE receive |
| 1 | 1 | 1 | x | Enable PAUSE transmit Enable PAUSE receive | Enable PAUSE transmit Enable PAUSE receive |

19.5.4.3.4 AN Link Partner Base Page Ability Register (ANLPBPA)

Figure 19-121 describes the definition for the ANLPBPA register.




Figure 19-121. AN Link Partner Base Page Ability Register Definition

Table 19-127 describes the fields of the ANLPBPA register.

Table 19-127. ANLPBPA Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | | | | |
|--------------|----------------|---|--------------|----------------|-------------|---|---|-------------------|---|---|--------------------------------------|---|---|-----------------|---|---|---|
| 0 | Next Page | Next page. This bit is read-only. The link partner sets or clears this bit. 0 Link partner has no subsequent next pages or is not capable of receiving next pages. 1 Link partner either requesting next page transmission or indicating the capability to receive next pages. | | | | | | | | | | | | | | | |
| 1 | — | Reserved. (Ignore on read) | | | | | | | | | | | | | | | |
| 2–3 | Remote Fault | The link partner’s remote fault condition is encoded in bits 2 and 3 of the base page. Values are shown in the remote fault encoding field table below. This bit is read-only. <table border="1" style="margin: 10px auto; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 15%;">RF1 bit[3]</th> <th style="width: 15%;">RF2 bit[2]</th> <th style="width: 70%;">Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No error, link OK</td> </tr> <tr> <td>0</td> <td>1</td> <td>Offline</td> </tr> <tr> <td>1</td> <td>0</td> <td>Link_Failure</td> </tr> <tr> <td>1</td> <td>1</td> <td>Auto-Negotiation_Error</td> </tr> </tbody> </table> | RF1 bit[3] | RF2 bit[2] | Description | 0 | 0 | No error, link OK | 0 | 1 | Offline | 1 | 0 | Link_Failure | 1 | 1 | Auto-Negotiation_Error |
| RF1 bit[3] | RF2 bit[2] | Description | | | | | | | | | | | | | | | |
| 0 | 0 | No error, link OK | | | | | | | | | | | | | | | |
| 0 | 1 | Offline | | | | | | | | | | | | | | | |
| 1 | 0 | Link_Failure | | | | | | | | | | | | | | | |
| 1 | 1 | Auto-Negotiation_Error | | | | | | | | | | | | | | | |
| 4–6 | — | Reserved, should be cleared. | | | | | | | | | | | | | | | |
| 7–8 | Pause | Encoding of the link partner’s PAUSE capability is shown in the PAUSE encoding table below. For priority resolution information consult. This bit is read-only <table border="1" style="margin: 10px auto; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 15%;">PAUSE bit[8]</th> <th style="width: 15%;">ASM_DIR bit[7]</th> <th style="width: 70%;">Capability</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No PAUSE</td> </tr> <tr> <td>0</td> <td>1</td> <td>Asymmetric PAUSE toward link partner</td> </tr> <tr> <td>1</td> <td>0</td> <td>Symmetric PAUSE</td> </tr> <tr> <td>1</td> <td>1</td> <td>Both symmetric PAUSE and Asymmetric PAUSE toward local device</td> </tr> </tbody> </table> | PAUSE bit[8] | ASM_DIR bit[7] | Capability | 0 | 0 | No PAUSE | 0 | 1 | Asymmetric PAUSE toward link partner | 1 | 0 | Symmetric PAUSE | 1 | 1 | Both symmetric PAUSE and Asymmetric PAUSE toward local device |
| PAUSE bit[8] | ASM_DIR bit[7] | Capability | | | | | | | | | | | | | | | |
| 0 | 0 | No PAUSE | | | | | | | | | | | | | | | |
| 0 | 1 | Asymmetric PAUSE toward link partner | | | | | | | | | | | | | | | |
| 1 | 0 | Symmetric PAUSE | | | | | | | | | | | | | | | |
| 1 | 1 | Both symmetric PAUSE and Asymmetric PAUSE toward local device | | | | | | | | | | | | | | | |
| 9 | Half Duplex | Half-duplex capability. This bit is read-only. 0 Link partner is not capable of half-duplex mode. 1 Link partner is capable of half-duplex mode. | | | | | | | | | | | | | | | |

Table 19-129 describes the fields of the ANNPT register.

Table 19-129. ANNPT Field Descriptions

| Bits | Name | Description |
|------|--|--|
| 0 | Next Page | Next page indication. [Reference MII bit 7.15 in IEEE 802.3, 2000 Edition Clause 28.2.4] 0 Last page. 1 Additional next pages to follow. |
| 1 | — | Reserved. (Ignore on read) |
| 2 | Msg Page | Message page. [Reference MII bit 7.13] 0 Unformatted page. 1 Message page. |
| 3 | Ack2 | Acknowledge 2. Used by the next page function to indicate that the device has the ability to comply with the message. [Reference MII bit 7.12] 0 The local device cannot comply with message. 1 The local device complies with message. |
| 4 | Toggle | Toggle. Used to ensure synchronization with the link partner during next page exchange. This bit always takes the opposite value of the toggle bit of the previously exchanged link code word. The initial value in the first next page transmitted is the inverse of bit 11 in the base link code word. [Reference MII bit 7.11] This bit is read-only. 0 Toggle bit of the previously exchanged link code word was 1. 1 Toggle bit of the previously exchanged link code word was 0. |
| 5–15 | Message/ Un-formatted Code Field | Message pages are formatted pages that carry a pre-defined message code, which is enumerated in IEEE 802.3u/Annex 28C. Unformatted code fields take on an arbitrary value. [Reference MII field 7.10:0] |

19.5.4.3.7 AN Link Partner Ability Next Page Register (ANLPANP)

Figure 19-124 describes the definition for the ANLPANP register.




Figure 19-124. AN Link Partner Ability Next Page Register Definition

Table 19-130 describes the fields of the ANLPANP register.

Table 19-130. ANLPANP Field Descriptions

| Bits | Name | Description |
|------|-----------|---|
| 0 | Next Page | Next page. The link partner sets and clears this bit. 0 Last page from link partner 1 Additional next pages to follow |
| 1 | — | Reserved. (Ignore on read) |

Table 19-130. ANLPANP Field Descriptions (continued)

| Bits | Name | Description |
|------|--|---|
| 2 | Msg Page | Message page. 0 Unformatted page 1 Message page |
| 3 | Ack2 | Acknowledge 2. Indicates the link partner's ability to comply with the message. 0 Link partner cannot comply with message. 1 Link partner complies with message. |
| 4 | Toggle | Toggle. Used to ensure synchronization with the link partner during next page exchange. This bit always takes the opposite value of the toggle bit of the previously exchanged link code word. The initial value in the first next page transmitted is the inverse of bit 11 in the base link code word. This bit is read-only. 0 Toggle bit of the previously exchanged link code word was 1. 1 Toggle bit of the previously exchanged link code word was 0. |
| 5–15 | Message/ Un-formatted Code Field | Message pages are formatted pages that carry a pre-defined message code, which is enumerated in IEEE 802.3u/Annex 28C. Unformatted code fields take on an arbitrary value. |

19.5.4.3.8 Extended Status Register (EXST)

Figure 19-125 describes the definition for the EXST register.




Figure 19-125. Extended Status Register Definition

Table 19-131 describes the fields of the EXST register.

Table 19-131. EXST Field Descriptions

| Bits | Name | Description |
|------|------------|--|
| 0 | 1000X Full | 1000X full-duplex capability. Returns 1 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-X full-duplex mode. 1 PHY can operate in 1000BASE-X full-duplex mode. |
| 1 | 1000X Half | 1000X half-duplex capability. Returns 0 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-X half-duplex mode. 1 PHY can operate in 1000BASE-X half-duplex mode. |
| 2 | 1000T Full | 1000T full-duplex capability. Returns 1 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-T full-duplex mode. 1 PHY can operate in 1000BASE-T full-duplex mode. |
| 3 | 1000T Half | 1000T half-duplex capability. Returns 0 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-T half-duplex mode. 1 PHY can operate in 1000BASE-T half-duplex mode. |
| 4–15 | — | Reserved |

19.5.4.3.9 Jitter Diagnostics Register (JD)

Annex 36A in IEEE 802.3z describes several jitter test patterns. These can be configured to be sent by writing the jitter diagnostics register. See the register description for more information. It may be wise to auto-negotiate and advertise a remote fault signaling of offline prior to beginning the test patterns.

Figure 19-126 describes the definition for the JD register.




Figure 19-126. Jitter Diagnostics Register Definition

Table 19-132 describes the fields of the JD register.

Table 19-132. JD Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|--|-----------------------|--|-----------------------|--------|--------|--------|--|---|---|---|---|---|---|---|--|---|---|---|--|---|---|---|--|---|---|---|--|---|---|---|----------|---|---|---|----------|---|---|---|
| 0 | Jitter Enable | Jitter enable. This bit is cleared by default. 0 Normal transmit operation. 1 Enable the TBI to transmit the jitter test patterns defined in IEEE 802.3z 36A. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 1–3 | Jitter Select | <p>Selects the jitter pattern to be transmitted in diagnostics mode. Encoding of this field is shown in the following table. Default is 00.</p> <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 10px;"> <thead> <tr> <th style="width: 70%;">Jitter Pattern Select</th> <th style="width: 10%;">bit[1]</th> <th style="width: 10%;">bit[2]</th> <th style="width: 10%;">bit[3]</th> </tr> </thead> <tbody> <tr> <td>User defined uses custom jitter pattern, bits 6–15</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> </tr> <tr> <td>High frequency (+/- D21.5) 101010101010101010101010101010101010...</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> </tr> <tr> <td>Mixed frequency (+/- K28.5) 1111101011000001010011111010110000010100...</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> </tr> <tr> <td>Low frequency 1111100000111110000011111000001111100000...</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> </tr> <tr> <td>Complex pattern (10'h17c,10'h0c9,10'h0e5,10'h2a3, 10'h17c,...)</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> </tr> <tr> <td>Square Wave (- K28.7) 0011111000001111100000111110000011111000...</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> </tr> <tr> <td>Reserved</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> </tr> <tr> <td>Reserved</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> </tr> </tbody> </table> | Jitter Pattern Select | bit[1] | bit[2] | bit[3] | User defined uses custom jitter pattern, bits 6–15 | 0 | 0 | 0 | High frequency (+/- D21.5) 101010101010101010101010101010101010... | 0 | 0 | 1 | Mixed frequency (+/- K28.5) 1111101011000001010011111010110000010100... | 0 | 1 | 0 | Low frequency 1111100000111110000011111000001111100000... | 0 | 1 | 1 | Complex pattern (10'h17c,10'h0c9,10'h0e5,10'h2a3, 10'h17c,...) | 1 | 0 | 0 | Square Wave (- K28.7) 0011111000001111100000111110000011111000... | 1 | 0 | 1 | Reserved | 1 | 1 | 0 | Reserved | 1 | 1 | 1 |
| Jitter Pattern Select | bit[1] | bit[2] | bit[3] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| User defined uses custom jitter pattern, bits 6–15 | 0 | 0 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| High frequency (+/- D21.5) 101010101010101010101010101010101010... | 0 | 0 | 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Mixed frequency (+/- K28.5) 1111101011000001010011111010110000010100... | 0 | 1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Low frequency 1111100000111110000011111000001111100000... | 0 | 1 | 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Complex pattern (10'h17c,10'h0c9,10'h0e5,10'h2a3, 10'h17c,...) | 1 | 0 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Square Wave (- K28.7) 0011111000001111100000111110000011111000... | 1 | 0 | 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Reserved | 1 | 1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Reserved | 1 | 1 | 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4–5 | — | Reserved | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 6–15 | Custom Jitter Pattern | Used in conjunction with jitter (pattern) select and jitter (diagnostic) enable; set this field to the desired custom pattern which is continuously transmitted. Its default is 0x000. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

19.5.4.3.10 TBI Control Register (TBICON)

Figure 19-127 describes the definition for the TBICON register.




Figure 19-127. TBI Control Register Definition

Table 19-133 describes the fields of the TBICON register.

Table 19-133. TBICON Field Descriptions

| Bits | Name | Description |
|------|----------------|---|
| 0 | Soft_Reset | Soft reset. This bit is cleared by default. 0 Normal operation. 1 Resets the functional modules in the TBI. |
| 1 | — | Reserved. (Ignore on read) |
| 2 | Disable Rx Dis | Disable receive disparity. This bit is cleared by default. 0 Normal operation. 1 Disables the running disparity calculation and checking in the receive direction. |
| 3 | Disable Tx Dis | Disable transmit disparity. This bit is cleared by default. 0 Normal operation. 1 Disables the running disparity calculation and checking in the transmit direction. |
| 4–6 | — | Reserved |
| 7 | AN Sense | Auto-negotiation sense enable. This bit is cleared by default. 0 IEEE 802.3z Clause 37 behavior is desired, which results in the link not completing. 1 Allow the auto-negotiation function to sense either a Gigabit MAC in auto-negotiation bypass mode or an older Gigabit MAC without auto-negotiation capability. If sensed, auto-negotiation complete becomes true; however, the page received is low, indicating no page was exchanged. Management can then act accordingly. |
| 8–9 | — | Reserved |
| 10 | Clock Select | Clock select. This bit selects how the on-chip TBI PHY is clocked. This bit is cleared by default. 0 The TBI PHY is clocked by dual split-phase 62.5 MHz receive clocks. These external signals must be provided via TBI receive clock 0 (TSEC _n _RX_CLK) and TBI receive clock 1 (TSEC _n _TX_CLK). If operating in SGMII mode, clearing this bit effectively disables the TBI PHY clock. 1 The TBI PHY is clocked by a single 125 MHz receive clock (required for SGMII operation). This single clock, if operating in a non-SGMII (parallel) mode, must be provided via the TBI receive clock 0 (TSEC _n _RX_CLK) external signal. If operating in SGMII mode, this clock is provided on-chip by the SerDes block. |

Table 19-133. TBICON Field Descriptions (continued)

| Bits | Name | Description |
|-------|---------|---|
| 11 | MI Mode | This bit describes the configuration mode of the TBI. The user reads a 1 while the TBI is configured in GMII/MII mode (connected to a GMII/MII PHY) and a 0 while configured in TBI mode (connected to a 1000BASE-X SerDes). Its value is the inverse of ECNTRL[TBIM]. 0 TBI mode. 1 GMII mode. |
| 12–15 | — | Reserved |

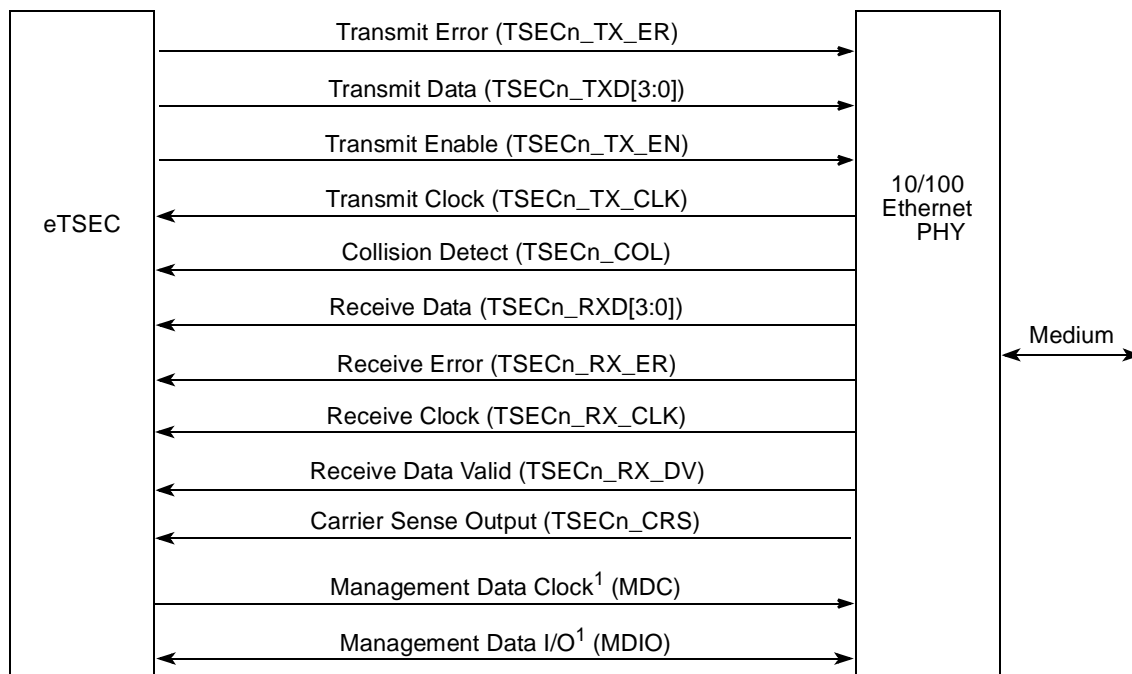
19.6 Functional Description

19.6.1 Connecting to Physical Interfaces on Ethernet

This section describes how to connect the eTSEC to various interfaces: MII, RMII, RGMII, and RTBI. To avoid confusion, all of the buses follow the bus conventions used in the IEEE 802.3 specification because the PHYs follow the same conventions. (For instance, in the bus TSEC_n_TXD[3:0], bit 3 is the msb and bit 0 is the lsb). If a mode does not use all input signals available to a particular eTSEC, those inputs that are not used must be pulled low on the board.

19.6.1.1 Media-Independent Interface (MII)

This section describes the media-independent interface (MII) intended to be used between the PHYs and the eTSEC. Figure 19-128 depicts the basic components of the MII including the signals required to establish eTSEC module connection with a PHY.



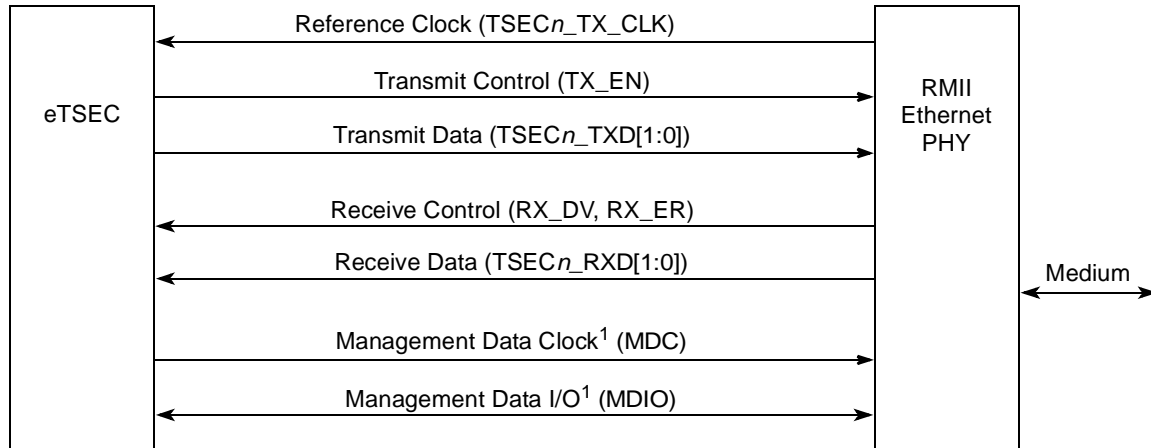
¹ The management signals (MDC and MDIO) are common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

Figure 19-128. eTSEC-MII Connection

An MII interface has 18 signals (including the MDC and MDIO signals), as defined by the IEEE 802.3u standard, for connecting to an Ethernet PHY.

19.6.1.2 Reduced Media-Independent Interface (RMII)

This section describes the reduced media-independent interface (RMII) intended to be used between the PHYs and the GMII MAC. The RMII is a reduced-pin alternative to the IEEE 802.3u MII. The RMII reduces the number of signals required to interconnect the MAC and the PHY from a maximum of 18 signals (MII) to 10 signals. To accomplish this objective, the data paths are halved in width and clocked at twice the MII clock frequency, while clocks, carrier sense and error signals have been partly combined. For 100 Mbps operation, the reference clock operates at 50 MHz, whereas for 10 Mbps operation, the clock remains at 50MHz, but only every 10th cycle is used. Figure 19-129 depicts the basic components of the reduced media-independent interface and the signals required to establish an eTSEC's connection with a PHY. The RMII is implemented as defined by the RMII Specification of the RMII Consortium, as of March 20, 1998.

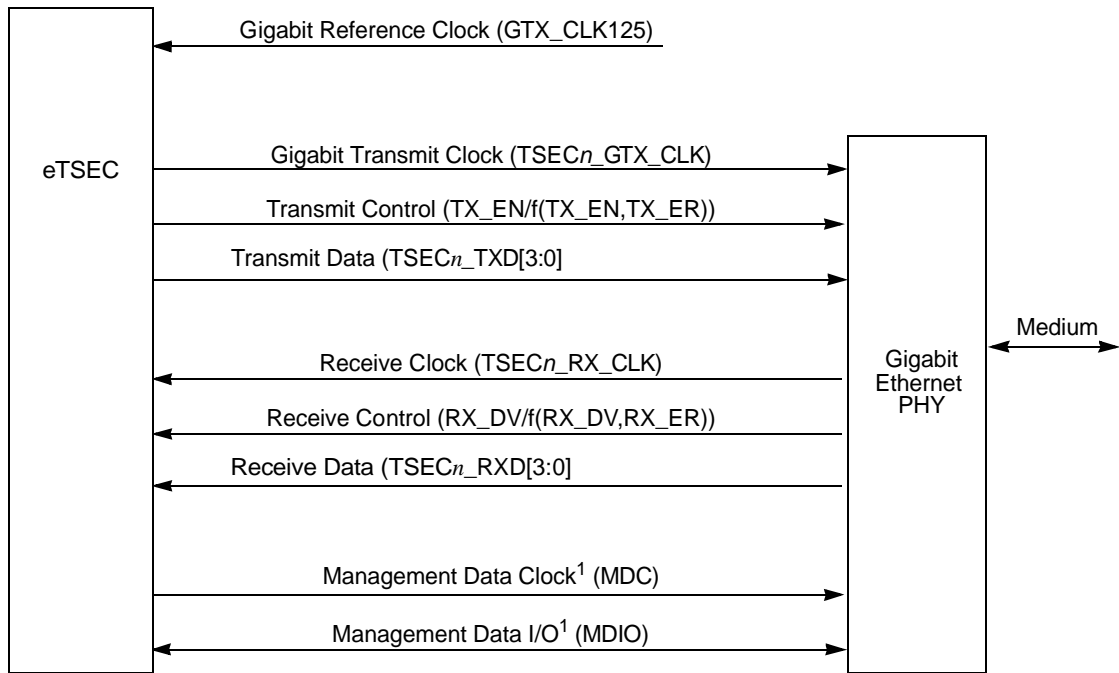


¹ The management signals (MDC and MDIO) are common to all of the Ethernet controllers module connections in the system, assuming that each PHY has a different management address.

Figure 19-129. eTSEC-RMII Connection

19.6.1.3 Reduced Gigabit Media-Independent Interface (RGMI)

This section describes the reduced gigabit media-independent interface (RGMI) intended to be used between the PHYs and the GMII MAC. The RGMI is an alternative to the IEEE802.3u MII, the IEEE802.3z GMII. The RGMI reduces the number of signals required to interconnect the MAC and the PHY from a maximum of 28 signals (GMII) to 15 signals (GTX_CLK125 included) in a cost effective and technology independent manner. To accomplish this objective, the data paths and all associated control signals are multiplexed using both edges of the clock. For gigabit operation, the clocks operate at 125MHz, and for 10/100 operation, the clocks operate at 2.5 MHz or 25 MHz, respectively. Note that the GTX_CLK125 input must be provided at 125 MHz for an RGMI interface, regardless of operation speed (1 Gbps, 100 Mbps, or 10 Mbps). [Figure 19-130](#) depicts the basic components of the gigabit reduced media-independent interface and the signals required to establish the gigabit Ethernet controllers' module connection with a PHY. The RGMI is implemented as defined by the RGMI specification Version 1.2a 9/22/00.

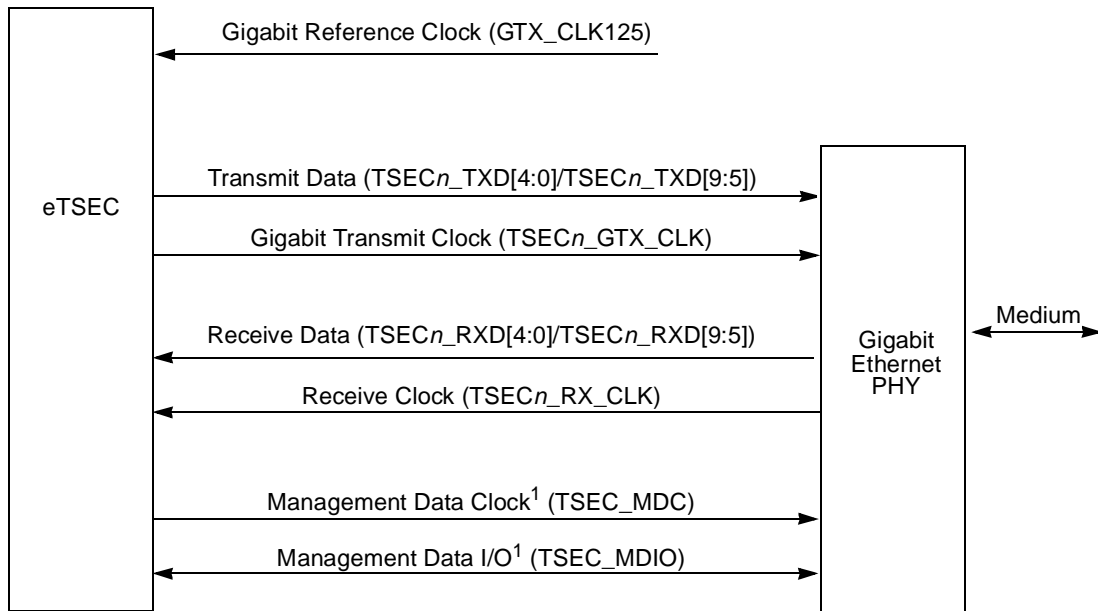


¹ The management signals (MDC and MDIO) are common to all of the gigabit Ethernet controllers module connections in the system, assuming that each PHY has a different management address.

Figure 19-130. eTSEC-RGMII Connection

19.6.1.4 Reduced Ten-Bit Interface (RTBI)

This section describes the reduced ten-bit interface (RTBI) intended to be used between the PHYs and the eTSEC to implement a reduced-pin count version of a SerDes interface for optical-fiber devices in 1000BASE-SX/LX applications. [Figure 19-131](#) depicts the basic components of the RTBI including the signals required to establish eTSEC module connection with a PHY. Note that in RTBI the eTSEC immediately begins auto-negotiation with the SerDes.



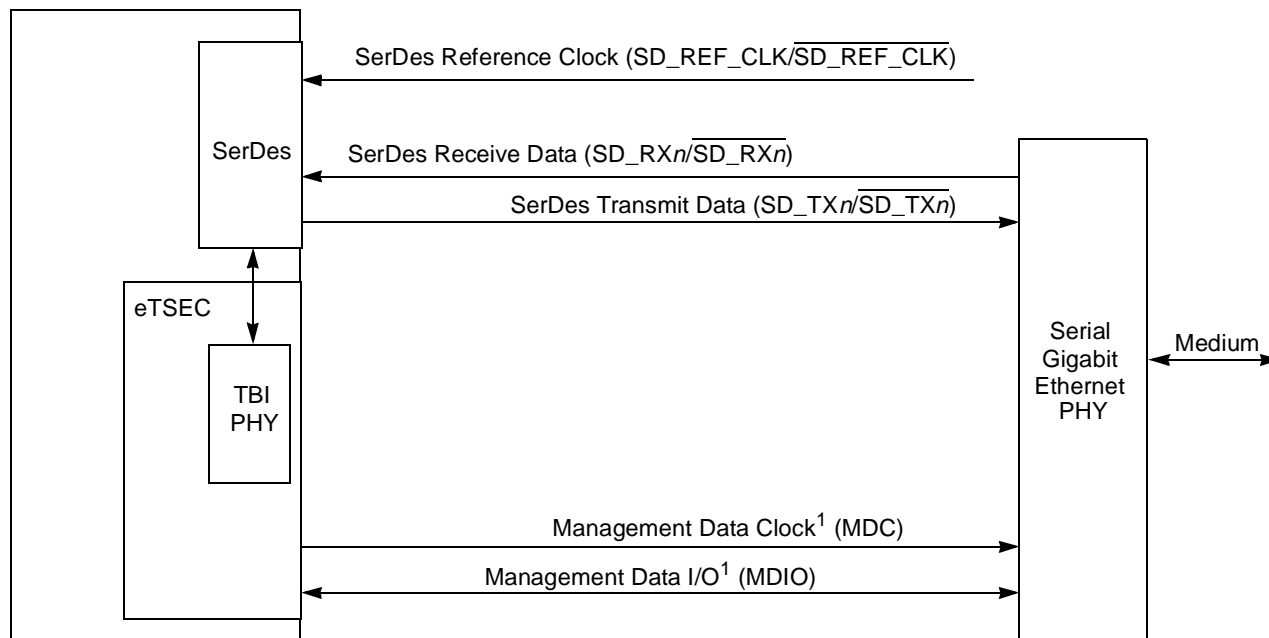
¹ The management signals (MDC and MDIO) are common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

Figure 19-131. eTSEC-RTBI Connection

A RTBI interface has 15 signals (GE_GTX_CLK125 included), as defined by the RGMII specification Version 1.2a 9/22/00, and is intended to be an alternative to the IEEE 802.3u MII, the IEEE 802.3z GMII and the TBI standard for connecting to an Ethernet PHY.

19.6.1.5 Serial Gigabit Media-Independent Interface (SGMII)

This section describes the serial gigabit media-independent interface (SGMII) intended to be used between a SerDes PHY and the eTSEC to implement a serial gigabit version of a media-independent interface. [Figure 19-132](#) depicts the basic components of the SGMII including the signals required to establish eTSEC module connection with a PHY. Note that in SGMII the eTSEC utilizes the on-chip TBI PHY in addition to the SerDes interface.



¹ The management signals (MDC and MDIO) may be common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

Figure 19-132. eTSEC-SGMII Connection

19.6.1.6 Ethernet Physical Interfaces Signal Summary

Table 19-134 describes the signal multiplexing for MII and RMII interfaces.

Table 19-134. MII and RMII Signals Multiplexing

| eTSEC Signals | | | MII Interface | | | RMII Interface | | |
|---------------------|-----|----------------|--------------------|-----|----------------|--------------------|-----|----------------|
| Frequency [MHz] 125 | | | Frequency [MHz] 25 | | | Frequency [MHz] 50 | | |
| Voltage[V] 3.3/2.5 | | | Voltage[V] 3.3 | | | Voltage[V] 3.3 | | |
| Signals (TSECn_) | I/O | No. of Signals | Signals (TSECn_) | I/O | No. of Signals | Signals (TSECn_) | I/O | No. of Signals |
| GTX_CLK | O | 1 | | | | | | |
| TX_CLK | I | 1 | TX_CLK | I | 1 | REF_CLK | I | 1 |
| TxD[0] | O | 1 | TxD[0] | O | 1 | TxD[0] | O | 1 |
| TxD[1] | O | 1 | TxD[1] | O | 1 | TxD[1] | O | 1 |
| TxD[2] | O | 1 | TxD[2] | O | 1 | | | |
| TxD[3] | O | 1 | TxD[3] | O | 1 | | | |
| TX_EN | O | 1 | TX_EN | O | 1 | TX_EN | O | 1 |
| TX_ER | O | 1 | TX_ER | O | 1 | | | |
| RX_CLK | I | 1 | RX_CLK | I | 1 | | | |
| RxD[0] | I | 1 | RxD[0] | I | 1 | RxD[0] | I | 1 |

Table 19-134. MII and RMIi Signals Multiplexing (continued)

| eTSEC Signals | | | MII Interface | | | RMIi Interface | | |
|---------------------|---|----|--------------------|---|----|--------------------|---|---|
| Frequency [MHz] 125 | | | Frequency [MHz] 25 | | | Frequency [MHz] 50 | | |
| Voltage[V] 3.3/2.5 | | | Voltage[V] 3.3 | | | Voltage[V] 3.3 | | |
| RxD[1] | I | 1 | RxD[1] | I | 1 | RxD[1] | I | 1 |
| RxD[2] | I | 1 | RxD[2] | I | 1 | | | |
| RxD[3] | I | 1 | RxD[3] | I | 1 | | | |
| RX_DV | I | 1 | RX_DV | I | 1 | CRS_DV | I | 1 |
| RX_ER | I | 1 | RX_ER | I | 1 | RX_ER | I | 1 |
| COL | I | 1 | COL | I | 1 | | | |
| CRS | I | 1 | CRS | I | 1 | | | |
| Sum | | 25 | Sum | | 16 | Sum | | 8 |

Table 19-135 describes the signal multiplexing for RGMII and RTBI interfaces.

Table 19-135. RGMII and RTBI Signals Multiplexing

| eTSEC Signals | | | RGMII Interface | | | RTBI Interface | | |
|------------------------------|-----|----------------|------------------------------|-----|----------------|------------------------------|-----|----------------|
| Frequency [MHz] 125 | | | Frequency [MHz] 125 | | | Frequency [MHz] 62.5 | | |
| Voltage[V] 3.3/2.5 | | | Voltage[V] 2.5 | | | Voltage[V] 2.5 | | |
| Signals (TSEC _n) | I/O | No. of Signals | Signals (TSEC _n) | I/O | No. of Signals | Signals (TSEC _n) | I/O | No. of Signals |
| GTX_CLK | O | 1 | GTX_CLK | O | 1 | GTX_CLK | O | 1 |
| TX_CLK | I | 1 | | | | | | |
| TxD[0] | O | 1 | TxD[0] | O | 1 | TCG[0]/TCG[5] | O | 1 |
| TxD[1] | O | 1 | TxD[1] | O | 1 | TCG[1]/TCG[6] | O | 1 |
| TxD[2] | O | 1 | TxD[2] | O | 1 | TCG[2]/TCG[7] | O | 1 |
| TxD[3] | O | 1 | TxD[3] | O | 1 | TCG[3]/TCG[8] | O | 1 |
| TX_EN | O | 1 | TX_CTL (TX_EN/TX_ERR) | O | 1 | TCG[4]/TCG[9] | O | 1 |
| TX_ER | O | 1 | | | | | | |
| RX_CLK | I | 1 | RX_CLK | I | 1 | RX_CLK | I | 1 |
| RxD[0] | I | 1 | RxD[0]/RxD[4] | I | 1 | RCG[0]/RCG[5] | I | 1 |
| RxD[1] | I | 1 | RxD[1]/RxD[5] | I | 1 | RCG[1]/RCG[6] | I | 1 |
| RxD[2] | I | 1 | RxD[2]/RxD[6] | I | 1 | RCG[2]/RCG[7] | I | 1 |
| RxD[3] | I | 1 | RxD[3]/RxD[7] | I | 1 | RCG[3]/RCG[8] | I | 1 |
| RX_DV | I | 1 | RX_CTL (RX_DV/RX_ERR) | I | 1 | RCG[4]/RCG[9] | I | 1 |
| RX_ER | I | 1 | | | | | | |

Table 19-135. RGMII and RTBI Signals Multiplexing (continued)

| eTSEC Signals | | | RGMII Interface | | | RTBI Interface | | |
|---------------------|-----|----------------|---------------------|-----|----------------|----------------------|-----|----------------|
| Frequency [MHz] 125 | | | Frequency [MHz] 125 | | | Frequency [MHz] 62.5 | | |
| Voltage[V] 3.3/2.5 | | | Voltage[V] 2.5 | | | Voltage[V] 2.5 | | |
| Signals (TSECn_) | I/O | No. of Signals | Signals (TSECn_) | I/O | No. of Signals | Signals (TSECn_) | I/O | No. of Signals |
| COL | I | 1 | | | | | | |
| CRS | I | 1 | | | | | I | |
| Sum | | 25 | Sum | | 12 | Sum | | 12 |

Table 19-136 describes the signalling for SGMII interfaces. SGMII communication using the eTSEC is accomplished through the SerDes interface.

Table 19-136. SGMII Signalling

| Signals | I/O | No. of Signals | Function |
|-------------|-----|----------------|------------------------------------|
| SDn_RX[n] | I | 2 | SGMII receive data (differential) |
| SDn_TX[n] | O | 2 | SGMII transmit data (differential) |
| SDn_REF_CLK | I | 2 | Reference clock (differential) |
| Sum | | 6 | — |

Table 19-137 describes the signals shared by all interfaces.

Table 19-137. Shared Signals

| Signals | I/O | No. of Signals | Function |
|------------|-----|----------------|----------------------------|
| MDIO | I/O | 1 | Management interface I/O |
| MDC | O | 1 | Management interface clock |
| GTX_CLK125 | I | 1 | Reference clock |
| Sum | | | — |

19.6.2 Gigabit Ethernet Controller Channel Operation

This section describes the operation of the eTSEC. First, the software initialization sequence is described. Next, the software (Ethernet driver) interface for transmitting and receiving frames is reviewed. Frame filtering and receive filing algorithm features are also discussed. The section concludes with interrupt handling, inter-packet gap time, and loop back descriptions.

19.6.2.1 Initialization Sequence

This sections describes which registers are reset due to a hard or software reset and what registers the user must initialize prior to enabling the eTSEC.

19.6.2.1.1 Hardware Controlled Initialization

A hard reset occurs when the system powers up. All eTSEC's registers and control logic are reset to their default states after a hard reset has occurred. In this state, each eTSEC behaves like a PowerQUICC II Pro device, except for the absence of out-of-sequence TxBD features. That is, initially TCP/IP off-load is disabled and only single RxBD and TxBD rings are accessible.

19.6.2.1.2 User Initialization

After the system has undergone a hard reset, software must initialize certain basic eTSEC registers. Other registers can also be initialized during this time, but they are optional and must be determined based on the requirements of the system. See [Table 19-3](#) for the register list. [Table 19-138](#) describes the minimum steps for register initialization.

Table 19-138. Steps for Minimum Register Initialization

| Description |
|--|
| 1. Set and clear MACCFG1 [Soft_Reset] |
| 2. Initialize MACCFG2 |
| 3. Initialize MAC station address |
| 4. Set up the PHY using the MII Mgmt Interface |
| 5. Configure the GMII |
| 6. Clear IEVENT |
| 7. Initialize IMASK |
| 8. Initialize RCTRL |
| 9. Initialize DMACTRL |

After the initialization of registers is performed, the user must execute the following steps in the order described below to bring the eTSEC into a functional state (out of reset):

1. Write to the MACCFG1 register and set the appropriate bits. These need to include RX_EN and TX_EN. To enable flow control, Rx_Flow and Tx_Flow should also be set.
2. For the transmission of Ethernet frames, TxBDs must first be built in memory, linked together as a ring, and pointed to by the TBASE_n registers. A minimum of two buffer descriptors per ring is required, unless the ring is disabled. Setting the ring to a size of one causes the same frame to be transmitted twice. If TCP/IP off-load is to be enabled, the TxBD[TOE] bit must be set for each frame.

3. Likewise, for the reception of Ethernet frames, the receive queue (or queues) must be ready, with its RxBD pointed to by the RBASE n registers. If TCP/IP off-load is to be enabled, RCTRL[PRSDPEP] must be set to the required off-load level. Both transmit and receive can be gracefully stopped after transmission and reception begins.
4. Clearing DMACTRL[GTS] triggers the transmission of frame data if the transmitter had been previously stopped. The DMACTRL[GRS] must be cleared if the receiver had been previously stopped. Refer to the DMACTRL register section, and [Section 19.6.7.1, “Data Buffer Descriptors,”](#) for more information.

19.6.2.2 Soft Reset and Reconfiguring Procedure

Before issuing a soft-reset to and/or reconfiguring the MAC with new parameters, the user must properly shutdown the DMA and make sure it is in an idle state for the entire duration. User must gracefully stop the DMA by setting both GRS and GTS bits in the DMACTRL register, then wait for both GRSC and GTSC bits to be set in the IEVENT register before resetting the MAC or changing parameters. Both GRS and GTS bits must be cleared before re-enabling the MAC to resume the DMA.

During the MAC configuration, if a new set of Tx buffer descriptors are used, the user must load the pointers into the TBASE registers. Likewise if a new set of Rx buffer descriptors are used, the RBASE registers must be written with new pointers.

Following is a procedure to gracefully reset and reconfigure the MAC:

1. Set GRS/GTS bits in DMACTRL register
2. Poll GRSC/GTSC bits in IEVENT register until both are set
3. Set SOFT_RESET bit in MACCFG1 register (Note that SOFT_RESET must remain set for at least 3 TX clocks before proceeding.)
4. Clear SOFT_RESET bit in MACCFG1 register
5. Load TBASE0–TBASE7 with new Tx BD pointers
6. Load RBASE0–RBASE7 with new Rx BD pointers
7. Setup other MAC registers (MACCFG2, MAXFRM, and so on)
8. Setup group address hash table (GADDR0–GADDR15) if address filtering is required
9. Setup receive frame filter table (through RQFAR, RQFCR, and RQFPR) if filtering to multiple RxBD rings is required
10. Setup WWR, WOP, TOD bits in DMACTRL register
11. Enable transmit queues in TQUEUE, and ensure that the transmit scheduling mode is correctly set in TCTRL.
12. Enable receive queues in RQUEUE, and optionally set TOE functionality in RCTRL.
13. Clear THLT and TXF bits in TSTAT register by writing 1 to them
14. Clear QHLT and RXF bits in RSTAT register by writing 1 to them.
15. Clear GRS/GTS bits in DMACTRL (do not change other bits)
16. Enable Tx_EN/Rx_EN in MACCFG1 register

19.6.2.3 Gigabit Ethernet Frame Transmission

The Ethernet transmitter requires little core intervention. After the software driver initializes the system, the eTSEC begins to poll the first transmit buffer descriptor (TxBD) in TxBD ring 0 every 512 transmit clocks. If TxBD[R] is set, and the TxBD ring is scheduled for transmission, the eTSEC begins copying the associated transmit buffer from memory to its Tx FIFO. The transmitter takes data from the Tx FIFO and transmits data to the MAC. The MAC transmits the data through the GMII interface to the physical media. The transmitter, once initialized, runs until the end-of-frame (EOF) condition is detected unless a collision within the collision window occurs (half-duplex mode) or an abort condition is encountered.

If the user has a frame ready to transmit, setting the DMACTRL[TOD] eliminates waiting for the next poll and a DMA transfer of the transmit data buffers can begin immediately. The transmission begins once all data for the frame is loaded into the Tx FIFO or sufficient transmit data (determined by the Tx FIFO threshold register) is in the Tx FIFO. If the line is not busy, the MAC transmit logic asserts TX_EN and sends the 7-octet preamble sequence, 1-octet start of frame delimiter, and frame information in that order. If the line is busy, the controller waits for the carrier sense signal, CRS, to remain inactive for 60 bit times (60 clocks) and transmission begins after an additional 36 bit times (96 bit times after CRS became active). In full-duplex mode, because collisions are ignored, frame transmission maintains only the interframe gap (96 bit times) regardless of CRS.

In half-duplex mode (MACCFG2[Full Duplex] is cleared) the MAC defers transmission if the line is busy (CRS asserted). Before transmitting, the MAC waits for carrier sense to become inactive, at which point it then determines if CRS remains negated for 60 clocks. If so, transmission begins after an additional 36 bit times (96 bit times after CRS originally became negated). If CRS continues to be asserted, the MAC follows a specified back-off procedure and tries to retransmit the frame until the retry limit is reached. Data stored in the Tx FIFO is re-transmitted in case of a collision. This avoids unnecessary memory traffic.

The transmitter also monitors for an abort condition and terminates the current frame if an abort condition is encountered. In full-duplex mode the protocol is independent of network activity, and only the transmit inter-frame gap must be enforced.

The transmitter implements full-duplex flow control. If a flow control frame is received, the MAC does not service the transmitter's request to send data until the pause duration is over. If the MAC is currently sending data after a pause frame has been received and processed, the MAC finishes sending the current frame, then suspends subsequent frames (except a pause frame) until the pause duration is over. In addition, the transmitter supports transmission of flow control frames through TCTRL[TFC_PAUSE]. The transmit pause frame is generated internally based on the PAUSE register that defines the pause value to be sent. Note that it is possible to send a pause frame while the pause timer has not expired.

The MAC automatically appends FCS (32-bit CRC) bytes to the frame if any of the following values are set:

- TxBD[PAD/CRC] is set in first TxBD
- TxBD[TC] is set in first TxBD
- MACCFG2[PAD/CRC] is set
- MACCFG2[CRC] is set

The TX_EN is negated after the FCS is sent. This notifies the PHY of the need to generate the illegal Manchester encoding that signifies the end of an Ethernet frame. Following the transmission of the FCS,

the Ethernet controller writes the frame status bits into the BD and clears TxBD[R]. If the end of the current buffer is reached and TxBD[L] is cleared (a frame is comprised of multiple buffer descriptors), only TxBD[R] is cleared.

For both half- and full-duplex modes, an interrupt can be issued depending on TxBD[I]. The Ethernet controller then proceeds to the next TxBD in the table. In this way, the core can be interrupted after each frame, after each buffer, or after a specific buffer is sent. If TxBD[PAD/CRC] is set, the Ethernet controller pads any frame shorter than 64 bytes with zero bytes to make up the minimum length.

To pause transmission, or rearrange the transmit queue, set DMACTRL[GTS]. This can be useful for transmitting expedited data ahead of previously linked buffers or for error situations. If this bit is set, the eTSEC transmitter performs a graceful transmit stop. The Ethernet controller stops immediately if no transmission is in progress or continues transmission until all queued frames in the Tx FIFO have been disposed of. The IEVENT[GTSC] interrupt occurs once the graceful transmit stop operation is completed. After the DMACTRL[GTS] is cleared, the eTSEC resumes transmission with the next frame.

While the eTSEC is in 10/100Mbps mode it sends bytes least-significant nibble first and each nibble is sent lsb first. While it is in 1000Mbps mode it sends bytes LSB first.

19.6.2.4 Gigabit Ethernet Frame Reception

The eTSEC Ethernet receiver is designed to work with little core intervention and can perform data extraction, address recognition, CRC checking, short frame checking, and maximum frame-length checking.

After a hardware reset, the software driver clears the RSTAT register and sets MACCFG1[RX_EN]. The Ethernet receiver is enabled and immediately starts processing receive frames. The MAC checks for when TSEC_n_RX_DV is asserted and as long as TSEC_n_COL remains negated (full-duplex mode ignores TSEC_n_COL), the MAC looks for the start of a frame by searching for a valid preamble/SFD (start of frame delimiter) header, which is stripped (unless MACCFG2[PreAM RxEN] is set) and the frame begins to be processed. If a valid header is not found, the frame is ignored.

If the receiver detects the first bytes of a frame, the eTSEC controller begins to perform the frame recognition function through destination address (DA) recognition (see [Section 19.6.2.7, “Frame Recognition”](#)). Based on this match the frame can be accepted or rejected. The receiver can filter frames based on individual (unicast), group (multicast), and broadcast addresses. Because Ethernet receive frame data is not written to memory until the internal frame recognition algorithm is complete, system bus usage is not wasted on frames unwanted by this station.

If a frame is accepted, the Ethernet controller fetches the receive buffer descriptor (RxBD) from either queue 0 or the queue determined by the filer. If the RxBD is not being used by software (RxBD[E] is set), the eTSEC starts transferring the incoming frame. RxBD[F] is set for the first RxBD used for any particular receive frame. If the current RxBD is not available for the received frame, a receive busy error condition is raised in IEVENT[BSY].

After the buffer is filled, the eTSEC clears RxBD[E] and, if RxBD[I] is set, generates an interrupt. If the incoming frame is larger than the buffer, the Ethernet controller fetches the next RxBD in the table. If it is empty, the controller continues receiving the rest of the frame. In half-duplex mode, if a collision is

detected during the frame, no RxBDs are used; thus, no collision frames are presented to the user except late collisions, which indicate LAN problems.

The RxBD length is determined by the MRBL field in the maximum receive buffer length register (MRBLR). The smallest valid value is 64 bytes, with larger values being some integral multiple of 64 bytes. During reception, the Ethernet controller checks for frames that are too short or too long. After the frame ends (CRS is negated), the receive CRC field is checked and written to the data buffer. The data length written to the last RxBD in the Ethernet frame is the length of the entire frame, which enables the software to recognize an oversized frame condition.

Receive frames are not truncated when they exceed maximum frame bytes in the MAC's maximum frame register if MACCFG2[Huge Frame] is set, yet the babbling receiver error interrupt occurs (IEVENT[BABR] is set) and RxBD[LG] is set.

After the receive frame is complete, the Ethernet controller sets RxBD[L], updates the frame status bits in the RxBD, and clears RxBD[E]. If RxBD[I] is set, the Ethernet controller next generates an interrupt (that can be masked) indicating that a frame was received and is in memory. The Ethernet controller then waits for a new frame.

To interrupt reception or rearrange the receive queue, DMACTRL[GRS] must be set. If this bit is set, the eTSEC receiver performs a graceful receive stop. The Ethernet controller stops immediately if no frames are being received or continues receiving until the current frame either finishes or an error condition occurs. The IEVENT[GRSC] interrupt event is signaled after the graceful receive stop operation is completed. While in this mode the user can write to registers that are accessible to both the user and the eTSEC hardware without fear of conflict, and finally clear IEVENT[GRSC]. After DMACTRL[GRS] is cleared, the eTSEC scans the input data stream for the start of a new frame (preamble sequence and start of frame delimiter), it resumes receiving, and the first valid frame received is placed in the next available RxBD.

19.6.2.5 Ethernet Preamble Customization

By default eTSEC generates a standard Ethernet preamble sequence prior to transmitting frames. However, the user can substitute a custom preamble sequence for the purpose of controlling switching equipment at the receiver, particularly at 100/1000Mbps speeds.; in any RMII mode only the standard preamble can be transmitted

eTSEC normally searches for and discards the standard Ethernet preamble sequence upon receiving frames. Part of the received preamble sequence can be optionally recovered and returned as part of the frame data, making it visible to user software. Note however, that, preamble cannot be recovered in any RMII mode. Note that it is also possible for the first two bytes of custom preamble (PreOct0 and PreOct1) to be lost in during conversion to ten-bit code groups in the PCS sub-layer. Thus is it recommended that any custom preamble start at PreOct2.

19.6.2.5.1 User-Defined Preamble Transmission

To substitute a custom preamble, the user must ensure that:

- MACCFG2[PreAm TxEN] bit is set
- The first TxBD of every frame containing a custom preamble has its PRE bit set

- An 8-byte custom preamble sequence appears before the Ethernet DA field in the first transmit data buffer

The definition of the 8-byte custom preamble sequence is shown in [Figure 19-133](#).

| Byte Offsets | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|--------------|---------|---|---|---|---|---|---|---------|---|---|----|----|----|----|----|----|
| 0–1 | PreOct0 | | | | | | | PreOct1 | | | | | | | | |
| 2–3 | PreOct2 | | | | | | | PreOct3 | | | | | | | | |
| 4–5 | PreOct4 | | | | | | | PreOct5 | | | | | | | | |
| 6–7 | PreOct6 | | | | | | | | | | | | | | | |

Figure 19-133. Definition of Custom Preamble Sequence

The fields of the custom preamble sequence are described in [Table 19-139](#). It should be noted that use of preamble octets matching the standard start of frame delimiter (0xD5) can be expected to trigger premature frame reception by the receiving station.

Table 19-139. Custom Preamble Field Descriptions

| Bytes | Bits | Name | Description |
|-------|------|---------|--|
| 0–1 | 0–7 | PreOct0 | Octet #0 of custom transmit preamble. This is the first octet of preamble sent. |
| | 8–15 | PreOct1 | Octet #1 of custom transmit preamble. This is the second octet of preamble sent. |
| 2–3 | 0–7 | PreOct2 | Octet #2 of custom transmit preamble. This is the third octet of preamble sent. |
| | 8–15 | PreOct3 | Octet #3 of custom transmit preamble. This is the fourth octet of preamble sent. |
| 4–5 | 0–7 | PreOct4 | Octet #4 of custom transmit preamble. This is the fifth octet of preamble sent. |
| | 8–15 | PreOct5 | Octet #5 of custom transmit preamble. This is the sixth octet of preamble sent. |
| 6–7 | 0–7 | PreOct6 | Octet #6 of custom transmit preamble. This is the seventh octet of preamble sent. The last octet (the start of frame delimiter) is generated by the MAC automatically. |
| | 8–15 | — | Reserved; should be cleared. |

19.6.2.5.2 User-Visible Preamble Reception

To return the received preamble, the user must ensure that:

- MACCFG2[PreAm RxEN] bit is set
- Space for an 8-byte preamble sequence is allowed before the Ethernet DA field in the first receive data buffer of each frame

The definition of the 8-byte received preamble sequence is shown in [Figure 19-134](#).

| Byte Offsets | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|--------------|---------|---|---|---|---|---|---|---------|---|---|----|----|----|----|----|----|
| 0–1 | PreOct0 | | | | | | | PreOct1 | | | | | | | | |
| 2–3 | PreOct2 | | | | | | | PreOct3 | | | | | | | | |
| 4–5 | PreOct4 | | | | | | | PreOct5 | | | | | | | | |
| 6–7 | PreOct6 | | | | | | | | | | | | | | | |

Figure 19-134. Definition of Received Preamble Sequence

The fields of the received preamble sequence are described in [Table 19-140](#). Should the received preamble be shorter than the 7-octet sequence defined by IEEE Std. 802.3, initial bytes of the received preamble sequence hold undefined values. The standard start of frame delimiter (0xD5) is always omitted. Note that preamble extraction is not possible in RMII mode.

Table 19-140. Received Preamble Field Descriptions

| Bytes | Bits | Name | Description |
|-------|------|---------|--|
| 0–1 | 0–7 | PreOct0 | Octet #0 of received preamble. This is the first octet of preamble received. |
| | 8–15 | PreOct1 | Octet #1 of received preamble. This is the second octet of preamble received. |
| 2–3 | 0–7 | PreOct2 | Octet #2 of received preamble. This is the third octet of preamble received. |
| | 8–15 | PreOct3 | Octet #3 of received preamble. This is the fourth octet of preamble received. |
| 4–5 | 0–7 | PreOct4 | Octet #4 of received preamble. This is the fifth octet of preamble received. |
| | 8–15 | PreOct5 | Octet #5 of received preamble. This is the sixth octet of preamble received. |
| 6–7 | 0–7 | PreOct6 | Octet #6 of received preamble. This is the seventh octet of preamble received. The last octet (the start of frame delimiter) is discarded. |
| | 8–15 | — | Reserved |

19.6.2.6 RMON Support

Using promiscuous mode, the eTSEC can automatically gather network statistics required for remote network interface monitoring. The RMON MIB group 1, RMON MIB group 2, RMON MIB group 3, RMON MIB group 9, RMON MIB2, and the IEEE 802.3 Ethernet MIB are supported. For RMON statistics and their corresponding counters, see the memory map.

19.6.2.7 Frame Recognition

The Ethernet controller performs frame recognition using destination address (DA) recognition. A frame can be rejected or accepted based on the outcome.

19.6.2.7.1 Destination Address Recognition and Frame Filtering

The eTSEC can perform layer 2 frame filtering on the basis of destination Ethernet address (DA), as illustrated by the flowchart in [Figure 19-135](#).

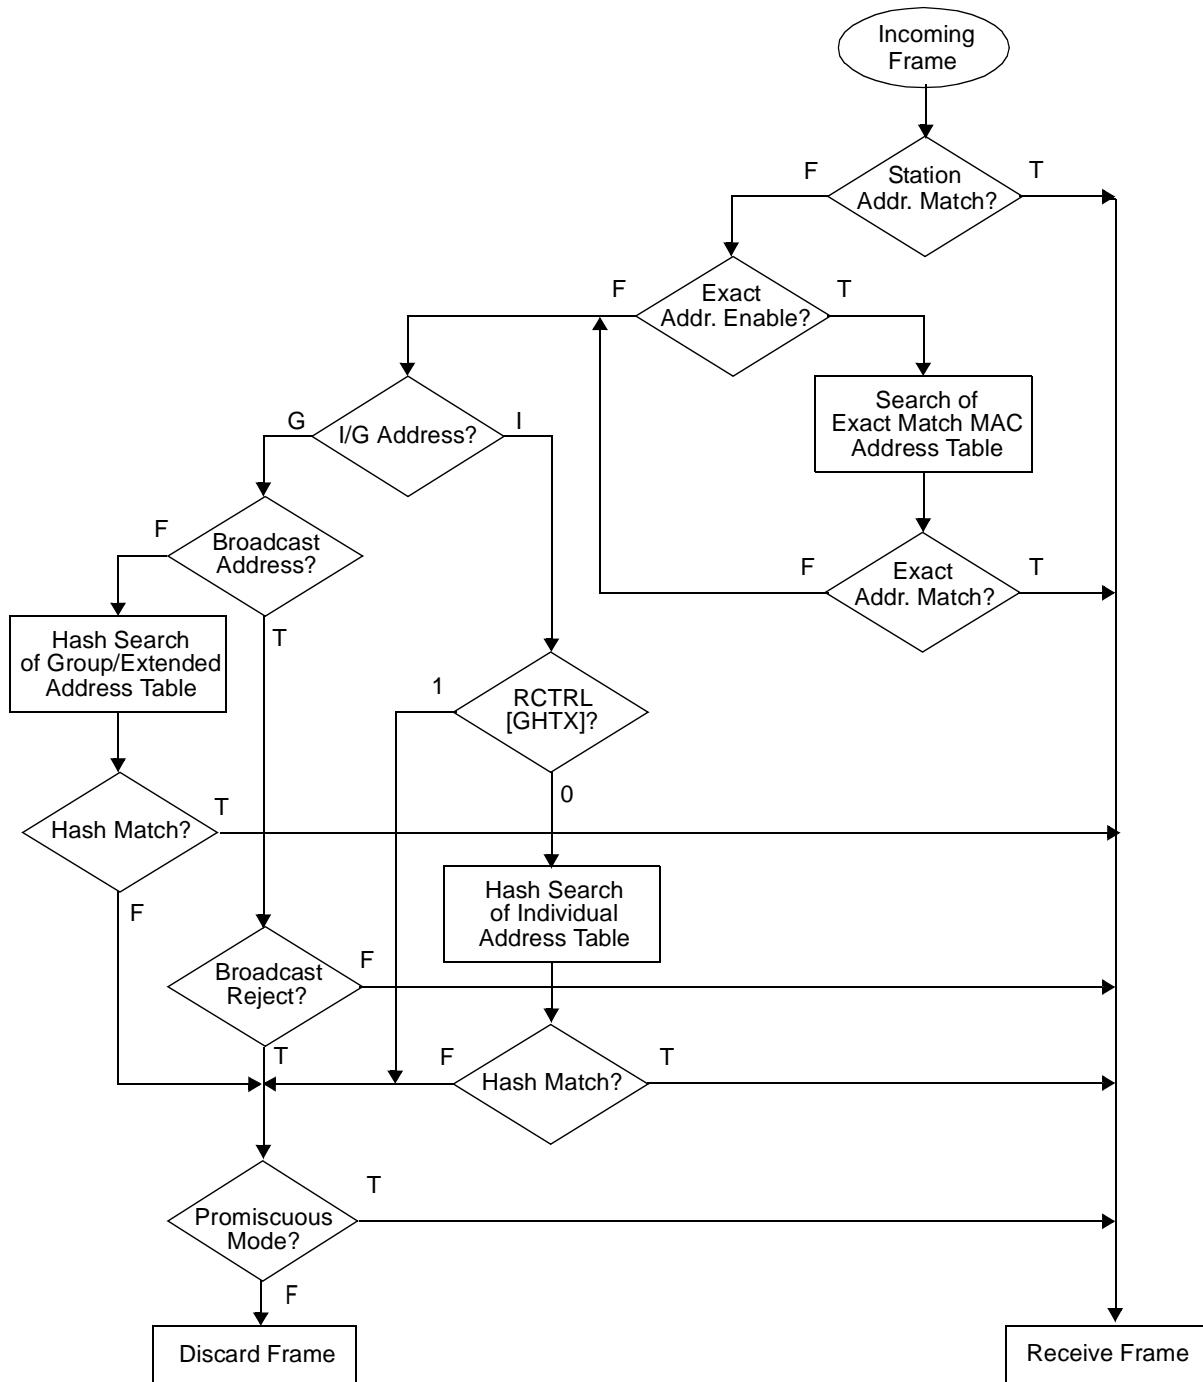


Figure 19-135. Ethernet Address Recognition Flowchart

In promiscuous mode, the eTSEC accepts all received frames regardless of DA. Note, however, that Ethernet frame filtering simply restricts the traffic seen by the receive queue filter. Therefore even in

promiscuous mode it remains possible to program the filter to reject frames based on their higher-layer header contents.

In the case of an individual address, the DA field of the received frame is compared with the physical address that the user programs in the station address registers (MACSTNADDR1 and MACSTNADDR2). If the DA does not match the station address, and exact MAC address matching is enabled through RCTRL[EMEN], the controller performs address recognition on the multiple MAC addresses written to the MACxADDR1 and MACxADDR2 registers. These virtual addresses give a particular eTSEC the ability to mirror other MACs on the network, which caters for router redundancy protocols, such as HSRP and VRRP.

If exact MAC address matching is not enabled, the eTSEC determines whether DA is a group or individual address. If DA is the standard broadcast address, and broadcast addresses are not rejected, the frame is accepted. If any other group address is received, the eTSEC looks-up the DA by means of the group hash table. The group hash table may be extended to 512 entries if RCTRL[GHTX] = 1. Otherwise, an individual address is hashed into the 256-entry individual hash table when RCTRL[GHTX] = 0.

19.6.2.7.2 Hash Table Algorithm

The hash table process used in the group hash filtering operates as follows. By default, the Ethernet controller maps any 48-bit destination address into one of 256 bins, represented by the 256 bits in IGADDR0–IGADDR7 for individual addresses, and the 256 bits in GADDR0–GADDR7 for group addresses. But in the case where RCTRL[GHTX] is set, both sets of registers are combined into an extended group-only hash table of 512 bits, where IGADDR0–IGADDR7 contain the first 256 bits and GADDR0–GADDR7 contain the last 256 bits. No individual-address table exists in extended mode.

The 48-bit destination address received by the MAC is passed through the Ethernet CRC-32 algorithm to produce a hash value. The CRC polynomial used is:

$$x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

The MAC initializes its CRC register to 0xFFFFFFFF before computing a CRC on the 6 bit-reversed octets of the DA. A non-optimized sample of C code for computing the DA hash is listed in [Figure 19-136](#). The 9 most significant bits of the raw, uninverted CRC are used as the hash table index, H[8:0]. If RCTRL[GHTX] = 0, bits H[8:6] select one of the 8 IGADDR or GADDR registers, while bits H[5:1] select a bit within the 32-bit register. If RCTRL[GHTX] = 1, bits H[8:5] select one of the 16 registers in the {IGADDR, GADDR} set, while bits H[4:0] select a bit within the 32-bit register. For example, if H[8:5] = 7, IGADDR7 is selected, whereas H[8:5] = 9 selects GADDR1.

```

/* Wrapper macros for 256-bucket and 512-bucket hash tables:
   Pass 6-byte Ethernet MAC address as parameter. */
#define TSEC_HASH256(macaddr) ((crc32(macaddr) >> 24) & 0xff)
#define TSEC_HASH512(macaddr) ((crc32(macaddr) >> 23) & 0x1ff)

/* CRC constants. Note: CRC-32 polynomial is bit-reversed. */
#define CRC_POLYNOMIAL 0xedb88320
#define CRC_INITIAL    0xffffffff
#define MAC_ADDRLEN    6
#define BITS_PER_BYTE  8

/* crc32() Takes the array of bytes, macaddr[], representing an
   Ethernet MAC address and returns the CRC-32 result over these bytes,
   where each byte is used in bit-reversed form (Ethernet bit order).
   Index 0 of macaddr[] is the first byte of the address on the wire.
   Test case: the result of crc32 on {0x00, 0x01, 0x02, 0x03, 0x04, 0x05}
   should be 0xad0c28f3.
*/
unsigned long crc32(unsigned char macaddr[MAC_ADDRLEN])
{
    unsigned long crc, result;
    int byte, i;

    /* CRC-32 algorithm starts by inverting first 4 bytes */
    crc = CRC_INITIAL;
    /* add each byte to running CRC accumulator */
    for (byte = 0; byte < MAC_ADDRLEN; ++byte) {
        crc ^= macaddr[byte];
        /* shift CRC right to perform but reversal on byte of address */
        for (i = 0; i < BITS_PER_BYTE; ++i)
            if (crc & 1)
                crc = (crc >> 1) ^ CRC_POLYNOMIAL;
            else
                crc >>= 1;
    }
    /* finally, reverse bits of result to get CRC in normal bit order */
    for (result = 0, i = 4*BITS_PER_BYTE-1; i >= 0; crc >>= 1, --i)
        result |= (crc & 1) << i;
    return result;
}

```

Figure 19-136. Sample C Code for Computing eTSEC Hash Table Indices

If the CRC hash table index selects a bit that is set in the hash table, the frame is accepted. If 32 group addresses are stored in the hash table and random group addresses are received, the extended hash table prevents roughly 480/512 (93.8%) of the group address frames from reaching memory. Software must further filter those that reach memory to determine if they contain the correct addresses. Alternatively, small multicast groups can be held in the exact match MAC address registers, which guarantees that only correct frames are admitted.

The effectiveness of the hash table declines as the number of addresses increases. For instance, as the number of addresses stored in the 512-bin hash table increases, the vast majority of the hash table bits are set, preventing only a small fraction of frames from reaching memory.

NOTE

The hash table cannot be used to reject frames that match a set of selected addresses because unintended addresses can map to the same bit in the hash table. The receive queue filer may be used to reject frames with unintended address hits in the hash table.

19.6.2.8 Magic Packet Mode

eTSEC implements the AMD Magic Packet specification for LAN-initiated power management. This mode is normally entered with the rest of the system in a low-power sleep mode. Software must enable normal receive function in the Ethernet MAC, and then finally set the MACCFG2[MPEN] bit to enable Magic Packet detection before the system enters a reduced mode. While the rest of the system is operating in low-power mode, the enabled eTSEC continues to receive Ethernet frames, but discards them immediately. Upon receipt of any frame whose contents contain the valid Magic Packet sequence, the eTSEC exits out of Magic Packet mode, thus clearing MACCFG2[MPEN], and raises an error/diagnostic interrupt through IEVENT[MAG], which causes the surrounding system to wake-up. Frames received after Magic Packet mode has exited are received into software buffers as usual. Software can abort Magic Packet mode by writing 0 to MACCFG2[MPEN] at any time.

AMD specify a Magic Packet to be any Ethernet frame containing a valid Ethernet header (Destination and Source Addresses) and valid FCS (CRC-32), and whose payload includes the specific Magic Packet byte sequence at any offset from the start of frame. The specific byte sequence comprises an unbroken stream of 102 bytes, the first 6 bytes of which are 0xFFFFFFFF_FFFFFFFF, followed by 16 copies of the MAC's unique IEEE station address in the normal byte order for Ethernet addresses. For example, if the station address were 0x112233_445566, then the MAC would have to receive 0xFFFFFFFF_FFFFFFFF, 0x112233_445566, ..., 0x112233_445566 in any payload to detect a Magic Packet. Only frames addressed specifically to the MAC's station address or a valid multicast or broadcast address can be examined for the Magic Packet sequence.

19.6.2.9 Flow Control

Because collisions cannot occur in full-duplex mode, gigabit Ethernet can operate at the maximum rate. If the rate becomes too fast for a station's receiver, the station's transmitter can send flow-control frames to reduce the rate. Flow-control instructions are transferred by special frames of minimum frame size. The length/type fields of these frames have a special value.

Table 19-141 lists the flow-control frame structure.

Table 19-141. Flow Control Frame Structure

| Size [Octets] | Description | Value | Comment |
|---------------|---------------------|-------------------|---|
| 7 | Preamble | | — |
| 1 | SFD | | Start frame delimiter |
| 6 | Destination address | 01-80-C2-00-00-01 | Multicast address reserved for use in MAC frames (or MAC station address) |
| 6 | Source address | | — |
| 2 | Length/type | 88-08 | Control frame type |

Table 19-141. Flow Control Frame Structure (continued)

| Size [Octets] | Description | Value | Comment |
|---------------|------------------------|-------|--|
| 2 | MAC opcode | 00-01 | Pause command |
| 2 | MAC parameter | | Pause time as defined by the PTV[PT] field. The pause period is measured in pause_quanta, a speed independent constant of 512 bit-times (unlike slot time). The most-significant octet is transmitted first. |
| 2 | Extended MAC parameter | | Pause time extended as defined by the PTV[PTE] field. The most significant octet is transmitted first. |
| 40 | Reserved | — | — |
| 4 | FCS | | Frame check sequence (CRC) |

If flow-control mode is enabled (MACCFG1[Rx_Flow] is set) and the receiver identifies a pause-flow control frame, transmission stops for the time specified in the control frame. The controller completes any frame in progress before stopping transmission and does not commence counting the pause time until transmit is idle. During a pause, only a control frame can be sent (TCTRL[TFC_PAUSE] is set). Normal transmission resumes after the pause timer stops counting, or resumes immediately if a pause frame with a zero time-out is received. If another pause-control frame is received during the pause, the period changes to the new value received.

19.6.2.10 Interrupt Handling

The following describes what usually occurs within a eTSEC interrupt handler:

- If an interrupt occurs, read IEVENT to determine interrupt sources. IEVENT bits to be handled in this interrupt handler are normally cleared at this time. There are three kinds of interrupts:
 - Receive data frame interrupts, when bits RXB or RXF in IEVENT are set
 - Transmit data frame interrupts, when bits TXB or TXF in IEVENT are set
 - Error, diagnostic, and special interrupts (all bits in IEVENT other than RXB, RXF, TXB, or TXF)
- Process the TxBDs to reuse them if the IEVENT[TXB, TXF or TXE] were set. Consult register bits TSTAT[TXF0–TXF7] to determine which TxBD rings gave rise to the transmit interrupt in the case of TXF. If the transmit speed is fast or the interrupt delay is long, more than one transmit buffer may have been sent by the eTSEC; thus, it is important to check more than just one TxBD during the interrupt handler. One common practice is to process all TxBDs in the interrupt handler until one is found with R set.
- Obtain data from RxBD rings if IEVENT[RXC, RXB or RXF] is set. Consult register bits RSTAT[RXF0–RXF7] to determine which RxBD rings gave rise to the receive interrupt in the case of RXF. If the receive speed is fast or the interrupt delay is long, the eTSEC may have received more than one RxBD; thus, it is important to check more than just one RxBD during interrupt handling. Typically, all RxBDs in the interrupt handler are processed until one is found with E set. Because the eTSEC pre-fetches BDs, the BD table must be big enough so that there is always another empty BD to pre-fetch, otherwise a BSY error occurs.

- Clear any set halt or frame interrupt bits in TSTAT and RSTAT registers, or DMACTRL[GTS] and DMACTRL[GRS] by writing 1s to these bits.
- Continue normal execution.

Table 19-142. Non-Error Transmit Interrupts

| Interrupt | Description | Action Taken by the eTSEC |
|-----------|---|---|
| GTSC | Graceful transmit stop complete: transmitter is put into a pause state after completion of the frame currently being transmitted. | None |
| TXC | Transmit control: Instead of the next transmit frame, a control frame was sent. | None |
| TXB | Transmit buffer: A transmit buffer descriptor, that is not the last one in the frame, was updated in one of the enabled TxBD rings. | Programmable 'write with response' TxBD to memory before setting IEVENT[TXB]. |
| TXF | Transmit frame: A frame from an enabled TxBD ring was transmitted and the last transmit buffer descriptor (TxBD) of that frame was updated. | Programmable 'write with response' to memory on the last TxBD before setting IEVENT[TXF]. |

Table 19-143. Non-Error Receive Interrupts

| Interrupt | Description | Action Taken by the eTSEC |
|-----------|---|---|
| GRSC | Graceful receive stop complete: Receiver is put into a pause state after completion of the frame currently being received. | None |
| RXC | Receive control: A control frame was received. As soon as the transmitter finishes sending the current frame, a pause operation is performed. | None |
| RXB | Receive buffer: A receive buffer descriptor, that is not the last one of the frame, was updated in one of the enabled RxBD rings. | Programmable 'write with response' RxBD to memory before setting IEVENT[RXB]. |
| RXF | Receive frame: A frame was received to an enabled RxBD ring and the last receive buffer descriptor (RxBD) of that frame was updated. | Programmable 'write with response' to memory on the last RxBD before setting IEVENT[RXF]. |

19.6.2.10.1 Interrupt Coalescing

Interrupt coalescing offers the user the ability to contour the behavior of the eTSEC with regard to frame interrupts. Separate but identical mechanisms exist for both transmitted frames and received frames. In either case, frame interrupts require that software set the I-bit in RxBDs or TxBDs, and disable buffer interrupts (IEVENT[RXB] or IEVENT[TXB]). Particular rings can remain free of interrupts by ensuring that the I-bit is consistently cleared in all BDs. While interrupt coalescing is enabled, a transmit or receive frame interrupt is raised either when a counter threshold-defined number of frames is received/transmitted or the timer threshold-defined period of time has elapsed, whichever occurs first. Disabling and then re-enabling interrupt coalescing forces reset of the coalescing timers and counters to reflect changes made to the threshold registers.

19.6.2.10.2 Interrupt Coalescing By Frame Count Threshold

To avoid interrupt bandwidth congestion due to frequent, consecutive interrupts, the user may enable and configure interrupt coalescing to deliberately group frame interrupts, reducing the total number of

interrupts raised. The number of frames received or transmitted prior to an interrupt being raised is determined by the frame threshold field (ICFT) in the appropriate interrupt coalescing configuration register (RXIC or TXIC). The frame threshold field may be assigned a value between 1 and 255. The internal transmit or receive frame counter decrements from this initial value each time a frame is transmitted or received. Upon reaching zero, an interrupt is raised, the appropriate threshold counter is reset to the value in the ICFT field, and then eTSEC continues counting frames while the interrupt is active. The appropriate threshold counter is also reset to the value in the ICFT field if an interrupt is raised subject to the corresponding threshold timer.

19.6.2.10.3 Interrupt Coalescing By Timer Threshold

To avoid stale frame interrupts, the user may also assign a timer threshold, beyond which any frame interrupts not yet raised are forced. The timer threshold fields of the receive and transmit interrupt coalescing configuration registers (RXIC[ICTT] and TXIC[ICTT]) are defined in units equivalent to 64 interface clocks or system clocks, depending on the setting of the ICCS field in RXIC and TXIC.

After transmitting a frame, the transmit interrupt coalescing threshold time begins counting down from the value in TXIC[ICTT]. An interrupt is raised when the counter reaches zero. In the event of graceful transmit stop completion before the coalescing timer expires, the eTSEC issues two interrupts, the first for GTS, the second for TXF (due to timer expiration of a pending event). To prevent the second interrupt from affecting servicing of the GTS event, it is recommended that the user mask out the TXF event during execution of the service routine. After receiving a frame, the receive interrupt coalescing threshold time begins counting down from the value in RXIC[ICTT]. An interrupt is raised when the counter reaches zero. In the event of graceful receive stop completion before the coalescing timer expires, the eTSEC issues two interrupts, the first for GRS, the second for RXF (due to timer expiration of a pending event). To prevent the second interrupt from affecting servicing of the GRS event, it is recommended that the user mask out the RXF event during execution of the service routine.

The interrupt coalescing timer thresholds (transmit and receive, operating independently) may be values ranging from 0x0001 to 0xFFFF. [Table 19-144](#) specifies the range of possible timing thresholds subject to timer clock source, the interface or system frequency, and the value of the RXIC[ICTT] or TXIC[ICTT] field.

Table 19-144. Interrupt Coalescing Timing Threshold Ranges

| ICCS (Clock Source) | eTSEC Interface Format and Frequency or eTSEC System Frequency | Interrupt Coalescing Threshold Time | |
|------------------------|--|-------------------------------------|-------------------------|
| | | Minimum (ICTT = 0x0001) | Maximum (ICTT = 0xFFFF) |
| 0 (I/F clock) | 10Base-T at 2.5 MHz | 25.6 μ s | 1.68 s |
| 0 (I/F clock) | 100Base-T at 25 MHz | 2.56 μ s | 168 ms |
| 0 (I/F clock) | 1000Base-T at 125 MHz | 0.51 μ s | 33.6 ms |
| 1 (sys. clock) | eTSEC operating at 266 MHz | 0.24 μ s | 15.7 ms |
| 1 (sys. clock) | eTSEC operating at 333 MHz | 0.19 μ s | 12.6 ms |

The transmit timer threshold counter is reset to the value in TXIC[ICTT] and begins counting down on transmission of the frame following an interrupt.

The receive timer threshold counter is reset to the value in RXIC[ICTT] and begins counting down on receiving the frame following an interrupt.

19.6.2.11 Inter-Frame Gap Time

If a station must transmit, it waits until the LAN becomes silent for a specified period (inter-frame gap, or IFG). The minimum inter-packet gap (IPG) time for back-to-back transmission is set by IPGIFG[Back-to-Back Inter-Packet-Gap]. The receiver receives back-to-back frames with the minimum interframe gap (IFG) as set in IPGIFG[Minimum IFG Enforcement]. If multiple frames are ready to transmit, the Ethernet controller follows the minimum IPG as long as the following restrictions are met:

- The first TxBD pointer, TBPTR n , of any given frame is located at a 16-byte aligned address.
- All BDs for any multiple-BD frame reside in the same cache line.
- TCP/UDP and IP Checksum generation are disabled in each frame's TxFCB, or in TCTRL, or frames are limited to 1200 bytes in length.
- Each TxBD[Data Length] is greater-than or equal to 64 bytes.

If the first TxBD alignment restriction is not met, the back-to-back IPG may be as many as 32 cycles. If the TxBD size restriction is not met, the back-to-back IPG may be significantly longer.

In half-duplex mode, after a station begins sending, it continually checks for collisions on the LAN. If a collision is detected, the station forces a jam signal (all ones) on its frame and stops transmitting. Collisions usually occur close to the beginning of a packet. The station then waits a random time period (back-off) before attempting to send again. After the back-off completes, the station waits for silence on the LAN (carrier sense negated) and then begins retransmission (retry) on the LAN. Retransmission begins 36 bit times after carrier sense is negated for at least 60 bit times. If the frame is not successfully sent within a specified number of retries, an error is indicated (collision retry limit exceeded).

19.6.2.12 Internal and External Loop Back

Setting MACCFG1[Loop Back] causes the MAC transmit outputs to be looped back to the MAC receive inputs. Clearing this bit results in normal operation. This bit is cleared by default. Clearing this bit results in normal operation.

19.6.2.13 Error-Handling Procedure

The eTSEC reports frame reception and transmission error conditions using the channel BDs, the error counters, and the IEVENT register.

Transmission errors are described in [Table 19-145](#).

Table 19-145. Transmission Errors

| Error | Response |
|----------------------|--|
| Transmitter underrun | Transmitter underrun can occur either after frame transmission has commenced, or in response to an incomplete sequence of TxBDs. In the former case, the controller sends 32 bits that ensure a CRC error, and terminates buffer transmission. In the latter case, the relevant transmit queue is halted. In all cases, the eTSEC closes the buffer, sets TxBD[UN], IEVENT[XFUN], and IEVENT[TXE]. The controller resumes transmission after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared). |

Table 19-145. Transmission Errors (continued)

| Error | Response |
|---------------------------------------|---|
| Retransmission attempts limit expired | The controller terminates buffer transmission, sets TxB[RL], closes the buffer, IEVENT[CRL], and IEVENT[TXE]. Transmission resumes after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared). |
| Late collision | The controller terminates buffer transmission, sets TxB[LC], closes the buffer, IEVENT[LC], and IEVENT[TXE]. The controller resumes transmission after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared). |
| Memory read error | A system bus error occurred during a DMA transaction. The controller sets IEVENT[EBERR], DMA stops sending data to the FIFO which causes an underrun error, and therefore TxB[UN] is set, but IEVENT[XFUN] is not set. The TSTAT[THLT] is set. Transmits are continued once TSTAT[THLT] is cleared. |
| Data parity error | Data in the transmit FIFO was potentially corrupted. The controller sets IEVENT[DPE], but otherwise continues transmission until halted explicitly. |
| Babbling transmit error | A frame is transmitted which exceeds the MAC's Maximum Frame Length and MACCFG2[Huge Frame] is a 0. The controller sets IEVENT[BABT] and continues without interruption. |

Reception errors are described in [Table 19-146](#).

Table 19-146. Reception Errors

| Error | Description |
|------------------------------------|---|
| Overrun error | The Ethernet controller maintains an internal FIFO buffer for receiving data. If a receiver FIFO buffer overrun occurs, the controller sets RxB[OV], sets RxB[L], closes the buffer, increments the discarded frame counter (RDRP), and sets IEVENT[RXF]. The receiver then enters hunt mode (seeking start of a new frame). |
| Busy error | A frame is received and discarded due to a lack of buffers. The controller sets IEVENT[BSY] and increments the discarded frame counter (RDRP). In addition, the RSTAT[QHLT n] bit is set. RDRP increments for each frame that is received while the receiver is halted due to a busy condition. The halted queue resumes reception once the RSTAT[QHLT n] bit is cleared. |
| Filed frame to invalid queue error | A frame is received and discarded as a result of the filer directing it to an RxB ring that is currently not enabled. The controller sets IEVENT[FIQ] and increments the discarded frame counter (RDRP). |
| Parser error | If the receive frame parser is enabled, a parse error can be flagged as a result of inconsistencies discovered between fields of the embedded packet headers. For example, the L2 header may indicate an IPv4 header, but the IP version number fails to match. In the event of a parse error, parsing is terminated at the inconsistent header, and the RxFCB[PERR] field indicates at which layer of the protocol stack the error was discovered. Receiver function continues regardless of parse errors, but IEVENT[PERR] is set. The receive queue filer may operate with reduced or default information in some cases; therefore, filer rule sets should be constructed so as to be tolerant of malformed frames. Note: Any values in the length/type field between 1500 and 1536 is treated as a length, however, only illegal packets exist with this length/type since these are not valid lengths and not valid types. These are treated by the MAC logic as out of range. Software must confirm the parser and filer results by checking the type/length field after the packet has been written to memory to see if it falls in this range. |

Table 19-146. Reception Errors (continued)

| Error | Description |
|----------------------------------|--|
| Non-octet error (dribbling bits) | The Ethernet controller handles a nibble of dribbling bits if the receive frame terminates as non-octet aligned and it checks the CRC of the frame on the last octet boundary. If there is a CRC error, the frame non-octet aligned (RxB[NO]) error is reported, IEVENT[RXF] is set, and the alignment error counter increments. The eTSEC relies on the statistics collector block to increment the receive alignment error counter (RALN). If there is no CRC error, no error is reported. |
| CRC error | If a CRC error occurs, the controller sets RxB[CR], closes the buffer, and sets IEVENT[RXF]. This eTSEC relies on the statistics collector block to record the event. After receiving a frame with a CRC error, the receiver then enters hunt mode. |
| Memory read error | A system bus error occurred during a DMA transaction. The controller sets IEVENT[EBERR] and discards the frame and increments the discarded frame counter (RDRP). In addition the RSTAT[QHLT n] bit is set. The halted queue resumes reception once the RSTAT[QHLT n] bit is cleared. |
| Data parity error | Data in the receive FIFO or filter table was potentially corrupted. The controller sets IEVENT[DPE], but otherwise continues reception until halted explicitly. |
| Babbling receive error | A frame is received that exceeds the MAC's maximum frame length. The controller sets IEVENT[BABR] and continues. |

19.6.3 TCP/IP Off-Load

Each eTSEC provides hardware support for accelerating the basic functions of TCP/IP packet transmission and reception. By default, these features are disabled and must be explicitly enabled through RCTRL and TCTRL. In this configuration, the eTSEC processes frames as vanilla Ethernet frames and none of the multi-ring QoS/CoS receive services or per-frame VLAN insertion and deletion are available. Operate eTSEC in this default configuration when using existing TCP/IP stack software that has not been modified to take advantage of TOE.

TOE can be enabled independently for Rx and Tx and at various levels. Receive TOE functions are controlled by RCTRL and transmit functions through a combination of TCTRL[TUCSEN] and the Tx frame control block.

On receive, according to RCTRL[PRSDPE], eTSEC can parse frames at layer 2 of the stack only (Ethernet headers and switching headers), layers 2 to 3 (including IPv4 or IPv6), or layers 2 to 4 (including TCP and UDP). TOE provides protocol header recognition, header verification (IPv4 header checksum verification), and TCP/UDP payload checksum verification including verification of associated pseudo-header checksums. For large frames off-load of checksum verification saves a significant fraction of the CPU cycles that would otherwise be spent by the TCP/IP stack. IP packet fragmentation and re-assembly, and TCP stream establishment and tear-down are not performed in hardware. The frame parser sets RQFPR[IPF] status flag encountering a fragmented frame. The frame parser in eTSEC searches a maximum of 512 bytes from the start of a received frame when attempting to locate headers; headers deeper than 512 bytes are assumed not to exist, and any associated receive status flags in the frame control block remain cleared.

On transmit, TOE provides IPv4 and TCP/UDP header checksum generation. Like receive TOE, checksum generation reduces CPU load significantly for TCP/IP stacks modified to exploit eTSEC TOE functions. The eTSEC does not checksum transmitted packets with IPv6 routing headers or calculate TCP/UDP checksums from IP fragments. If a transmitted TCP segment requires checksum generation but

IPv6 extension headers would prevent eTSEC from calculating the pseudo-header checksum, software can calculate just the pseudo-header checksum in advance and supply it to the eTSEC as part of per-frame TOE configuration.

19.6.3.1 Frame Control Blocks

Frame control blocks (FCBs) are 8-byte blocks of TOE control and/or status data that are passed between software (driver and TCP/IP stack) and each eTSEC. A FCB always precedes the frame it applies to, and is present only when TOE functions are being used. As Figure 19-137 shows, the first BD of each frame points to the initial data buffer and the FCB. The initial data buffer must be at least 8 bytes long to contain the FCB without breaking it. Custom or received Ethernet preamble sequences also follow the FCB if preambles are visible.




Figure 19-137. Location of Frame Control Blocks for TOE Parameters

For TxBD rings, FCBs are assumed present when the TxBD[TOE/UN] bit is set by user software. The eTSEC ignores the TxBD[TOE/UN] bit in all BDs other than those pointing to initial data buffers, therefore FCBs must not be inserted in second and subsequent data buffers. Since TxBD[TOE/UN] can be set under software discretion, TOE acceleration for transmit may be applied on a frame-by-frame basis.

In the case of RxBD rings, FCBs are inserted by the eTSEC whenever RCTRL[PRSDEP] is set to a non-zero value. Only one FCB is inserted per frame, in the buffer pointed to by the RxBD with bit F set. TOE acceleration for receive is enabled for all frames in this case.

19.6.3.2 Transmit Path Off-Load and Tx PTP Packet Parsing

TOE functions for transmit are defined by the contents of the Tx FCB. [Figure 19-138](#) describes the definition for the Tx FCB.

| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|------------|-------|----|-----|-----|-----|-----|-----|------|---|---|----|----|----|----|----|-----|
| Offset + 0 | VLN | IP | IP6 | TUP | UDP | CIP | CTU | NPH | | | | | | | | PTP |
| Offset + 2 | L4OS | | | | | | | L3OS | | | | | | | | |
| Offset + 4 | PHCS | | | | | | | | | | | | | | | |
| Offset + 6 | VLCTL | | | | | | | | | | | | | | | |

Figure 19-138. Transmit Frame Control Block

The user instructs the Tx packet to be timestamped via setting bit 15 in the TxFCB to mark a PTP packet. TxFCB[VLCTL] can be translated as the Tx PTP packet identification number. BD[TOE] has to be set to enable transmit PTP packet time stamping. TxFCB[PTP] bit takes precedence over TxFCB[VLN] bit. It disables per packet VLAN tag insertion. On a PTP packet, VLAN tag can be inserted from the DFVLAN register. A proposed TxFCB update for the PTP packet is shown in [Figure 19-145](#).

The contents of the Tx FCB are defined in [Table 19-147](#).

Table 19-147. Tx Frame Control Block Description

| Bytes | Bits | Name | Description |
|-------|------|------|---|
| 0–1 | 0 | VLN | VLAN control word valid. This bit is ignored when the PTP bit is set. VLAN tag is read from the DFVLAN register if PTP=1. 0 Ignore VLCTL field. 1 If VLAN tag insertion is enabled for eTSEC, use the VLCTL field as the VLAN control word. |
| | 1 | IP | Layer 3 header is an IP header. 0 Ignore layer 3 and higher headers. 1 Assume that the layer 3 header is an IPv4 or IPv6 header, and take L3OS field as valid. |
| | 2 | IP6 | IP header is IP version 6. Valid only if IP = 1. 0 IP header version is 4. 1 IP header version is 6. |
| | 3 | TUP | Layer 4 header is a TCP or UDP header. 0 Do not process any layer 4 header. 1 Assume that the layer 4 header is either TCP or UDP (see UDP bit), and offload checksumming on the basis that the IP header has no extension headers. |
| | 4 | UDP | UDP protocol at layer 4. 0 Layer 4 protocol is either TCP (if TUP = 1) or undefined. 1 Layer 4 protocol is UDP if TUP = 1. |

Table 19-147. Tx Frame Control Block Description (continued)

| Bytes | Bits | Name | Description |
|-------|------|------------------|---|
| 0–1 | 5 | CIP | Checksum IP header enable. 0 Do not generate an IP header checksum. 1 Generate an IPv4 header checksum. |
| | 6 | CTU | Checksum TCP or UDP header enable. 0 Do not generate a TCP or UDP header checksum. RFC 768 advises that UDP packets not requiring checksum validation should have their checksum field set to zero. 1 Generate a TCP header checksum if IP = 1 and TUP = 1 and UDP = 0. |
| | 7 | NPH | Disable calculation of TCP or UDP pseudo-header checksum. This bit should be set if IP options need to be consulted in forming the pseudo-header checksum, as eTSEC does not examine IP options or extension headers for TCP/IP offload on transmit. 0 Calculate TCP or UDP pseudo-header checksum as normal, assuming that the IP header has no options. 1 Do not calculate a TCP or UDP pseudo-header checksum, but instead use the value in field PHCS when determining the overall TCP or UDP checksum. |
| | 8–14 | — | Reserved |
| | 15 | PTP | Indication to the transmitter that this is a PTP packet. Enabling PTP disables per packet VLAN tag insertion. Instead, VLAN tag will be read from the DFVLAN when the PTP field is true. 0 Do not attempt to capture transmission event time 1 Valid PTP_ID field. When this packet is transmitted, capture the time of transmission. Must be clear if TMR_CTRL[TE] is clear. |
| 2–3 | 0–7 | L4OS | Layer 4 header offset from start of layer 3 header. The layer 4 header starts L4OS octets after the layer 3 header if it is present. The maximum layer 3 header length supported is thus 255 bytes, which may prevent TCP/IP offload on particularly large IPv6 headers. |
| | 8–15 | L3OS | Layer 3 header offset from start of frame not including the 8 bytes for this FCB. The layer 3 header starts L3OS octets from the start of the frame including any custom preamble header that may be present. The maximum layer 2 header length supported is thus 255 bytes. |
| 4–5 | 0–15 | PHCS | Pseudo-header checksum (16-bit one's complement sum with carry wraparound, but without result inversion) for TCP or UDP packets, calculated by software. Valid only if NPH = 1. |
| 6–7 | 0–15 | VLCTL/ PTP_ID | VLAN control word for insertion in the transmitted VLAN tag. Valid only if VLN = 1. Tx PTP packet identification number. This number will be copied into the Tx PTP packet time stamp identification field. PTP field takes precedence over VLN field. |

19.6.3.3 Receive Path Off-Load

Upon receive, the Rx FCB returns the status of frame parse and TOE functions applied to the accompanying frame. [Figure 19-139](#) describes the definition for the Rx FCB.

| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|------------|-------|----|-----|-----|-----|-----|-----|-----|-----|---|----|------|----|----|------|----|
| Offset + 0 | VLN | IP | IP6 | TUP | CIP | CTU | EIP | ETU | — | | | PERR | | — | GPFP | |
| Offset + 2 | | | | | RQ | | | | PRO | | | | | | | |
| Offset + 4 | | | | | | | | | | | | | | | | |
| Offset + 6 | VLCTL | | | | | | | | | | | | | | | |

Figure 19-139. Receive Frame Control Block

The contents of the Rx FCB are defined in [Table 19-148](#).

Table 19-148. Rx Frame Control Block Descriptions

| Bytes | Bits | Name | Description |
|-------|-------|------|--|
| 0–1 | 0 | VLN | VLAN tag recognized. This bit is set only if RCTRL[VLEX] is set. 0 No VLAN tag recognized. 1 IEEE Std. 802.1Q VLAN tag found; VLAN control word in VLCTL is valid. |
| | 1 | IP | IP header found at layer 3. RCTRL[PRSDEP] must be set to 10 or 11 in order to enable IP discovery. See also IP6 bit of FCB. 0 No layer 3 header recognized. 1 An IP header was recognized at layer 3; the IANA protocol identifier for the next header can be found in PRO; see PRO for more information. If S/W is relying on the RxFCB for the parse results, any RxFCB[IP] bits set with the corresponding RxFCB[PRO] = 0xFF indicates a fragmented packet (or that this packet had a back-to-back IPv6 routing extension header). Additionally, RQFPR[IPF] (see Section 19.5.3.3.8, “Receive Queue Filer Table Property Register (RQFPR)”) indicates that the packet was fragmented. |
| | 2 | IP6 | IP version 6 header found at layer 3. 0 No IPv6 header was found. 1 The layer 3 header was an IPv6 header provided IP = 1. |
| | 3 | TUP | TCP or UDP header found at layer 4. RCTRL[PRSDEP] must be set to 10 or 11 in order to enable TCP/UDP discovery. 0 No layer 4 header recognized. 1 The layer 4 header was recognized as either TCP (PRO = 0x06) or UDP (PRO = 0x11). |
| | 4 | CIP | IPv4 header checksum checked. RCTRL[PRSDEP] must be set to 10 or 11 in order to enable IPv4 checksum verification. 0 IPv4 header checksum not verified, either because verification was disabled or a valid IPv4 header could not be located. 1 IPv4 header checksum was verified by the eTSEC, and bit EIP indicates result. |
| | 5 | CTU | TCP or UDP header checksum checked. RCTRL[PRSDEP] must be set to 11 in order to enable layer 4 checksum verification. 0 TCP or UDP header checksum not verified, either because verification was disabled or a valid TCP or UDP header could not be located. If a UDP header with zero checksum was located, this bit is cleared in accordance with RFC 768. 1 TCP or UDP header checksum was verified by the eTSEC, and ETU indicates result. |
| | 6 | EIP | IPv4 header checksum verification error. Not valid unless CIP = 1. 0 No checksum error in IPv4 header. 1 Error in header checksum only if IP = 1 and IP6 = 0. |
| 0–1 | 7 | ETU | TCP or UDP header checksum verification error. Not valid unless CTU = 1. 0 No checksum error in TCP or UDP header. 1 Error in header checksum only if PRO = 0x06 or PRO = 0x11. |
| | 8–11 | — | Reserved |
| | 12–13 | PER | Parse error. 00 No error in L2 to L4 parse 01 Reserved 10 Inconsistent or unsupported L3 header sequence 11 Reserved |
| | 14 | — | Reserved |
| | 15 | GFPF | General-purpose filer event packet. This packet was filed based on matching a GPI rule sequence. |

Table 19-148. Rx Frame Control Block Descriptions (continued)

| Bytes | Bits | Name | Description |
|-------|------|-------|--|
| 2–3 | 0–1 | — | Reserved |
| | 2–7 | RQ | Receive queue index. This index was selected by the eTSEC Rx Filer (from a matching Filer rule's RQCTRL[Q] field) when it accepted the associated frame. If filing is not enabled, RQ is zero. Note that the 3 least significant bits of RQ correspond with the RxBD ring index whenever RCTRL[FSQEN] = 0. |
| | 8–15 | PRO | <p>If IP = 1, PRO is set as follows:</p> <ul style="list-style-type: none"> • PRO=0xFF for a fragment header or a back to back route header • PRO=0xnn for an unrecognized header, where nn is the next protocol field • PRO=(TCP/UDP header), as defined in the IANA specification, if TCP or UDP header is found <p>If IP = 0, PRO is undefined.</p> <p>Note that the eTSEC parser logic stops further parsing when encountering an IP datagram that has indicated that it has fragmented the upper layer protocol. This in general means that there is likely no layer 4 header following the IP header and extension headers. eTSEC leaves the RxFCB[PRO] and RQFPR[L4P] fields 0xFF in this case, which usually means that there was no IP header seen. In this case RxFCB[IP] and optionally RxFCB[IP6] is set. IP header checksumming operates and performs as intended. Most of the time, the eTSEC updates the RxFCB[PRO] field and RQFPR[L4P] fields with whatever value was found in the protocol field of the IP header. See Section 19.5.3.3.8, “Receive Queue Filer Table Property Register (RQFPR),” for a description of RQFPR.</p> |
| 4–5 | 0–15 | — | Reserved |
| 6–7 | 0–15 | VLCTL | VLAN control word as per IEEE Std. 802.1Q. The lower 12 bits comprise the VLAN identifier. Valid only if VLN = 1. |

19.6.4 Quality of Service (QoS) Provision

This section describes the quality of service support features of this device. It includes a parser which extracts vital packet properties and passes them to the filer which essentially acts as a frame classifier.

19.6.4.1 Receive Parser

The receive parser parses the incoming frame data and generates filer properties and frame control block (FCB). The receive parser composes of the Ethernet header parser and L3/L4 parser.

The Ethernet header parser parses only L2 (ethertype) headers. It is enabled by RCTRL[PRSDEP] != 0. It has the following key features:

- Extraction of 48-bit MAC destination and source addresses
- Extraction and recognition of the first 2-byte ethertype field
- Extraction and recognition of the final 2-byte ethertype field
- Extraction of 2-byte VLAN control field
- Walk through MPLS stack and find layer 3 protocol
- Walk through VLAN stack and find layer 3 protocol
- Recognition of the following ethertypes for inner layer parsing
 - LLC and SNAP header

- JUMBO and SNAP header
- IPV4
- IPV6
- VLAN
- MPLSU/MPLSM
- PPOES

For stack L2 (that is, more than one ethertypes) header, the Ethernet parser traverses through the header until it finds the last valid ethertype or the ethertype is unsupported. [Table 19-149](#) describes what the Ethernet header parser recognizes for stack L2 header.

Table 19-149. Supported Stack L2 Ethernet Headers

| Column—Current L2 Ethertype Row—Next Supported L2 Ethertype | LLC/ SNAP | JUMBO/ SNAP | IPV4 | IPV6 | VLAN | MPLSU | MPLSM | PPOES |
|--|--------------|----------------|------|------|------|-------|-------|-------|
| LLC/SNAP | N | N | Y | Y | Y | Y | Y | Y |
| JUMBO/SNAP | N | N | Y | Y | Y | Y | Y | Y |
| IPV4 | N | N | N | N | N | N | N | N |
| IPV6 | N | N | N | N | N | N | N | N |
| VLAN | Y | Y | Y | Y | Y | Y | Y | Y |
| MPLSU | N | N | Y* | Y* | N | y | Y | N |
| MPLSM | N | N | Y* | Y* | N | Y | Y | N |
| PPOES | N | N | Y | Y | N | Y | Y | N |

Note: * means that it is the next protocol

The L3 parser is enabled by `RCTRL[PRSDEP] = 10` or `11`. It begins when the Ethernet parser ends and a valid IPv4/v6 ethertype is found. The L4 header is enabled by `RCTRL[PRSDEP] = 11`. It begins when the L3 parser ends and a valid TCP/UDP next protocol is found and no fragment frame is found. The primary functionalities of L3(IPv4/6) and L4(TCP/UDP) parsers are as follows:

- IP recognition (v4/v6, encapsulated protocol)
- IP header checksum verification
- IPv4/6 over IPv4/6 (tunneling)—parse headers and find layer 4 protocol
- IP layer 4 protocol/next header extraction
- Stop parsing on unrecognized next header/protocol
- IPv4 support
 - IPv4 source and destination addresses

- 8-bit IPv4 type of service
- IP layer 4 protocol / next header support
 - IPV4
 - IPV4 Fragment. Parser stops after a fragment is found
 - TCP/UDP
- IPv6 support
 - The first 4 bytes of the IPv6 source address extraction
 - The first 4 bytes of the IPv6 destination address extraction
 - IPv6 source address hash for pseudo header calculation
 - IPv6 destination address hash for pseudo header calculation
 - 8-bit IPv6 traffic class field extraction
 - Payload length field extraction
 - IP layer 4 protocol/next header support
 - IPV6
 - IPV6 fragment. Parser stops after a fragment is found
 - IPV6 route
 - IPV6 hop/destination
 - TCP/UDP
- L4 (TCP/UDP) support
 - Extraction of 16-bit source port number extraction
 - Extraction of 16-bit destination port number extraction
 - TCP checksum calculation (including pseudo header)
 - UDP checksum calculation if the checksum field is not zero (including pseudo header)

19.6.4.2 Receive Queue Filer

The receive queue filer receives protocol header properties extracted from the incoming frame by the eTSEC frame parse engine. A property is defined to be a field extracted from a packet header, such as a TCP port number or VLAN identifier. As soon as the last identifiable header has been recognized, the filer commences searching the receive queue filer table, comparing properties in the table against properties extracted from the frame. This table is illustrated in [Figure 19-140](#). Software populates the table with property values, stored to the RQPROP field, and indicates how to match and interpret the properties by setting flags in the RQCTRL field. The eTSEC memory map provides access to these fields by way of an address register (RQFAR) and two porthole registers (RQFCR and RQFPR).

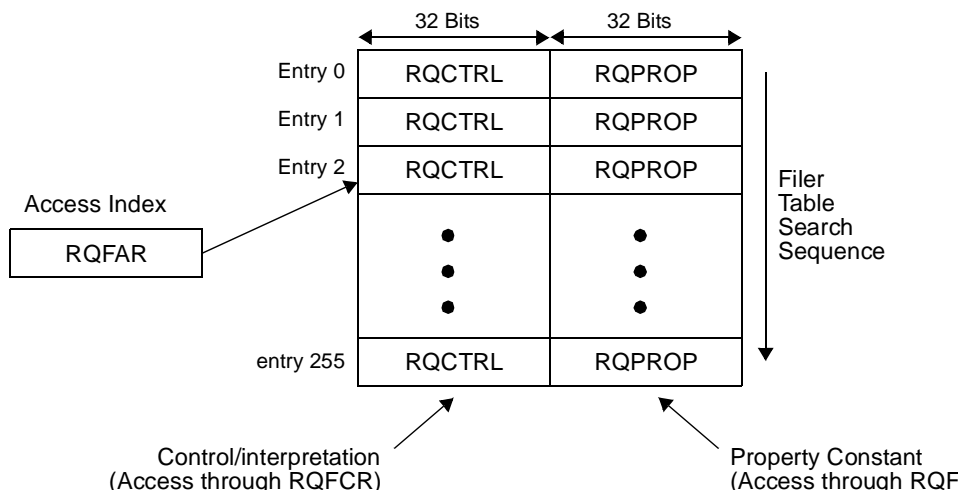


Figure 19-140. Structure of the Receive Queue Filer Table

19.6.4.2.1 Filing Rules

Unless the filer is disabled, every received frame from the Ethernet MAC initiates a search of the receive queue filer table, starting at entry 0. The table search is terminated as soon as an entry is found whose contents match a property of the frame. Accordingly, software must guarantee that at least one entry results in a match—even if only to set a default receive queue index.

Since eTSEC searches the table at a rate of two entries every system clock cycle, all 256 entries can be searched in the time taken to receive a 64-byte Ethernet frame. Note that in non-FIFO modes such as GMII running at slow platform frequency, MaxRules equation also applies. For example, using 20 for IPG, a minimum of 190.5 MHz eTSEC sysclk (381 MHz platform) is needed to process all 256 entries in GMII for a 64-byte receive frame.

Each entry of the receive queue filer table specifies a simple match rule for determining how to process the received frame. The elements of a filing rule, expressed in the RQCTRL and RQPROP fields, are summarized as follows:

- The PID field in RQCTRL identifies what property is being matched against RQPROP. The eTSEC supports 16 properties, some of which are different portions of the same header field. Reserved or unused bits in RQPROP are read as zero. See [Section 19.5.3.3.8, “Receive Queue Filer Table Property Register \(RQFPR\),” on page 19-57](#) for a list of all properties and their associated PID values.
- The Q field in RQCTRL identifies which one of 64 virtual receive queues the frame should be filed to (sent through DMA) in the event of a filing rule match that accepts the frame. The physical RxBD ring this queue maps to is controlled by the RCTRL[FSQEN] bit. If RCTRL[FSQEN] = 0, the three least significant bits of the Q field indicate which physical RxBD ring hosts the queue. If RCTRL[FSQEN] = 1, RxBD ring 0 hosts all receive queues, but the RxFCB[RQ] field allows software to distinguish queues by ID. In all cases if Q maps to a RxBD ring that is not currently enabled, the frame is discarded with an IEVENT[FIQ] error.

- The REJ field in RQCTRL controls whether the frame is to be rejected (REJ = 1) or filed (REJ = 0) upon a filing rule match. Rejected frames occupy Rx FIFO space, but do not consume memory bus cycles.
- The CMP field in RQCTRL determines how property PID is compared against RQPROP. Equality, inequality, greater-or-equal, and less-than compares are available.
- The AND field in RQCTRL allows more than one comparison in a sequence to be chained together as a Boolean AND condition. Setting AND = 1 defers evaluation of the rule until the next entry has been matched, which may, in turn, have AND set. If any comparison involving AND = 1 fails, the entire chained sequence fails. A typical use for AND is to combine a pair of comparisons in a range match; the first such entry has AND = 1, the second has AND = 0 and its values of Q and REJ take effect.
- The CLE field in RQCTRL offers a way to bracket a set of consecutive—perhaps related—rules into a rule cluster. A cluster must be preceded by a guard rule, which simply determines whether the cluster rules can be evaluated. If the guard rule succeeds and its last entry has both CLE = 1 and AND = 1, the cluster rules that follow are enabled. The cluster ends at the first entry where CLE = 1 and AND = 0, which may also belong to a rule that files or rejects a frame. If the guard rule fails, all rules in the cluster are skipped, including mask_register assignments. Clusters must not be nested.
- The GPI field offers the user the ability to interrupt the core upon matching a rule that causes a frame to be filed to memory. Once the last RxBD corresponding to that frame is written to memory, the IEVENT[FGPI] event will be asserted. This bit will be set regardless of any interrupt coalescing that may be set.

19.6.4.2.2 Comparing Properties with Bit Masks

By default, extracted properties are compared arithmetically according to the CMP field in each RQCTRL word. This permits point value matches in each table entry, and range checks across a pair of table entries combined with the AND attribute in RQCTRL. However, inspection of the parse flags, Ethernet preamble, and IP addresses typically requires “don’t care” bit fields in the properties to be cleared as part of the comparison. The eTSEC provides a dedicated 32-bit register, known as the mask_register, for performing such masking operations. At the start of each table search by the filer, mask_register is reset to 0xFFFF_FFFF, which ensures that no masking occurs.

Filer rules may be configured to assign specific bit patterns to mask_register. Such rules can be configured to either match always (useful for implementing a default rule and specifying an associated receive queue), or fail always (which prevents termination of the filer table search). Once mask_register has been assigned, it retains its value until it is reassigned or the table search terminates. All properties are non-destructively bit-wise ANDed with mask_register prior to comparison in subsequent rules, which allows an entire cluster of rules to make use of a common mask. Individual masks for specific rules can also be created simply by combining a mask_register assignment (match always form) with a regular rule using the AND attribute.

To create a mask_register assignment rule, it is necessary to select PID = 0 in RQCTRL, and choose CMP such that the rule either matches (CMP = 01) or fails (CMP = 11). In this entry, RQPROP is then considered to be the assigned bit vector.

19.6.4.2.3 Special-Case Rules

It is frequently useful to create rules that are guaranteed to succeed or fail, specifically to enforce a default filing decision or act as null entries. Suggested constructions for such rules are shown in [Table 19-150](#).

Table 19-150. Special Filer Rules

| Rule Description | RQCTRL Fields | | | | | | | RQPROP Word | RQCTRL Word ¹ |
|--|---------------|------------------|-----|-----|---------|-----|------|-------------|--------------------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | |
| Default file—Always file frame to ring Q | 0 | 0 | 0 | 0 | Q | 01 | 0000 | All zeros | 0x0000_0q20 |
| Default reject—Always discard frame | 0 | 0 | 1 | 0 | 000_000 | 01 | 0000 | All zeros | 0x0000_0120 |
| Empty rule in AND—Always matches | 0 | 0/1 ² | 0 | 1 | 000_000 | 01 | 0000 | 0xFFFF_FFFF | 0x0000_00A0 |
| Empty rule in rule set—Always fails | 0 | 0/1 ³ | 0 | 0 | 000_000 | 11 | 0000 | 0xFFFF_FFFF | 0x0000_0060 |

¹ Hexadecimal digits *qq* denotes field Q shifted left 2 bits.

² Set CLE = 1 if the empty rule guards a cluster.

³ Set CLE = 1 if the empty rule occurs at the end of a cluster.

19.6.4.2.4 Filer Interrupt Events

The filer can produce three interrupt events in IEVENT. Event FIR indicates an error condition where the filer was unable to provide a definite result, either because no rule in the table succeeded, or because frames arrived too rapidly to complete searching of the table. Event FIQ indicates that the filer accepted a frame to a RxBD ring that was not enabled in RQCTRL (this can also occur if the filer is disabled, but RxBD ring 0—default queue or FSQEN mode queue—is not enabled). FIQ is also asserted in the case where no rule in the entire table succeeded. The various combinations of these interrupt events and their interpretation appear in [Table 19-151](#).

Table 19-151. Receive Queue Filer Interrupt Events

| IEVENT[FIR] | IEVENT[FIQ] | Description |
|-------------|-------------|---|
| 0 | 0 | No error. The filer successfully rejected or filed a frame. |
| 0 | 1 | Illegal queue error. The filer accepted a frame to a RxBD ring that is disabled (including ring 0 if filing is disabled). |
| 1 | 0 | Partial search error. The filer did not have sufficient time to complete its search of the filer table. |
| 1 | 1 | No matching rule error. The filer searched all 256 entries of the filer table without finding a rule that succeeds. |

A functional interrupt is provided via use of the general purpose interrupt (GPI) bit in the filer table. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the filer will set IEVENT[FGPI] when the corresponding receive frame is written to memory. This allows the user to set up a filer rule where the core will be interrupted upon the reception of ‘special’ frames.

If the timer is enabled (TMR_CTRL[TE] = 1), then the interrupt dedicated for timer events (in addition to the usual receive, transmit and error interrupts) will be asserted.

19.6.4.2.5 Setting Up the Receive Queue Filer Table

The eTSEC frame parser always provides values for all properties, even where the relevant headers are not available. In the latter case, the filer is given default properties that can be used to avoid conflict with normal, defined property values. Accordingly, the rules in the filer table can be partitioned into rule sets such that if all rules in a given set fail (due to headers being unavailable), lower priority rule sets can be subsequently searched until either a rule set provides a match or a single default—catch-all—rule specifies a definite receive queue. For example, an IEEE 802.1p priority rule set may be followed by an IP TOS rule set, followed by a default rule; thus, if no VLAN tag appears in the received frame, the TOS rules are checked, or the default is activated should no IP header be present.

The rule cluster feature is used to conditionalize evaluation of rule sets. Typically, this avoids evaluating rules based on properties that may not be valid or relevant to the filing or filtering decision. For example, TCP-related rules might be clustered behind a guard rule that checks that a TCP header has appeared and the IP address matches our home address. Property 1—the parse flags property—is provided specifically to check the characteristics of the received frame and the parser error status. The mask_register is typically assigned beforehand to extract specific flags, in which case care should be taken that mask_register be reassigned an appropriate mask vector for following comparisons.

In many cases it is possible to write the entire filer table before using eTSEC, as the rule set is static. However, dynamic rule updates can be supported by pre-allocating partially instantiated rule sets, which software rewrites as necessary. Rules that are not instantiated should be composed of empty entries, as indicated in [Table 19-150](#). In many cases empty entries can be overwritten by software without stopping eTSEC’s receive function.

19.6.4.2.6 Filer Example—802.1p Priority Filing

This example, shown in [Table 19-152](#), illustrates how to file frames according to layer 2 802.1p priority. This matches against property 1001, comparing each specific priority level in order to associate them with a RxBD ring index. Note that if a VLAN tag does not appear in the frame, the parser passes priority 0 to the filer, which always matches the rule at entry 7 and terminate the table search.

Table 19-152. Filer Table Example—802.1p Priority Filing

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---------------------------|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 0 | 0 | 0 | 0 | 0 | 000_000 | 00 | 1001 | 0x0000_0007 | File priority 7 to ring 0 | 0x0000_0009 |
| 1 | 0 | 0 | 0 | 0 | 000_001 | 00 | 1001 | 0x0000_0006 | File priority 6 to ring 1 | 0x0000_0409 |
| 2 | 0 | 0 | 0 | 0 | 000_010 | 00 | 1001 | 0x0000_0005 | File priority 5 to ring 2 | 0x0000_0809 |
| 3 | 0 | 0 | 0 | 0 | 000_011 | 00 | 1001 | 0x0000_0004 | File priority 4 to ring 3 | 0x0000_0C09 |
| 4 | 0 | 0 | 0 | 0 | 000_100 | 00 | 1001 | 0x0000_0003 | File priority 3 to ring 4 | 0x0000_1009 |
| 5 | 0 | 0 | 0 | 0 | 000_101 | 00 | 1001 | 0x0000_0002 | File priority 2 to ring 5 | 0x0000_1409 |

Table 19-152. Filer Table Example—802.1p Priority Filing (continued)

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|--|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 6 | 0 | 0 | 0 | 0 | 000_110 | 00 | 1001 | 0x0000_0001 | File priority 1 to ring 6 | 0x0000_1809 |
| 7 | 0 | 0 | 0 | 0 | 000_111 | 00 | 1001 | All zeros | File undefined 802.1p or priority 0 to ring 7—Default always matches | 0x0000_1C09 |

19.6.4.2.7 Filer Example—IP Diff-Serv Code Points Filing

This example demonstrates use of rule priority for determining class selector codepoints (RFC 2474) from the IP TOS property. An example filer table is shown in [Table 19-153](#). The example relies on the fact that the first rule matched terminates the search, hence successively lower Diff-Serv codepoint ranges can be compared in each step until the default (zero or greater) range is reached. By default, property 1010 (IP TOS) takes the value 0x00 if no IP headers were recognized, therefore the table search always terminates.

Table 19-153. Filer Table Example—IP Diff-Serv Code Points Filing

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 0 | 0 | 0 | 0 | 0 | 001_000 | 01 | 1010 | 0x0000_00E0 | File class 7 to queue 8 (TOS >= 0xE0) | 0x0000_202A |
| 1 | 0 | 0 | 0 | 0 | 001_001 | 01 | 1010 | 0x0000_00C0 | File class 6 to queue 9 (TOS >= 0xC0) | 0x0000_242A |
| 2 | 0 | 0 | 0 | 0 | 001_010 | 01 | 1010 | 0x0000_00A0 | File class 5 to queue 10 (TOS >= 0xA0) | 0x0000_282A |
| 3 | 0 | 0 | 0 | 0 | 001_011 | 01 | 1010 | 0x0000_0080 | File class 4 to queue 11 (TOS >= 0x80) | 0x0000_2C2A |
| 4 | 0 | 0 | 0 | 0 | 000_100 | 01 | 1010 | 0x0000_0060 | File class 3 to queue 4 (TOS >= 0x60) | 0x0000_102A |
| 5 | 0 | 0 | 0 | 0 | 001_100 | 01 | 1010 | 0x0000_0040 | File class 2 to queue 12 (TOS >= 0x40) | 0x0000_302A |
| 6 | 0 | 0 | 0 | 0 | 010_100 | 01 | 1010 | 0x0000_0020 | File class 1 to queue 20 (TOS >= 0x20) | 0x0000_502A |
| 7 | 0 | 0 | 0 | 0 | 011_100 | 01 | 1010 | All zeros | File class 0 to queue 28 (TOS >= 0x00) or file to ring 4 by default | 0x0000_702A |

19.6.4.2.8 Filer Example—TCP and UDP Port Filing

This example demonstrates rule clusters and AND-combined entries for filing packets based on transport protocol and well-known port numbers in a termination application. An example filer table is shown in [Table 19-154](#). The example contains two clusters; the first is entered only for TCP packets, the second is entered only for UDP packets. A default filing rule catches the case where neither TCP nor UDP headers are found. Each cluster compares source port number (property 1111) against a list of server ports, and files the packets accordingly. Note that entries 1 and 2 form an AND rule for checking that the port number >= 20 and port number < 22. Entries 4 and 5 are initially set up to always fail (zero port number), and thus comprise empty entries that can be used at a later time.

Table 19-154. Filer Table Example—TCP and UDP Port Filing

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 0 | 0 | 1 | 0 | 1 | 000_000 | 00 | 1011 | 0x0000_0006 | Enter cluster if layer 4 is TCP | 0x0000_028B |
| 1 | 0 | 0 | 0 | 1 | 000_000 | 01 | 1111 | 0x0000_0014 | AND rule—FTP from TCP ports 20 and 21: file to ring 2 | 0x0000_00AF |
| 2 | 0 | 0 | 0 | 0 | 000_010 | 11 | 1111 | 0x0000_0016 | | 0x0000_086F |
| 3 | 0 | 0 | 0 | 0 | 000_011 | 00 | 1111 | 0x0000_0017 | telnet from TCP port 23: file to ring 3 | 0x0000_0C0F |
| 4 | 0 | 0 | 0 | 0 | 000_000 | 00 | 1111 | All zeros | <i>empty entry reserved for future use</i> | 0x0000_000F |
| 5 | 0 | 0 | 0 | 0 | 000_000 | 00 | 1111 | All zeros | <i>empty entry reserved for future use</i> | 0x0000_000F |
| 6 | 0 | 1 | 0 | 0 | 000_001 | 01 | 0000 | All zeros | end cluster; default TCP: file to ring 1 | 0x0000_0620 |
| 7 | 0 | 1 | 0 | 1 | 000_000 | 00 | 1011 | 0x0000_0011 | Enter cluster if layer 4 is UDP | 0x0000_028B |
| 8 | 0 | 0 | 0 | 0 | 000_101 | 00 | 1111 | 0x0000_0801 | NFS from UDP port 2049 | 0x0000_140F |
| 9 | 0 | 0 | 0 | 0 | 000_111 | 00 | 1111 | 0x0000_0208 | Route from UDP port 520 | 0x0000_000F |
| 10 | 0 | 0 | 0 | 0 | 000_110 | 00 | 1111 | 0x0000_0045 | TFTP from UDP port 69 | 0x0000_180F |
| 11 | 0 | 1 | 0 | 0 | 000_100 | 01 | 0000 | All zeros | End cluster; default UDP: file to ring 4 | 0x0000_1220 |
| 12 | 0 | 0 | 0 | 0 | 000_000 | 01 | 0000 | All zeros | By default, file to ring 0 | 0x0000_0020 |

19.6.4.2.9 Filer Example—Interrupt from Deep Sleep

The example in [Table 19-155](#) shows how the filer can facilitate exit from deep sleep if any of the following packets arrive:

- ARP packet with Target IP address matching either of two IP addresses (either static or link local address)
- IPv4/UDP multicast DNS query
- IPv4/UDP SNMP broadcast query

These packets are also be stored in memory. All other packets are dropped.

Table 19-155. Filer Example—Interrupt from Deep Sleep

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|--|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | |
| 0 | 0 | 0 | 0 | 0 | 000_000 | 11 | 0000 | 0x0001_0000 | Check for ARP request; set mask register to mask off everything but the ARP request flag. |
| 1 | 0 | 1 | 0 | 1 | 000_000 | 00 | 0001 | 0x0001_0000 | Check to see if ARP request flag is set by doing a =1 comparison. Enter the “ARP Request Cluster” if true. |
| 2 | 0 | 0 | 0 | 0 | 000_000 | 11 | 0000 | 0xFFFF_FFFF | ARP Cluster: Set Mask to unmask everything (Reset mask to all F’s) |

Table 19-155. Filer Example—Interrupt from Deep Sleep (continued)

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | |
| 3 | 1 | 0 | 0 | 0 | 000_001 | 00 | 1100 | 0xXXXX_XXXX | Compare the ARP Target IP address against "MY_IP_1", which is indicated by the user-defined value of 0xXXXX_XXXX; if they match, accept the frame and generate an interrupt. |
| 4 | 1 | 0 | 0 | 0 | 000_001 | 00 | 1100 | 0xYYYY_YYYY | Compare the ARP Target IP address against "MY_IP_2", which is indicated by the user-defined value of 0xYYYY_YYYY; if they match, accept the frame. |
| 5 | 0 | 1 | 1 | 0 | 000_000 | 01 | 0000 | 0x0000_0000 | Default rule that will always discard the packet; inserted here because an ARP request was received, but the Target IP address did not match either local IP addresses; therefore drop the packet and exit the cluster. |
| 6 | 0 | 0 | 0 | 0 | 000_000 | 11 | 0000 | 0x0000_02D0 | Set Mask for IP4 Packet (2D0), with IPv4 checksum checked and verified, and UDP header located. |
| 7 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0001 | 0x0000_02D0 | Check to see if IP4 Packet, with IPv4 checksum checked and verified, and UDP header. |
| 8 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0000 | 0xFFFF_FFFF | Set Mask to unmask everything (Reset mask to all F's). |
| 9 | 0 | 1 | 0 | 1 | 000_000 | 00 | 1011 | 0x0000_0011 | Check against UDP protocol; if this passes, enter the cluster - all packets in the cluster are IPv4 packets with UDP protocol identified as the L4 protocol type. |
| 10 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0011 | 0x00XX_XXXX | Compare upper L2 DA bits to XX_XXXX (for multicast DNS query) |
| 11 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0100 | 0x00YY_YYYY | Compare lower L2 DA bits to YY_YYYY |
| 12 | 0 | 0 | 0 | 1 | 000_000 | 00 | 1100 | 0xZZZZ_ZZZZ | Compare L3 Destination IP address to ZZZZ_ZZZZ |
| 13 | 0 | 0 | 0 | 1 | 000_000 | 00 | 1110 | 0x0000_XXXX | Compare L4 destination port to XXXX |
| 14 | 1 | 0 | 0 | 0 | 000_001 | 00 | 1111 | 0x0000_YYYY | If all of the previously consecutive ANDed conditions pass, multicast DNS Query has matched. |
| 15 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0011 | 0x00XX_XXXX | Compare upper L2 DA bits to XX_XXXX (for SNMP broadcast query). |
| 16 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0100 | 0x00YY_YYYY | Compare lower L2 DA bits to YY_YYYY. |
| 17 | 1 | 0 | 0 | 0 | 000_001 | 00 | 1110 | 0x0000_ZZZZ | If all of the previously consecutive ANDed conditions pass, SNMP broadcast Query has matched. |

Table 19-155. Filer Example—Interrupt from Deep Sleep (continued)

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | |
| 18 | 0 | 1 | 1 | 0 | 000_000 | 01 | 0000 | 0x0000_0000 | Cluster End: IPv4, UDP Comparison Default rule that will always discard the packet; inserted here because an IPv4 packet with L4=UDP request was received, but the profiles didn't match anything "interesting"; therefore drop the packet and exit the cluster. |
| 19 | 0 | 0 | 1 | 0 | 000_000 | 01 | 0000 | 0x0000_0000 | Default rule that will always discard the packet; inserted here no matches for any "interesting" packets were received that are used to wake up the CPU. All packets that reach this rule are discarded. |

19.6.4.3 Transmission Scheduling

Each eTSEC can maintain multiple TxBD rings (or transmission queues) to satisfy QoS requirements. The ability to choose from a number of transmission streams dynamically is especially important during periods of network congestion. Certain application such as voice and video streaming are delay sensitive, but loss insensitive. For instance, VoIP applications require little bandwidth, but are highly sensitive to latency. Conversely, FTP or SMTP protocols are delay insensitive, but loss sensitive.

eTSEC has a transmission scheduler that implements a programmable QoS regime. The scheduler is responsible for choosing which of the prefetched TxBDs shall be processed next, and accordingly issuing DMA requests to service the data stream described by the chosen BD(s). The scheduler cycle is one of:

1. decide on a TxBD queue,
2. transmit exactly one frame from that queue, and
3. return to deciding on another queue, in step 1.

If TCTRL[TXSCHED] is set to 00, no transmission scheduling occurs, and only TxBD ring 0 is polled for new data to transmit, with DMACTRL controlling waiting or polling. TCTRL[TXSCHED], if not zero, can be programmed to invoke one of two scheduling algorithms, namely priority-based queuing (PBQ), and modified weighted round-robin queuing (MWRR). In all cases where TCTRL[TXSCHED] is not zero, the scheduler can choose from among 1 to 8 TxBD rings per eTSEC, with individual rings being enabled by the setting of TQUEUE[EN0–EN7] bits. For example, TxBD rings 3, 4, and 7 may be enabled for scheduling by setting EN3, EN4, and EN7, and clearing all other EN bits.

19.6.4.3.1 Priority-Based Queuing (PBQ)

PBQ is the simplest scheduler decision policy. The enabled TxBD rings are assigned a priority value based on their index. Rings with a lower index have precedence over rings with higher indices, with priority assessed on a frame-by-frame basis. For example, frames in TxBD ring 0 have higher priority than frames in TxBD ring 1, and frames in TxBD ring 1 have higher priority than frames in TxBD ring 2, and so on.

The scheduling decision is then achieved as follows:

```
loop
    # start or S/W clear of TSATn
```

```

ring = 0;
while ring <= 7 loop
    if enabled(ring) and not ring_empty(ring) then
        transmit_frame(ring);
        ring = 0;
    else
        ring = ring + 1;
    endif
endloop
endloop
    
```

19.6.4.3.2 Modified Weighted Round-Robin Queuing (MWRR)

eTSEC implements a modified weighted round-robin scheduling algorithm across all enabled TxBD rings when TCTRL[TXSCHEM] = 10. In MWRR, the weights in the TR03WT and TR47WT registers determine the ideal size of each transmit slot, as measured in multiples of 64 bytes. Thus, to set a transmit slot of 512 bytes, a weight of 512/64 or 8 needs to be set for the ring. In this mode TxBD rings 1–7 are selected in round-robin fashion, whereas TxBD ring 0, if enabled with ready data for transmission, is always selected in between other rings so as to expedite transmission from ring 0.

The scheduling decision is then achieved as follows:

```

for ring = 1..7 and enabled(ring) loop
    credit[ring] = 0;
endloop
for ring = 1..7 and enabled(ring) loop
    if not ring_empty(0) then
        credit[0] = credit[0] + weight[0];
        while credit[0] > 0 loop
            transmit_frame(0);
            credit[0] = credit[0] - frame_size;
            if ring_empty(0) then
                credit[0] = 0;
            endif
        endloop
    endif
    if not ring_empty(ring) then
        credit[ring] = credit[ring] + weight[ring];
    endif
    while credit[ring] > 0 loop
        transmit_frame(ring);
        credit[ring] = credit[ring] - frame_size;
        if ring_empty(ring) then
            credit[ring] = 0;
        endif
    endloop
endloop
endloop
    
```

The algorithm checks registers TQUEUE[EN0–EN7] for `enabled()`, TSTAT[THLT0–THLT7] for `ring_empty()`, and TRxWT for `weight()`. For TxBD ring k , having a weight WT_k , the long term average throughput for that ring is:

$$\text{rate of queue}[k] \text{ (} K = 1 \text{ to } 7) = (\text{available bandwidth}) * WT_k / (\text{sum}(WT_i) + 6WT_0)$$

$$\text{rate of queue}(0) = (\text{available bandwidth}) * 7 * WT_0 / (\text{sum}(WT_i) + 6WT_0)$$

where $i = 0$ to 7

19.6.5 Lossless Flow Control

The eTSEC DMA subsystem is designed to be able to support simultaneous receive and transmit traffic at gigabit line rates. If the host memory has sufficient bandwidth to support such line rates, then the principle cause of overflow on receive traffic is due to a lack of Rx BDs. Thus, the long term receive throughput is determined by the rate at which software can process receive traffic. If a user desires to prevent dropped packets, they can inform the far-end link to stop transmission while the software processing catches up with the backlog.

To avoid overflow in the latter case, back pressure must be applied to the far-end transmitter before the Rx descriptor controller encounters a non-empty BD and halts with a BSY error. As there is lag between application of back-pressure and response of the far-end, the pause request must be issued while there are still BDs free in the ring. In the traditional eTSEC descriptor ring programming model, there is no way for hardware to know how many free BDs are available, so software must initiate any pause requests required during operation. If software is backlogged, the request may not be issued in time to prevent BSY errors. To allow the eTSEC to generate the pause request automatically, additional information (a pointer to the last free BD and ring length) is required.

19.6.5.1 Back Pressure Determination through Free Buffers

Ultimately, the rate of data reception is determined by how quickly software can release buffers back into the receive ring(s). Each time a buffer is freed, the associated BD has its empty bit set and hardware is free to consume both. Thus the number of free BDs in a given Rx ring indicates how close hardware is to the end of that ring. To prevent data loss, back pressure should be applied when the number of free BDs drops below some critical level. The number of BDs that can be consumed by an incoming packet stream while back-pressure takes effect is determined by several factors, such as: receive traffic profile, transmit traffic profile, Rx buffer size, physical transmission time between eTSEC and far-end device and intra-device latency. Theoretically, the worst case is as follows:

$$\text{FreeBDsRequired} = \frac{\text{MaxFrameSize}}{\text{MinFrameSize} + \text{IFG}} + \frac{\text{MaxFrameSize}}{\text{RxBufferSize}} + \text{LinkDelay}$$

This case comes about when:

- The eTSEC has just started transmitting a large frame and thus cannot send out a pause frame
- Upon reception of the pause request the far-end has just started transmission of a large frame
- The eTSEC receives a burst of short frames with minimum inter-frame-gap (96 bit times for Ethernet)

Once the user has determined the worst case scenario for their application, they program the required free BD threshold into the eTSEC (through RQPRM[PBTHR]). Since different BD rings may have different sizes and expected packet arrival rates, a separate threshold is provided for each active ring. It is recommended that a threshold of at least four BDs is the practical minimum for gigabit Ethernet links.

For the Rx descriptor controller to determine the number of free BDs remaining in the ring, it needs to know the following:

1. The location of the current BD being used by hardware
2. The location of the last BD that was released (freed) by software
3. The length of the Rx BD ring.

For each active ring, the current BD pointer ($RBPTR_n$) is maintained by the eTSEC. Software knows both the size of the Rx ring and the location of the last freed BD. By providing the eTSEC with those values (through $RQPRM[LEN]$ and $RFBPTR$ respectively) the eTSEC always know how many receive buffers are available to be consumed by incoming data.

The number of guaranteed free BDs in the ring is then determined by:

When $RFBPTR_n < RBPTR_n$

$$\text{FreeBDs} = RQPRM_n[LEN] - RBPTR_n + RFBPTR_n$$

When $RFBPTR_n > RBPTR_n$

$$\text{FreeBDs} = RFBPTR_n - RBPTR_n$$

When $RBPTR_n = RFBPTR_n$ the number of free BDs in the ring is either one (since $RFBPTR_n$ points to a free BD) or equal to the ring length. Since the BD pointed to by $RBPTR_n$ may be either in use or about to be used, it is not considered in the free BD count. To resolve the case where the two pointers collide, the following logic applies:

If $RBASE_n$ was updated and thus initializes both $RBPTR_n$ and $RFBPTR_n$, the ring is deemed empty.

If $RFBPTR_n$ is updated by a software write and matches $RBPTR_n$, the ring is deemed empty.

If HW updates $RBPTR_n$ and the result matches $RFBPTR_n$, the ring is deemed to have one BD remaining. Upon writing this BD back to memory (indicating the buffer is occupied) the ring is deemed to be full.

Important. There is a possibility that if software is severely backlogged in updating $RFBPTR_n$, the hardware could wrap around the ring entirely, consume exactly the remaining number of BDs and not halt with a BSY error. If software then increments $RFBPTR_n$ to the next address (thereby equalling $RBPTR_n$), the hardware assumes the ring is now empty (when in fact there is only a single BD freed up). This results in the hardware failing to maintain back pressure on the far end. Upon software incrementing $RFBPTR_n$ a subsequent time, the wrap condition is successfully detected and hardware recognizes a nearly full ring (rather than a nearly empty one). Since software can increment $RFBPTR_n$ by any amount, it is not possible for hardware to determine in this case whether the user has cleared the entire ring or just one BD. Users can eliminate the possibility of this condition occurring by ensuring that $RFBPTR_n$ is incremented by at least two BDs each time (that is, clear at least two buffers whenever the RxBD unload routine is called).

Once the eTSEC determines that this threshold has been reached, back pressure is applied accordingly. The type of back pressure that is applied varies according to the physical interface that is used.

- **Half duplex Ethernet:** No support in this mode.

- **Full duplex Ethernet:** An IEEE 802.3 PAUSE frame (see sect. 19.6.2.9/19-151) is issued as if the TCTRL[TFC_PAUSE] bit was set. An internal counter tracks the time the far end controller is expected to remain in pause (based on the setting of PTV[PT]). When that counter reaches half the value of PTV[PT], the eTSEC reissues a pause frame if the free BD calculation for any ring is below the threshold for that ring. For example, if PTV[PT] is set to 10 quanta, a pause frame is re-issued when five quanta have elapsed if the free BD threshold is still not met. A practical minimum for PTV[PT] of 4 quanta is recommended.
- **FIFO packet interface:** Link layer flow control is asserted through use of the RFC signal (CRS pin). Flow control is asserted for the entire time that free BD threshold is not met. The same mechanism is used for both GMII-style and encoded packet modes.

19.6.5.2 Software Use of Hardware-Initiated Back Pressure

19.6.5.2.1 Initialization

Software configures RBASE n and RQPRM n [LEN] according to the parameters for that ring. Then the number of free BDs that are required to prevent the eTSEC from automatically asserting flow control are programmed in RQPRM[FBTHR]. The receiver is then enabled.

Note: the act of programming RBASE n initializes RFBPTR n to the start of the of the ring. When the ring is in this initial empty state, there is no concept of a last freed BD. In this case, the calculated number of free BDs is the size of the ring. Since the BD that the hardware is currently pointing to is to be considered in-use, the free BD count is actually one higher than the total available. As soon as the hardware consumes a BD (by writing it back to memory), RBPTR n advances and the free BD count reflects the correct number of available free BDs.

19.6.5.2.2 Operation

As software frees BDs from the ring, it writes the physical address of the BD just freed to RFBPTR n . The eTSEC asserts flow control if the distance (using modulo arithmetic) between RBPTR n and RFBPTR n is $< RQPRM[FBTHR]$. In multi-ring operation, if the free BD count of **any** active ring drops below the threshold for that ring, flow control is asserted. Once enough BDs are freed for **all** active rings to meet their respective free BD thresholds, application of back pressure cases.

Note: The eTSEC does not issue an exit pause frame (that is, pause frame with PTV of 0x0000) once all active rings have sufficient BDs. Instead, it waits for the far-end pause timer to expire and start re-transmission.

19.6.6 Hardware Assist for IEEE 1588-Compatible Timestamping

There is a push in industrial control applications to use Ethernet as the principal link layer for communications. This requires Ethernet to be used for both data transfer and real-time control. For real-time systems, each node is required to be synchronized to a master clock. The precision of this clock is dictated by the application, but generally needs to be of the order of $<1\mu\text{Sec}$ for high-speed machinery (for example, printing presses).

IEEE 1588 [1588] specifies a mechanism for synchronizing multiple nodes to a master clock. Support for 1588 can be done entirely in software running on a host CPU, but applications that require sub 10uSec accuracy will need hardware support for accurate timestamping of incoming packets.

The eTSEC includes a new timer clock module to support the IEEE Std. 1588 timer standard. The following sections describe the features, programming model, and implementation information.

NOTE

IEEE 1588 timestamping is not supported in conjunction with the SGMII 10/100 interface mode.

19.6.6.1 Features

- 64-bit free running timer running from an external oscillator or internal clock
- Programmable timer oscillator clock selection
- Self-correcting precision timer with nano-second resolution
- Time stamp all incoming packets inline
 - Maskable interrupts on received PTP packet's filter rule match
- Time stamp transmit packets when instructed in the TxFCB
 - Maskable interrupts on transmit timestamp capture
- Two Tx time stamp registers per eTSEC with 16-bit tag for each of them to support burst mode.
- Time stamp capture on two general-purpose external triggers
 - Maskable interrupts on GPIO timestamp trigger
 - Programmable polarity of external trigger (GPIO) edge
- Two 64-bit alarm (future time) registers for future time comparison
 - Maskable interrupts on alarm
- Three programmable timer output pulse period phase aligned with 1588 timer clock
 - Maskable interrupts associated with each pulse
- Separate maskable timer interrupt event register
- Recognition of incoming PTP packet through filter rule match
- Phase aligned adjustable (divide by N) clock output
- Supports all Ethernet modes supported by the eTSEC, including full- and half-duplex modes
- Supports both master and slave modes
- Supports timestamp of nano-second resolution

19.6.6.2 Timer Logic Overview

The 1588 timer module can be partitioned into four different sub-modules as shown in [Figure 19-141](#).

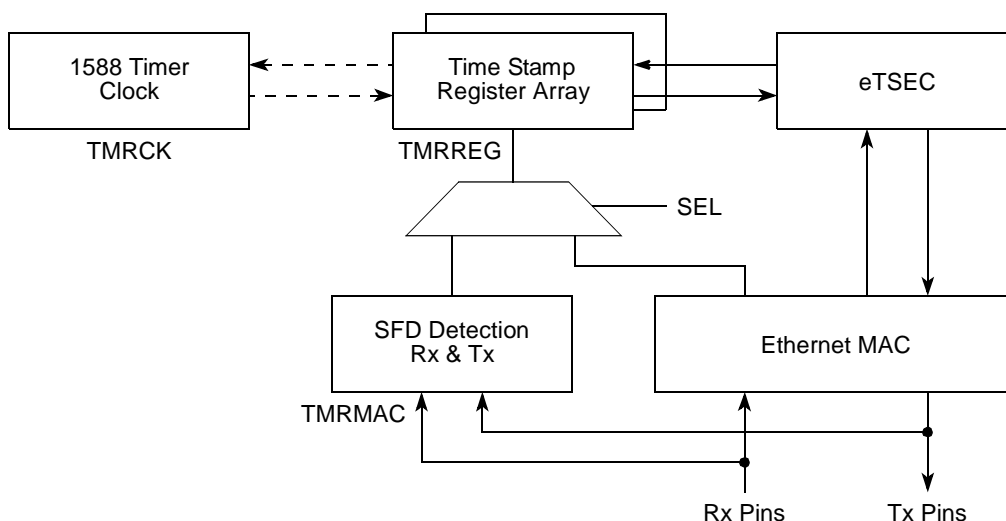


Figure 19-141. 1588 Timer Design Partition

19.6.6.3 Time-Stamp Insertion on the Received Packets

Every incoming packet's 8-byte time stamp is inserted into the packet data buffer as padding alignment bytes. Time-stamp insertion into the data buffer requires RCTRL[PAL] to be set to a value greater than or equal to 8 and the control bit RCTRL[TS] bit to be set.

19.6.6.3.1 Timestamp Point

The required timestamp point, as specified in the IEEE 1588 Specification Sep-2004 (IEC 61588 First Edition), is shown in [Figure 19-142](#). From this, it is clear that the end of the SFD is the critical point in the MII data stream.

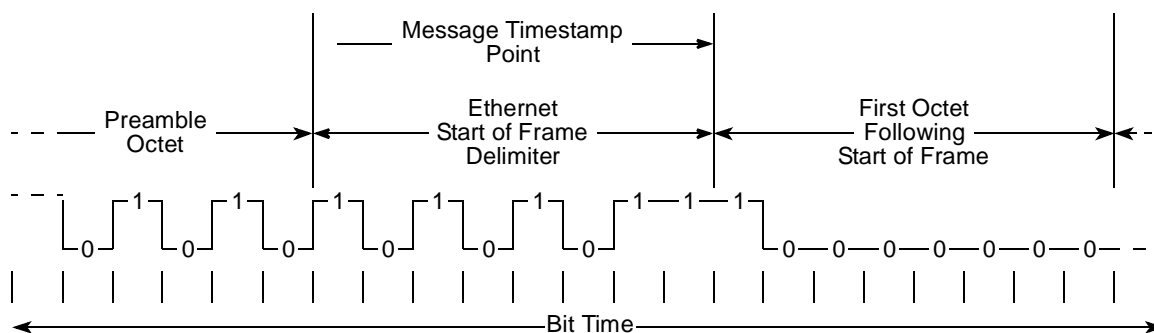


Figure 19-142. Ethernet Sampling Points for 1588

The sample point coincides with the cycle after the SFD (Start of Frame Delimiter) detection by the MAC. For received frames, this will be at least 4 bit times (MII) or 8 bit times (GMII) after the message timestamp point specified in [1588]. For transmission, the eTSEC sample point precedes the sample point specified in [1588] by at least 4-bit times (MII) or 8-bit times (GMII). For a particular mode, the eTSEC

sample point is a consistent number of bit times relative to the SFD detection. Thus, the offset from the [1558] specified sample point can be accounted for in the PTP software implementation.

19.6.6.4 PTP Packet Parsing

PTP packets are typically embedded within a UDP payload with special IP source and destination address and special source and destination ports numbers. Special fields of interest of a PTP packet are listed in [Table 19-156](#).

Table 19-156. PTP Payload Special Fields

| Layer | Octet (Offset from the SFD) | Field | Value | eTSEC filer PID | Comments |
|------------|-----------------------------|---|--|---|---|
| Ethernet | 12-13 | Length/Packet | 0x0800 | ETY-RQPFR[PID=0111] | IPv4 |
| IP header | 22 | Time to live | 0x00 | RBIFX-choose an arbitrary extraction byte | Must be 0 |
| IP header | 23 | IP Protocol | 0x11 | L4P-RQPFR[PID=1011] | UDP |
| IP header | 26-29 | Source IP Address IANA defines 4 multicast address for the PTP packet | | SIA-RQPFR[PID=1101] | |
| IP header | 30-33 | Destination IP Address IANA defines 4 multicast address for the PTP packet | 224.0.1.129 224.0.1.130 224.0.1.131 224.0.1.132 | DIA-RQPFR[PID=1100] | DefaultPTPdomain AlternatePTPdomain1 AlternatePTPdomain2 AlternatePTPdomain3 |
| UDP header | 34-35 | Source port number | | SPT-RQPFR[PID=1011] | |
| UDP header | 36-37 | Destination port number | 319 320 | DPT-RQPFR[PID=1011] | EventPort GeneralPort |
| UDP data | 74 | Control | 0x0 0x1 0x2 0x3 0x4 | RBIFX-choose an arbitrary extraction byte | Sync Delay_req Follow_up Delay_resp Management |

A representation of the PTP packet is shown in [Figure 19-143](#).




Figure 19-143. PTP Packet Format

19.6.6.4.1 General Purpose Filer Rule

The eTSEC receive filer has been enhanced with the addition of a general-purpose event bit. This event bit can be used in conjunction with filing table rules to identify 1588 packets and indicate these packets by setting special timer status register bits (TMR_STAT). Additionally, 1588 packets can be easily identified by upper-layer software by using the filer to queue all PTP packets to one or more predefined virtual queues. See Section 19.6.4.2.1, “Filing Rules” for further information.

19.6.6.5 Time-Stamp Insertion on Transmit Packets

Software has the option to write the time stamp of the transmitted frame to memory in the padding alignment bytes (PAL) located between the TxFCB and the frame data. It is required that a minimum of two TxBDs are used. The first points to the start of the 8 byte TxFCB. The second points to the start of frame data. In memory, the TxFCB, and at least the first 16 bytes of the TxPAL must be adjacent, that is, located in contiguous memory locations, as depicted in Figure 19-144.

The first TxBD[TOE] bit is set. When the TMR_CTRL[Record Time-stamp In PAL Enable] and TxFCB[PTP] bits are set, the timestamp is written to memory location TxBD[Data Buffer Pointer]+16.

The second TxBD’s Data Length must either contain the full frame length, or a value greater than the TxThreshold setting. Refer to Table 19-157. When time-stamps are inserted into the TxPAL, the TMR_TXTSn_H/L and TMR_TXTSn_ID registers still function normally.

19.6.6.5.1 Interrupts

The TxPAL is updated with a time-stamp before closing the second TxBD. The TxBD[I] bit can be set for the second TxBD frame to cause an interrupt (via IEVENT[TXF]) after the time-stamp has been written to the TxPAL.

When time-stamps are inserted into the TxPAL, the TMR_TXTSn_H/L and TMR_TXTSn_ID registers still function normally. Therefore, the 1588 interrupt can be triggered by using the TMR_PEVENT register bits TXP1, and TXP2.

Table 19-157. Time-Stamp Insertion Programming Requirements

| Requirement | Behavior if requirement is not met |
|------------------|--|
| TMR_CTRL[RTPE]=1 | If TMR_CTRL[RTPE]=0, then no time-stamp is written to a TxPAL. |
| TxBD[TOE]=1 | If TxBD[TOE]=0, then no time-stamp is written to a TxPAL. |

Table 19-157. Time-Stamp Insertion Programming Requirements

| Requirement | Behavior if requirement is not met |
|--|--|
| First TxBD[Data Buffer Pointer] is 8-byte aligned | The time-stamp will be written to address First TxBD[Data Buffer Pointer] + 0x10 rounded down to the nearest 8-byte aligned address, except at 4K page boundaries, in which case the time-stamp may be invalid, and the Second TxBD close status will be lost. |
| First TxBD[Data Length]=8, 8 bytes for TxFCB | If L2 or frame data is included in the Length, the buffer immediately following the FCB is transmitted on the line and the frame data stored in memory will be overwritten with a time-stamp value after the frame is transmitted. |
| TxFCB[PTP]=1 | If TxBD[PTP]=0, then no time-stamp is written to a TxPAL. |
| The TxFCB is followed immediately by a minimum of 16 bytes for the TxPAL | The time-stamp will be written to address First TxBD[Data Buffer Pointer] + 0x10. |
| Second TxBD[Data Buffer Pointers] points to start of L2 or frame data | If there is only one TxBD used to transfer a PTP frame, then no time-stamp is written to a TxPAL. |
| Second TxBD[Data Length] ≥ FIFO_TX_THR or includes the entire frame | If this condition is not true, the time-stamp in TxPAL is invalid. |

Figure 19-144 depicts the buffer format requirements for time-stamp insertion on transmit packets.




Figure 19-144. Buffer Format for Transmit Time-Stamp Insertion

19.6.6.5.2 Error Condition

When an error is encountered after a PTP packet has begun to be processed, the time-stamp written to the TxPAL is zero. Subsequent frames may be flushed by eTSEC. There will be no time-stamp update to TxPAL for the subsequent flushed frames.

19.6.6.6 Tx PTP Packet Parsing

Software instructs the Tx packet to be timestamped via setting bit 15 in the TxFCB to mark a PTP packet. TxFCB[VLCTL] can be translated as the Tx PTP packet identification number. BD[TOE] must be set to enable transmit PTP packet time stamping. TxFCB[PTP] bit takes precedence over TxFCB[VLN] bit. It disables per packet VLAN tag insertion. On a PTP packet, a VLAN tag can be inserted from the DFVLAN register. The TxFCB for the PTP packet is shown in [Figure 19-145](#).

| | | | | | | | | | | | | | | | | | |
|------------|--------------|----|-----|-----|-----|-----|-----|------|---|---|----|----|----|----|----|----|-----|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | |
| Offset + 0 | VLN | IP | IP6 | TUP | UDP | CIP | CTU | NPH | | | | | | | | | PTP |
| Offset + 2 | L4OS | | | | | | | L3OS | | | | | | | | | |
| Offset + 4 | PHCS | | | | | | | | | | | | | | | | |
| Offset + 6 | VLCTL/PTP_ID | | | | | | | | | | | | | | | | |

Figure 19-145. Transmit Frame Control Block

The contents of the Tx FCB are defined in [Table 19-158](#).

Table 19-158. Tx Frame Control Block Description

| Bytes | Bits | Name | Description |
|-------|------|------|---|
| 0–1 | 0 | VLN | VLAN control word valid. This bit is ignored when the PTP bit is set. VLAN tag is read from the DFVLAN register if PTP=1. 0 Ignore VLCTL field. 1 If VLAN tag insertion is enabled for eTSEC, use the VLCTL field as the VLAN control word. |
| | 1 | IP | Layer 3 header is an IP header. 0 Ignore layer 3 and higher headers. 1 Assume that the layer 3 header is an IPv4 or IPv6 header, and take L3OS field as valid. |
| | 2 | IP6 | IP header is IP version 6. Valid only if IP = 1. 0 IP header version is 4. 1 IP header version is 6. |
| | 3 | TUP | Layer 4 header is a TCP or UDP header. 0 Do not process any layer 4 header. 1 Assume that the layer 4 header is either TCP or UDP (see UDP bit), and offload checksumming on the basis that the IP header has no extension headers. |
| | 4 | UDP | UDP protocol at layer 4. 0 Layer 4 protocol is either TCP (if TUP = 1) or undefined. 1 Layer 4 protocol is UDP if TUP = 1. |

Table 19-158. Tx Frame Control Block Description (continued)

| Bytes | Bits | Name | Description |
|-------|------|------------------|---|
| 0–1 | 5 | CIP | Checksum IP header enable. 0 Do not generate an IP header checksum. 1 Generate an IPv4 header checksum. |
| | 6 | CTU | Checksum TCP or UDP header enable. 0 Do not generate a TCP or UDP header checksum. RFC 768 advises that UDP packets not requiring checksum validation should have their checksum field set to zero. 1 Generate a TCP header checksum if IP = 1 and TUP = 1 and UDP = 0. |
| | 7 | NPH | Disable calculation of TCP or UDP pseudo-header checksum. This bit should be set if IP options need to be consulted in forming the pseudo-header checksum, as eTSEC does not examine IP options or extension headers for TCP/IP offload on transmit. 0 Calculate TCP or UDP pseudo-header checksum as normal, assuming that the IP header has no options. 1 Do not calculate a TCP or UDP pseudo-header checksum, but instead use the value in field PHCS when determining the overall TCP or UDP checksum. |
| | 8–14 | — | Reserved |
| | 15 | PTP | Indication to the transmitter that this is a PTP packet. Enabling PTP disables per packet VLAN tag insertion. Instead, VLAN tag will be read from the DFVLAN when the PTP field is true. 0 Do not attempt to capture transmission event time 1 Valid PTP_ID field. When this packet is transmitted, capture the time of transmission. Must be clear if TMR_CTRL[TE] is clear. |
| 2–3 | 0–7 | L4OS | Layer 4 header offset from start of layer 3 header. The layer 4 header starts L4OS octets after the layer 3 header if it is present. The maximum layer 3 header length supported is thus 255 bytes, which may prevent TCP/IP offload on particularly large IPv6 headers. |
| | 8–15 | L3OS | Layer 3 header offset from start of frame not including the 8 bytes for this FCB. The layer 3 header starts L3OS octets from the start of the frame including any custom preamble header that may be present. The maximum layer 2 header length supported is thus 255 bytes. |
| 4–5 | 0–15 | PHCS | Pseudo-header checksum (16-bit one's complement sum with carry wraparound, but without result inversion) for TCP or UDP packets, calculated by software. Valid only if NPH = 1. |
| 6–7 | 0–15 | VLCTL/ PTP_ID | VLAN control word for insertion in the transmitted VLAN tag. Valid only if VLN = 1. Tx PTP packet identification number. This number will be copied into the Tx PTP packet time stamp identification field. PTP field takes precedence over VLN field. |

19.6.7 Buffer Descriptors

The eTSEC buffer descriptor (BD) is modeled after the MPC8260 Fast Ethernet controller BD for ease of reuse across the PowerQUICC network processor family. Drawing from the MPC8260 FEC BD programming model, the eTSEC descriptor base registers point to the beginning of BD rings. The eTSEC BD also expands upon the MPC8260 BD model to accommodate the eTSEC's unique features. However, the 8-byte data BD format is designed to be compatible with the existing MPC8260 BD model.

19.6.7.1 Data Buffer Descriptors

Data buffers are used in the transmission and reception of Ethernet frames (see [Figure 19-146](#)). Data BDs encapsulate all information necessary for the eTSEC to transmit or receive an Ethernet frame. Within each data BD there is a status field, a data length field, and a data pointer. The BD completely describes an

Ethernet packet by centralizing status information for the data packet in the status field of the BD and by containing a data BD pointer to the location of the data buffer. Software is responsible for setting up the BDs in memory. Because of pre-fetching, a minimum of four buffer descriptors per ring are required. This applies to both the transmit and the receive descriptor rings. Transmit rings are limited to a maximum size of 65536 BDs due to BD and frame data prefetching. Software also must have the data pointer pointing to a legal memory location. Within the status field, there exists an ownership bit which defines the current state of the buffer (pointed to by the data pointer). Other bits in the status field of the buffer descriptor are used to communicate status/control information between the eTSEC and the software driver.

Because there is no next BD pointer in the transmit/receive BD (see [Figure 19-147](#)), all BDs must reside sequentially in memory. The eTSEC increments the current BD location appropriately to the next BD location to be processed. There is a wrap bit in the last BD that informs the eTSEC to loop back to the beginning of the BD chain. Software must initialize the TBASE and RBASE registers that point to the beginning transmit and receive BDs for eTSEC.

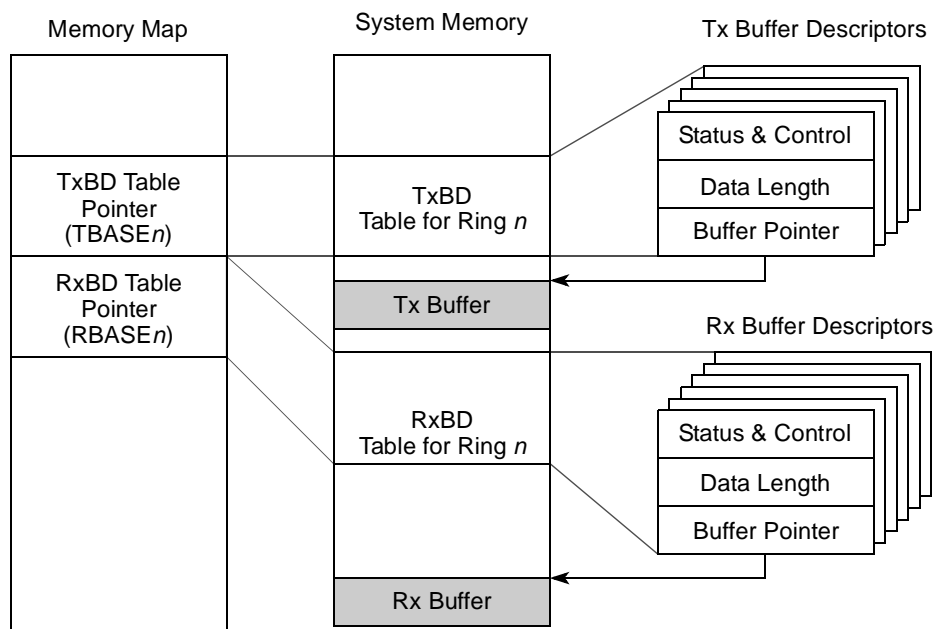


Figure 19-146. Example of eTSEC Memory Structure for BDs




Figure 19-147. Buffer Descriptor Ring

19.6.7.2 Transmit Data Buffer Descriptors (TxBD)

Data is presented to the eTSEC for transmission by arranging it in memory buffers referenced by the TxBDs. In the TxBD the user initializes the R, PAD, W, I, L, TC, PRE, HFE, CF, and TOE bits and the length (in bytes) in the first word, and the buffer pointer in the second word. Unused fields or fields written by the eTSEC must be initialized to zero.

The eTSEC clears the R bit in the first word of the BD after it finishes using the data buffer. The transfer status bits are then updated. Additional transmit frame status can be found in statistic counters in the MIB block.

Software must expect eTSEC to prefetch multiple TxBDs, and for TCP/IP checksumming an entire frame must be read from memory before a checksum can be computed. Accordingly, the R bit of the first TxBD in a frame must not be set until at least one entire frame can be fetched from this TxBD onwards. If eTSEC prefetches TxBDs and fails to reach a last TxBD (with bit L set), it halts further transmission from the current TxBD ring and report an underrun error as IEVENT[XFUN]; this indicates that an incomplete frame was fetched, but remained unprocessed. The relevant TBPTR register points to the next unread TxBD following the error.

Figure 19-148 defines the TxBD.

| | | | | | | | | | | | | | | | | |
|------------|------------------------|---------|---|---|---|----|---------|---|--------|-------|----|----|----|--------|----|----|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
| Offset + 0 | R | PAD/CRC | W | I | L | TC | PRE/DEF | 0 | HFE/LC | CF/RL | RC | | | TOE/UN | TR | |
| Offset + 2 | DATA LENGTH | | | | | | | | | | | | | | | |
| Offset + 4 | TX DATA BUFFER POINTER | | | | | | | | | | | | | | | |
| Offset + 6 | | | | | | | | | | | | | | | | |

Figure 19-148. Transmit Buffer Descriptor

The TxBD definition is interpreted by eTSEC hardware as if TxBDs mapped to C data structures in the manner illustrated by [Figure 19-149](#).

```
typedef unsigned short uint_16; /* choose 16-bit native type */
typedef unsigned int uint_32; /* choose 32-bit native type */
typedef struct txbd_struct {
    uint_16 flags;
    uint_16 length;
    uint_32 bufptr;
} txbd;
```

Figure 19-149. Mapping of TxBDs to a C Data Structure

The TxBD fields are detailed in [Table 19-159](#).

Table 19-159. Transmit Data Buffer Descriptor (TxBD) Field Descriptions

| Offset | Bits | Name | Description |
|--------|------|---------|--|
| 0-1 | 0 | R | Ready, written by eTSEC and user. 0 The data buffer associated with this BD is not ready for transmission. The user is free to manipulate this BD or its associated data buffer. The eTSEC clears this bit after the buffer is transmitted or after an error condition is encountered. 1 The data buffer, which is prepared for transmission by the user, was not transmitted or is currently being transmitted. No fields of this BD may be written by the user once this bit is set. |
| | 1 | PAD/CRC | Padding for frames. (Valid only while it is set in the first BD and MACCFG2[PAD enable] is cleared). If MACCFG2[PAD enable] is set, this bit is ignored. 0 Do not add padding to short frames. 1 Add PAD to frames. PAD bytes are inserted until the length of the transmitted frame equals 64 bytes. Unlike the MPC8260 which PADs up to MINFLR value, the eTSEC PADs always up to the IEEE minimum frame length of 64 bytes. CRC is always appended to frames. |
| | 2 | W | Wrap. Written by user. 0 The next buffer descriptor is found in the consecutive location. 1 The next buffer descriptor is found at the location defined in TBASE. |
| | 3 | I | Interrupt. Written by user. 0 No interrupt is generated after this buffer is serviced. 1 IEVENT[TXB] or IEVENT[TXF] are set after this buffer is serviced. These bits can cause an interrupt if they are enabled (That is, IEVENT[TXBEN] or IEVENT[TXFEN] are set). |
| | 4 | L | Last in frame. Written by user. 0 The buffer is not the last in the transmit frame. 1 The buffer is the last in the transmit frame. |

Table 19-159. Transmit Data Buffer Descriptor (TxBD) Field Descriptions (continued)

| Offset | Bits | Name | Description |
|--------|-------|------|---|
| 0–1 | 5 | TC | Tx CRC. Written by user. (Valid only while it is set in first BD and TxBD[PAD/CRC] is cleared and MACCFG2[PAD/CRC enable] is cleared and MACCFG2[CRC enable] is cleared.) If MACCFG2[PAD/CRC enable] is set or MACCFG2[CRC enable] is set, this bit is ignored in Ethernet modes. 0 End transmission immediately after the last data byte with no hardware generated CRC appended, unless TxBD[PAD/CRC] is set. 1 Transmit the CRC sequence after the last data byte. |
| | 6 | PRE | Transmit user-defined Ethernet preamble. Written by user. Valid only if set in the first BD of a frame, and MACCFG2[PreAm TxEN] is set. 0 This frame does not contain Ethernet preamble bytes for transmission. 1 This frame includes a user-defined Ethernet preamble sequence prior to the destination address in the data buffer. |
| | | DEF | Defer indication. The eTSEC updates this bit after transmitting a frame (TxBD[L] is set) 0 This frame was not deferred. 1 This frame did not have a collision before it was sent but it was sent late because of deferring |
| | 7 | — | Reserved |
| | 8 | HFE | Huge frame enable. Written by user. Valid only if set in the first BD of a frame and MACCFG2[Huge Frame] is cleared. If MACCFG2[Huge Frame] is set, this bit is ignored. 0 Truncate transmit frame if its length is greater than the MAC's maximum frame length. 1 Allow large frames to be transmitted without truncation. |
| | | LC | Late collision. Written by the eTSEC. 0 No late collision. 1 A collision occurred after 64 bytes are sent. The eTSEC terminates the transmission and updates LC. |
| | 9 | CF | Control Frame. Written by user. Valid only if set in the first BD of a frame. 0 Regular frame; transmission is deferred when eTSEC is in PAUSE. 1 Control frame; transmission starts even if eTSEC is in PAUSE. |
| | | RL | Retransmission Limit. Written by the eTSEC. 0 Transmission before maximum retry limit is hit. 1 The transmitter failed (max. retry limit + 1) attempts to successfully send a message due to repeated collisions. The eTSEC terminates the transmission and updates RL. |
| | 10–13 | RC | Retry Count. Written by the eTSEC. 0 The frame is sent correctly the first time. x One or more attempts where needed to send the transmit frame. If this field is 15, then 15 or more retries were needed. The Ethernet controller updates RC after sending the buffer. |

Table 19-159. Transmit Data Buffer Descriptor (TxBD) Field Descriptions (continued)

| Offset | Bits | Name | Description |
|--------|------|------------------------|---|
| 0–1 | 14 | UN | Underrun. Written by the eTSEC. 0 No underrun encountered (data was retrieved from external memory in time to send a complete frame). 1 The Ethernet controller encountered a transmitter underrun condition while sending the associated buffer. This could also have occurred in relation to a bus error causing IEVENT[EBERR]. The eTSEC terminates the transmission and updates UN. |
| | | TOE | TCP/IP off-load enable. Written by user. Valid only if set in the first BD of a frame. 0 No TCP/IP off-load acceleration is applied to the frame prior to transmission. 1 eTSEC looks for a TOE Frame Control Block preceding the frame, and applies TCP/IP off-load acceleration as controlled by the FCB. |
| | 15 | TR | Truncation. Written by the eTSEC. Set in the last TxBD (TxBD[L] is set) when IEVENT[BABT] occurs for a frame (a frame length greater than or equal to the value set in the maximum frame length register is encountered, the HFE bit in the BD is cleared, and MACCFG2[Huge Frame] is cleared). The frame is sent truncated. |
| 2–3 | 0–15 | Data Length | Data length is the number of octets the eTSEC should transmit from this BD's data buffer. It is never modified by the eTSEC. This field must be greater than zero, as zero indicates a BD not ready. |
| 4–7 | 0–31 | TX Data Buffer Pointer | The transmit buffer pointer contains the address of the associated data buffer. The data buffer pointer for the first BD of a TxPAL-enabled frame must be aligned on an 8-byte boundary. There are no alignment restrictions for the data buffer pointers of the second or subsequent BDs of a TxPAL-enabled frame, or for non-TxPAL frames. |

19.6.7.3 Receive Buffer Descriptors (RxBD)

In the RxBD the user initializes the E, I, and W bits in the first word and the pointer in second word. If the data buffer is used, the eTSEC modifies the E, L, F, M, BC, MC, LG, NO, CR, OV, and TR bits and writes the length of the used portion of the buffer in the first word. The M, BC, MC, LG, NO, CR, OV, and TR bits in the first word of the buffer descriptor are only modified by the eTSEC if the L (last BD in frame) bit is set. The first word of the RxBD contains control and status bits. Its formats are detailed below.

The number of buffer descriptors in a ring is set using the W bit to indicate that the next buffer wraps back to the beginning of the ring. See [Section 19.5.3.5.5, “Maximum Frame Length Register \(MAXFRM\),”](#) for information on setting the size of the buffer ring.

Figure 19-150 defines the RxBD.

| | | | | | | | | | | | | | | | | |
|------------|------------------------|-----|---|---|---|---|---|---|----|----|----|----|----|----|----|----|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
| Offset + 0 | E | RO1 | W | I | L | F | 0 | M | BC | MC | LG | NO | SH | CR | OV | TR |
| Offset + 2 | DATA LENGTH | | | | | | | | | | | | | | | |
| Offset + 4 | RX DATA BUFFER POINTER | | | | | | | | | | | | | | | |
| Offset + 6 | | | | | | | | | | | | | | | | |

Figure 19-150. Receive Buffer Descriptor

The RxBD definition is interpreted by eTSEC hardware as if RxBDs mapped to C data structures in the manner illustrated by [Figure 19-151](#).

```

typedef unsigned short uint_16; /* choose 16-bit native type */
typedef unsigned int uint_32; /* choose 32-bit native type */
typedef struct rxbd_struct {
    uint_16 flags;
    uint_16 length;
    uint_32 bufptr;
} rxbd;

```

Figure 19-151. Mapping of RxBDs to a C Data Structure

Table 19-160 describes the fields of the RxBD.

Table 19-160. Receive Buffer Descriptor Field Descriptions

| Offset | Bits | Name | Description |
|--------|------|------|--|
| 0-1 | 0 | E | Empty, written by the eTSEC (when cleared) and by the user (when set). 0 The data buffer associated with this BD is filled with received data, or data reception is aborted due to an error condition. The status and length fields have been updated as required. 1 The data buffer associated with this BD is empty, or reception is currently in progress. |
| | 1 | RO1 | Receive software ownership bit. This field is reserved for use by software. This read/write bit is not modified by hardware, nor does its value affect hardware. |
| | 2 | W | Wrap, written by user. 0 The next buffer descriptor is found in the consecutive location. 1 The next buffer descriptor is found at the location defined in RBASE. |
| | 3 | I | Interrupt, written by user. 0 No interrupt is generated after this buffer is serviced. 1 IEVENT[RXB] or IEVENT[RXF] are set after this buffer is serviced. This bit can cause an interrupt if enabled (IMASK[RXBEN] or IMASK[RXFEN]). If the user wants to be interrupted only if RXF occurs, then the user must disable RXB (IMASK[RXBEN] is cleared) and enable RXF (IMASK[RXFEN] is set). |
| | 4 | L | Last in frame, written by the eTSEC. 0 The buffer is not the last in a frame. 1 The buffer is the last in a frame. |
| | 5 | F | First in frame, written by the eTSEC. 0 The buffer is not the first in a frame. 1 The buffer is the first in a frame. |
| | 6 | — | Reserved |
| | 7 | M | Miss, written by the eTSEC. (This bit is valid only if the L-bit is set and eTSEC is in promiscuous mode.) This bit is set by the eTSEC for frames that were accepted in promiscuous mode, but were flagged as a “miss” by the internal address recognition; thus, while in promiscuous mode, the user can use the M-bit to quickly determine whether the frame was destined to this station. 0 The frame was received because of an address recognition hit. 1 The frame was received because of promiscuous mode. |

Table 19-160. Receive Buffer Descriptor Field Descriptions (continued)

| Offset | Bits | Name | Description |
|--------|------|------------------------|--|
| 0–1 | 8 | BC | Broadcast. Written by the eTSEC. (Only valid if L is set.) Is set if the DA is broadcast (FF-FF-FF-FF-FF-FF). |
| | 9 | MC | Multicast. Written by the eTSEC. (Only valid if L is set.) Is set if the DA is multicast and not BC. |
| | 10 | LG | Rx frame length violation, written by the eTSEC (only valid if L is set). A frame length greater than or equal to the maximum frame length was recognized; in this case LG is set regardless of the setting of MACCFG2[Huge Frame]. If MACCFG2[Huge Frame] is cleared, the frame is truncated to the value programmed in the maximum frame length register. This bit is valid only if the L bit is set. |
| | 11 | NO | Rx non-octet aligned frame, written by the eTSEC (only valid if L is set). A frame that contained a number of bits not divisible by eight was received. |
| | 12 | SH | Short frame, written by the eTSEC (only valid if L is set). A frame length less than the minimum 64B that is defined for Ethernet. was recognized, provided RCTRL[RSF] is set. |
| | 13 | CR | Rx CRC error, written by the eTSEC (only valid if L is set). This frame contains a CRC error and is an integral number of octets in length. This bit is also set if a receive code group error is detected. |
| | 14 | OV | Overflow, written by the eTSEC (only valid if L is set). A receive FIFO overflow occurred during frame reception. If this bit is set, the other status bits, M, LG, NO, CR and TR lose their normal meaning and are zero. |
| | 15 | TR | Truncation, written by the eTSEC (only valid if L is set). Set if the receive frame is truncated. This can happen if a frame length greater than the maximum frame length is received and MACCFG2[Huge Frame] is cleared. If this bit is set, the frame must be discarded and the other error bits must be ignored as they may be incorrect. |
| 2–3 | 0–15 | Data Length | Data length, written by the eTSEC. Data length is the number of octets written by the eTSEC into this BD's data buffer if L is cleared (the value is equal to MRBLR), or, if L is set, the length of the frame including CRC, FCB (if RCTRL[PRSDEP > 00]), preamble (if MACCFG2[PreAmRxEn]=1), timestamp (if RCTRL[TS]=1) and any padding (RCTRL[PAL]). |
| 4–7 | 0–31 | RX Data Buffer Pointer | Receive buffer pointer, written by the user. The receive buffer pointer, which always points to the first location of the associated data buffer, must be 8-byte aligned. For best performance, use 64-byte aligned receive buffer pointer addresses. The buffer must reside in memory external to the eTSEC. |

19.7 Initialization/Application Information

19.7.1 Interface Mode Configuration

This section describes how to configure the eTSEC in different supported interface modes. These include the following:

- MII
- RMII

- RGMII
- SGMII
- RTBI

The pinout, the data registers that must be initialized, as well as speed selection options are described.

19.7.1.1 MII Interface Mode

Table 19-161 describes the signal configurations required for MII interface mode.

Table 19-161. MII Interface Mode Signal Configuration

| eTSEC Signals | | | MII Interface | | |
|---------------|-----|----------------|--------------------|-----|----------------|
| | | | Frequency [MHz] 25 | | |
| | | | Voltage [V] 3.3 | | |
| Signals | I/O | No. of Signals | Signals | I/O | No. of Signals |
| GTX_CLK | O | 1 | leave unconnected | | |
| TX_CLK | I | 1 | TX_CLK | I | 1 |
| TxD[0] | O | 1 | TxD[0] | O | 1 |
| TxD[1] | O | 1 | TxD[1] | O | 1 |
| TxD[2] | O | 1 | TxD[2] | O | 1 |
| TxD[3] | O | 1 | TxD[3] | O | 1 |
| TX_EN | O | 1 | TX_EN | O | 1 |
| TX_ER | O | 1 | TX_ER | O | 1 |
| RX_CLK | I | 1 | RX_CLK | I | 1 |
| RxD[0] | I | 1 | RxD[0] | I | 1 |
| RxD[1] | I | 1 | RxD[1] | I | 1 |
| RxD[2] | I | 1 | RxD[2] | I | 1 |
| RxD[3] | I | 1 | RxD[3] | I | 1 |
| RX_DV | I | 1 | RX_DV | I | 1 |
| RX_ER | I | 1 | RX_ER | I | 1 |
| COL | I | 1 | COL | I | 1 |
| CRS | I | 1 | CRS | I | 1 |
| Sum | | 17 | Sum | | 16 |

Table 19-162 describes the shared signals of the MII interface.

Table 19-162. Shared MII Signals

| eTSEC Signals | I/O | No. of Signals | MII Signals | I/O | No. of Signals | Function |
|---------------|-----|----------------|-------------|-----|----------------|----------------------------|
| MDIO | I/O | 1 | MDIO | I/O | 1 | Management interface I/O |
| MDC | O | 1 | MDC | O | 1 | Management interface clock |
| ECGTX_CLK125 | I | 1 | not used | I | 0 | Reference clock |
| Sum | | | Sum | | | |

Table 19-163 describes the register initializations required to configure the eTSEC in MII mode.

Table 19-163. MII Mode Register Initialization Steps

| |
|---|
| Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000] |
| Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize MACCFG2, for MII, half duplex operation. Set I/F Mode bit, MACCFG2[0000_0000_0000_0000_0111_0001_0000_0100] (This example has Full Duplex = 0, Preamble count = 7, PAD/CRC append = 1) |
| Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has Statistics Enable = 1) |
| Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] Set station address to 02_60_8C_87_65_43, for example. |
| Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] Set station address to 02_60_8C_87_65_43, for example. |
| Reset the management interface. MIIMCFG[1000_0000_0000_0000_0000_0000_0000_0111] |
| Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] set source clock divide by 14 for example to insure that MDC clock speed is not less than 2.5 MHz |
| Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle. |
| Set up the MII Mgmt for a write cycle to the external PHY Auxiliary Control and Status Register to configure the PHY through the Management interface (overrides configuration signals of the PHY). MIIMADD[0000_0000_0000_0000_0000_0000_0001_1100] |
| Perform an MII Mgmt write cycle to the external PHY Writing to MII Mgmt Control with 16-bit data intended for the external PHY register, MIIMCON[0000_0000_0000_0000_0000_0000_0000_0100] |

Table 19-163. MII Mode Register Initialization Steps (continued)

| |
|---|
| <p>Check to see if MII Mgmt write is complete Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>Set up the MII Mgmt for a write cycle to the external PHY Extended PHY control register #1 to set up the interface mode selection. MIIMADD[0000_0000_0000_0000_0000_0000_0001_0111]</p> |
| <p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY register, MIIMCON[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>Set up the MII Mgmt for a write cycle to the external PHY Mode control register to set up the interface mode selection. MIIMADD[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY register, MIIMCON[0000_0000_0000_0000_00uu_00uu_0u00_0000] where u is user defined based on desired configuration.</p> |
| <p>Check to see if MII Mgmt write is complete Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>If auto-negotiation was enabled in the PHY, check to see if PHY has completed Auto-Negotiation. Set up the MII Mgmt for a read cycle to PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0000_0000_0000_0001] The PHY Status register is at address 0x1 and in this case the PHY Address is 0x00.</p> |
| <p>Perform an MII Mgmt read cycle of Status Register. Clear MIIMCOM[Read Cycle]. Set MIIMCOM[Read Cycle]. (Uses the PHY address (0) and Register address (1) placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MIIMSTAT register and check bit 10 (AN Done and Link is up) MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0100] Other information about the link is also returned.(Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p> |
| <p>Check auto-negotiation attributes in the PHY as necessary.</p> |
| <p>Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize GADDR_n (Optional) GADDR_n[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p> |

Table 19-163. MII Mode Register Initialization Steps (continued)

| |
|---|
| Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0–TBASE7, TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000] |
| Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0–RBASE7, RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000] |
| Enable Transmit Queues Initialize TQUEUE |
| Enable Receive Queues Initialize RQUEUE |
| Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101] |

19.7.1.2 RGMII Interface Mode

Table 19-164 shows the signals configurations required for RGMII interface mode.

Table 19-164. RGMII Interface Mode Signal Configuration

| eTSEC Signals | | | RGMII Interface | | |
|---------------|-----|----------------|-----------------------|-----|----------------|
| | | | Frequency [MHz] 125 | | |
| | | | Voltage [V] 2.5 | | |
| Signals | I/O | No. of Signals | Signals | I/O | No. of Signals |
| GTX_CLK | O | 1 | GTX_CLK | O | 1 |
| TX_CLK | I | 1 | not used | | |
| TxD[0] | O | 1 | TxD[0]/TxD[4] | O | 1 |
| TxD[1] | O | 1 | TxD[1]/TxD[5] | O | 1 |
| TxD[2] | O | 1 | TxD[2]/TxD[6] | O | 1 |
| TxD[3] | O | 1 | TxD[3]/TxD[7] | O | 1 |
| TX_EN | O | 1 | TX_CTL (TX_EN/TX_ERR) | O | 1 |
| TX_ER | O | 1 | leave unconnected | | |
| RX_CLK | I | 1 | RX_CLK | I | 1 |
| RxD[0] | I | 1 | RxD[0]/RxD[4] | I | 1 |
| RxD[1] | I | 1 | RxD[1]/RxD[5] | I | 1 |
| RxD[2] | I | 1 | RxD[2]/RxD[6] | I | 1 |
| RxD[3] | I | 1 | RxD[3]/RxD[7] | I | 1 |
| RX_DV | I | 1 | RX_CTL (RX_DV/RX_ERR) | I | 1 |
| RX_ER | I | 1 | not used | | |
| COL | I | 1 | not used | | |

Table 19-164. RGMII Interface Mode Signal Configuration (continued)

| eTSEC Signals | | | RGMII Interface | | |
|---------------|-----|----------------|---------------------|-----|----------------|
| | | | Frequency [MHz] 125 | | |
| | | | Voltage [V] 2.5 | | |
| Signals | I/O | No. of Signals | Signals | I/O | No. of Signals |
| CRS | I | 1 | not used | | |
| Sum | | 17 | Sum | | 12 |

Table 19-165 describes the shared signals for the RGMII interface.

Table 19-165. Shared RGMII Signals

| eTSEC Signals | I/O | No. of Signals | GMII Signals | I/O | No. of Signals | Function |
|---------------|-----|----------------|--------------|-----|----------------|----------------------------|
| MDIO | I/O | 1 | MDIO | I/O | 1 | Management interface I/O |
| MDC | O | 1 | MDC | O | 1 | Management interface clock |
| GTX_CLK125 | I | 1 | GTX_CLK125 | I | 1 | Reference clock |
| Sum | | | Sum | | | |

Table 19-166 describes the register initializations required to configure the eTSEC in RGMII mode.

Table 19-166. RGMII Mode Register Initialization Steps

| |
|--|
| Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000] |
| Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000] |
| Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1) |
| Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has RGMII 10Mbps mode, Statistics Enable = 1) |
| Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example. |
| Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example. |
| Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] Set source clock divide by 14, for example, to insure that TSEC_MDC clock speed is not less than 2.5 MHz. |

Table 19-166. RGMII Mode Register Initialization Steps (continued)

| |
|---|
| <p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.</p> |
| <p>Set up the MII Mgmt for a write cycle to external the PHY AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0001_0000_0100] The AN Advertisement register is at offset address 0x04 from the external PHY address. (in this case 0x11)</p> |
| <p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY AN Advertisement register, MIIMCON[0000_0000_0000_0000_u0uu_uuuu_uuuu_uuuu] Where u must be selected by the user for proper system configuration.</p> |
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>Set up the MII Mgmt for a write cycle to the external PHY Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0001_0000_0000] The control register (CR) is at offset address 0x00 from the external PHY address. (in this case 0x11)</p> |
| <p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY Control register, MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000] This enables the external PHY to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p> |
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>Check to see if PHY has completed Auto-Negotiation. Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0000_0010_0000_0001] The PHY Status register is at address 0x1 and in this case the PHY Address is 0x2.</p> |
| <p>Perform an MII Mgmt read cycle of Status Register. Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (2) placed in MIIMADD register) When MIIMIND[BUSY]=0, read the MIIMSTAT register and check bit 10. (AN Done) MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0000] Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p> |
| <p>Perform an MII Mgmt read cycle of AN Expansion Register. Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0110] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x11) and Register address (6) placed in MIIMADD register) When MIIMIND[BUSY]=0, read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd) MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p> |

Table 19-166. RGMII Mode Register Initialization Steps (continued)

| |
|---|
| Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional) Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0101] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x11) and Register address (5) placed in MIIMADD register) When MIIMIND[BUSY]=0, read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex) MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_000x_1x10_0000] |
| Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize GADDR _n (Optional) GADDR _n [0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0–TBASE7, TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000] |
| Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0–RBASE7, RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000] |
| Enable Transmit Queues Initialize TQUEUE |
| Enable Receive Queues Initialize RQUEUE |
| Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101] |

19.7.1.3 RMI Interface Mode

Table 19-167 shows the signals configurations required for RMI interface mode.

Table 19-167. RMI Interface Mode Signal Configuration

| eTSEC Signals | | | RMI Interface | | |
|---------------|-----|----------------|--------------------|-----|----------------|
| | | | Frequency [MHz] 50 | | |
| | | | Voltage [V] 3.3 | | |
| Signals | I/O | No. of Signals | Signals | I/O | No. of Signals |
| GTX_CLK | O | 1 | leave unconnected | | |
| TX_CLK | I | 1 | REF_CLK | I | 1 |
| TxD[0] | O | 1 | TxD[0] | O | 1 |
| TxD[1] | O | 1 | TxD[1] | O | 1 |
| TxD[2] | O | 1 | leave unconnected | | |
| TxD[3] | O | 1 | leave unconnected | | |
| TX_EN | O | 1 | TX_EN | O | 1 |
| TX_ER | O | 1 | leave unconnected | | |
| RX_CLK | I | 1 | leave unconnected | | |
| RxD[0] | I | 1 | RxD[0] | I | 1 |
| RxD[1] | I | 1 | RxD[1] | I | 1 |
| RxD[2] | I | 1 | not used | | |
| RxD[3] | I | 1 | not used | | |
| RX_DV | I | 1 | CRS_DV | I | 1 |
| RX_ER | I | 1 | RX_ER | I | 1 |
| COL | I | 1 | not used | | |
| CRS | I | 1 | not used | | |
| Sum | | 17 | Sum | | 8 |

Table 19-168 describes the shared signals for the RMI interface.

Table 19-168. Shared RMI Signals

| eTSEC Signals | I/O | No. of Signals | RMI Signals | I/O | No. of Signals | Function |
|---------------|-----|----------------|-------------|-----|----------------|----------------------------|
| MDIO | I/O | 1 | MDIO | I/O | 1 | Management interface I/O |
| MDC | O | 1 | MDC | O | 1 | Management interface clock |
| TX_CLK | I | 1 | REF_CLK | I | 1 | Reference clock |
| Sum | | 3 | Sum | | 3 | |

Table 19-169 describes the register initializations required to configure the eTSEC in RMII mode.

Table 19-169. RMII Mode Register Initialization Steps

| |
|--|
| <p>Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0001_0000_0101] (I/F Mode = 1, Full Duplex = 1)</p> |
| <p>Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0001_0000] (Used to setup Reduced-Pin mode = 1, and TBIM = 0,statistics enable = 1)</p> |
| <p>Initialize MAC Station Address MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543 for example</p> |
| <p>Initialize MAC Station Address MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543 for example</p> |
| <p>Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_1101] set system clock divide by 14 for example to insure that MDC clock speed = < 2.5 MHz</p> |
| <p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.</p> |
| <p>Set up the MII Mgmt for a write cycle to external the PHY AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0001_0000_0100] The AN Advertisement register is at offset address 0x04 from the external PHY address. (in this case 0x11)</p> |
| <p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY AN Advertisement register, MIIMCON[0000_0000_0000_0000_u0uu_uuuu_uuuu_uuuu] Where u must be selected by the user for proper system configuration.</p> |
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>Set up the MII Mgmt for a write cycle to the external PHY Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0001_0000_0000] The control register is at offset address 0x00 from the external PHY address. (in this case 0x11)</p> |
| <p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY Control register, MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000] This enables the external PHY to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p> |
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |

Table 19-169. RMII Mode Register Initialization Steps (continued)

| |
|---|
| <p>Check to see if PHY has completed Auto-Negotiation. Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0000_0010_0000_0001] The PHY Status register is at address 0x1 and in this case the PHY Address is 0x2.</p> |
| <p>Perform an MII Mgmt read cycle of Status Register. Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (1) placed in MIIMADD register) When MIIMIND[BUSY]=0, read the MIIMSTAT register and check bit 10. (AN Done) MIIMSTAT ---> [0000_0000_0000_0000_0000_0010_0000] Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p> |
| <p>Perform an MII Mgmt read cycle of AN Expansion Register. Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0110] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x11) and Register address (6) placed in MIIMADD register) When MIIMIND[BUSY]=0, read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd) MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p> |
| <p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional) Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0101] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x11) and Register address (5) placed in MIIMADD register) When MIIMIND[BUSY]=0, read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex) MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_000x_x110_0000]</p> |
| <p>Setting up the MII Mgmt for a write cycle to TBI MII Mgmt register (write the TBI's address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_1011] the TBI control register is at offset address 0x11 from TBIPA</p> |
| <p>Perform an MII Mgmt write cycle Writing to MII Mgmt Control with 16-bit data intended for TBI's MII Mgmt control register (TBI control), MIIMCON[0000_0000_0000_0000_0000_0010_0001_0000] This configures the TBI control to GMII mode and AN sense</p> |
| <p>Check to see if MII Mgmt write is complete Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicate that the write cycle was completed</p> |
| <p>Perform an MII Mgmt read cycle (Optional) Set MIIMCOM[Read Cycle] (Uses the TBI address and Register address placed in MIIMADD register), read the MIIMSTAT register and verify that MIIMSTAT ---> [0000_0000_0000_0000_0000_0010_0001_0000]</p> |
| <p>Check to see if PHY has completed Auto-Negotiation Setting up the MII Mgmt for a read cycle to PHY's MII Mgmt register (write the PHY's address and Register address), MIIMADD[0000_0000_0000_0000_0000_0010_0000_0010] the PHY Status control register is at address 0x2 and lets say the PHY Address is 0x2</p> |

Table 19-169. RMII Mode Register Initialization Steps (continued)

| |
|---|
| <p>Perform an MII Mgmt read cycle of Status Register Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (2) placed in MIIMADD register), read the MIIMSTAT register and check bit 10 (AN Done) MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0000] other information about the link is also returned (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p> |
| <p>Perform an MII Mgmt read cycle of AN Expansion Register MIIMADD[0000_0000_0000_0000_0000_0010_0000_0110] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (6) placed in MIIMADD register), read the MII Mgmt AN Expansion register and check bits 13 and 14 (NP Able and Page Rx'd) MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p> |
| <p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register (Optional) MIIMADD[0000_0000_0000_0000_0000_0010_0000_0101] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (5) placed in MIIMADD register), read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10 (Half and Full Duplex) MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_000X_1110_0000]</p> |
| <p>Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize GADDR_n (Optional) GADDR_n[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0–TBASE7, TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p> |
| <p>Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0–RBASE7, RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p> |
| <p>Enable Transmit Queues Initialize TQUEUE</p> |
| <p>Enable Receive Queues Initialize RQUEUE</p> |
| <p>Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]</p> |

19.7.1.4 RTBI Interface Mode

Table 19-170 describes the signal configurations required for RTBI interface mode.

Table 19-170. RTBI Interface Mode Signal Configuration

| eTSEC Signal s | | | RTBI Interface | | |
|----------------|-----|----------------|---------------------|-----|----------------|
| | | | Frequency [MHz] 125 | | |
| | | | Voltage [V] 2.5 | | |
| Signals | I/O | No. of Signals | Signals | I/O | No. of Signals |
| GTX_CLK | O | 1 | GTX_CLK | O | 1 |
| TX_CLK | I | 1 | not used | | |
| TxD[0] | O | 1 | TCG[0]/TCG[5] | O | 1 |
| TxD[1] | O | 1 | TCG[1]/TCG[6] | O | 1 |
| TxD[2] | O | 1 | TCG[2]/TCG[7] | O | 1 |
| TxD[3] | O | 1 | TCG[3]/TCG[8] | O | 1 |
| TX_EN | O | 1 | TCG[4]/TCG[9] | O | 1 |
| TX_ER | O | 1 | leave unconnected | | |
| RX_CLK | I | 1 | RX_CLK | I | 1 |
| RxD[0] | I | 1 | RCG[0]/RCG[5] | I | 1 |
| RxD[1] | I | 1 | RCG[1]/RCG[6] | I | 1 |
| RxD[2] | I | 1 | RCG[2]/RCG[7] | I | 1 |
| RxD[3] | I | 1 | RCG[3]/RCG[8] | I | 1 |
| RX_DV | I | 1 | RCG[4]/RCG[9] | I | 1 |
| RX_ER | I | 1 | not used | | |
| COL | I | 1 | not used | | |
| CRS | I | 1 | not used | | |
| Sum | | 17 | sum | | 12 |

Table 19-171 describes the shared signals for the RTBI interface.

Table 19-171. Shared RTBI Signals

| eTSEC Signals | I/O | No. of Signals | GMII Signals | I/O | No. of Signals | Function |
|---------------|-----|----------------|--------------|-----|----------------|----------------------------|
| MDIO | I/O | 1 | MDIO | I/O | 1 | Management interface I/O |
| MDC | O | 1 | MDC | O | 1 | Management interface clock |
| ECGTX_CLK125 | I | 1 | GTX_CLK125 | I | 1 | Reference clock |
| Sum | | 3 | Sum | | 3 | |

Table 19-172 describes the register initializations required to configure the eTSEC in RTBI mode.

Table 19-172. RTBI Mode Register Initialization Steps

| |
|--|
| Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000] |
| Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000] |
| Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1) |
| Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has Statistics Enable = 1) |
| Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example. |
| Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example. |
| Assign a Physical address to the TBI, TBIPA[0000_0000_0000_0000_0000_0000_0001_0000] set to 16, for example. |
| Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] Set source clock divide by 14, for example, to insure that TSEC_MDC clock speed is not less than 2.5 MHz. |
| Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle. |
| Set up the MII Mgmt for a read cycle to TBI's Control register (write the TBI's address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] The control register (CR) is at offset address 0x0 from TBIPA. |
| Perform an MII Mgmt read cycle to verify state of TBI Control Register(Optional) Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the TBI address and Register address placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MIIMSTAT and look for AN Enable and other bit information. |
| Set up the MII Mgmt for a write cycle to TBI's AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0100] The AN Advertisement register is at offset address 0x04 from the TBI's address. (in this case 0x10) |
| Perform an MII Mgmt write cycle to TBI. Write to MII Mgmt Control with 16-bit data intended for TBI's AN Advertisement register, MIIMCON[0000_0000_0000_0000_0000_0001_1010_0000] This advertises to the Link Partner that the TBI supports PAUSE and Full Duplex mode and does not support Half Duplex mode. |
| Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed. |

Table 19-172. RTBI Mode Register Initialization Steps (continued)

| |
|--|
| <p>Set up the MII Mgmt for a write cycle to TBI's Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] The control register (CR) is at offset address 0x00 from the TBI's address. (in this case 0x10)</p> |
| <p>Perform an MII Mgmt write cycle to TBI. Writing to MII Mgmt Control with 16-bit data intended for TBI's Control register, MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000] This enables the TBI to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p> |
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>Check to see if PHY has completed Auto-Negotiation. Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0001] The PHY Status control register is at address 0x1 and in this case the PHY Address is 0x10.</p> |
| <p>Perform an MII Mgmt read cycle of Status Register. Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (2) placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MIIMSTAT register and check bit 10 (AN Done) MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0000] Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p> |
| <p>Perform an MII Mgmt read cycle of AN Expansion Register. Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0110] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x10) and Register address (6) placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd) MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p> |
| <p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional) Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0101] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x10) and Register address (5) placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex) MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_0000_00x_x110_0000]</p> |
| <p>Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize GADDR_n (Optional) GADDR_n[0000_0000_0000_0000_0000_0000_0000_0000]</p> |

Table 19-172. RTBI Mode Register Initialization Steps (continued)

| |
|---|
| Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000] |
| Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0–TBASE7, TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000] |
| Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0–RBASE7, RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000] |
| Enable Transmit Queues Initialize TQUEUE |
| Enable Receive Queues Initialize RQUEUE |
| Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101] |

19.7.1.5 SGMII Interface Support

Table 19-173. SGMII Interface Signal Configuration (4-Wire)

| SerDes Signals | | | SGMII Interface | | |
|------------------------|-----|----------------|----------------------|-----|----------------|
| Frequency [MHz] 1250 | | | Frequency [MHz] 1250 | | |
| Voltage [V] LVDS | | | Voltage [V] LVDS | | |
| Signals | I/O | No. of Signals | Signals | I/O | No. of Signals |
| $TX_n/\overline{TX_n}$ | O | 2 | TXD | O | 2 |
| $RX_n/\overline{RX_n}$ | I | 2 | RXD | I | 2 |
| Sum | | 4 | Sum | | 4 |

SGMII mode initialization sequence requires additional initialization for the SerDes. An example of SGMII mode initialization sequence is shown in [Table 19-174](#).

NOTE

SGMII mode utilizes the internal TBI PHY. The internal TBI PHY only auto-negotiates at 1 Gbps. However, 10 Mbps and 100 Mbps speeds are supported in SGMII mode. It is recommended that the external PHY inform the MAC if the desired link speed is not 1 Gbps. Software can perform MII management cycles to determine the external PHY link speed and program ECNTRL and MACCFG2 accordingly.

Table 19-174. SGMII Mode Register Initialization Steps

| |
|--|
| <p><i>Initialize SerDes to select SGMII. The initialization sequence should be prepended with SerDes initialization.</i></p> |
| <p>Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1) (Set I/F mode = 1 in SGMII 10/100 Mbps speed)</p> |
| <p>Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0010_0010] (This example has Statistics Enable = 1, TBIM = 1, SGMII = 1) (Set R100M = 1 in SGMII 100 Mbps speed)</p> |
| <p>Initialize MAC Station Address MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example.</p> |
| <p>Initialize MAC Station Address MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example.</p> |
| <p>Assign a Physical address to the TBI, TBIPA[0000_0000_0000_0000_0000_0000_0001_0000] set to 16, for example.</p> |
| <p>Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_000_0000_0000_0101] set source clock divide by 14 for example to insure that MDC clock speed is not less than 2.5 MHz</p> |
| <p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.</p> |
| <p>Set up the MII Mgmt for a read cycle to TBI's Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] the control register (CR) is at offset address 0x00 from the TBI's address.</p> |
| <p>Perform an MII Mgmt read cycle to verify state of TBI Control Register (optional) Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the TBI address and Register address placed in MIIMADD register), When MIIMIND[BUSY] = 0, read the MIIMSTAT and look for AN Enable and other bit information.</p> |
| <p>Set up the MII Mgmt for a write cycle to TBICON register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0001_0001] The TBICON register is at offset address 0x11 from the TBI's address.</p> |
| <p>Perform an MII Mgmt write cycle to TBI. Writing to MII Mgmt Control with 16-bit data intended for TBICON register, MIIMCON[0000_0000_0000_0000_0000_0000_0010_0000] This sets TBI in single clock mode and MII Mode off to enable communication with SerDes.</p> |

Table 19-174. SGMII Mode Register Initialization Steps (continued)

| |
|--|
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>Set up the MII Mgmt for a write cycle to TBI's AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0100] The AN Advertisement register is at offset address 0x04 from the TBI's address.</p> |
| <p>Perform an MII Mgmt write cycle to TBI. Writing to MII Mgmt Control with 16-bit data intended for TBI's AN Advertisement register, MIIMCON[0000_0000_0000_0000_0000_0001_1010_0000] This advertises to the Link Partner that the TBI supports PAUSE and Full Duplex mode and does not support Half Duplex mode.</p> |
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p><i>Additional SerDes setup as required</i></p> |
| <p>Set up the MII Mgmt for a write cycle to TBI's Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] the control register (CR) is at offset address 0x00 from the TBI's address.</p> |
| <p>Perform an MII Mgmt write cycle to TBI. Writing to MII Mgmt Control with 16-bit data intended for TBI's Control register, MIIMCON[0000_0000_0000_0000_0001_0011_0100_0000] This enables the TBI to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p> |
| <p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p> |
| <p>Check to see if PHY has completed Auto-Negotiation. Set up the MII Mgmt for a read cycle to PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0001] The PHY Status control register is at address 0x1 and in this case the PHY Address is 0x10.</p> |
| <p>Perform an MII Mgmt read cycle of Status Register. Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (2) and Register address (2) placed in MIIMADD register), When MIIMIND[BUSY] = 0, read the MIIMSTAT register and check bit 10 (AN Done) MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110] Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p> |
| <p>Perform an MII Mgmt read cycle of AN Expansion Register. Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0110] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x10) and Register address (6) placed in MIIMADD register), When MIIMIND[BUSY] = 0, read the MII Mgmt AN Expansion register and check bits 13 and 14 (NP Able and Page Rx'd) MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p> |

Table 19-174. SGMII Mode Register Initialization Steps (continued)

| |
|--|
| <p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional) Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0101] Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the PHY address (0x10) and Register address (5) placed in MIIMADD register), When MIIMIND[BUSY] = 0, read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex) MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_000x_1110_0000]</p> |
| <p>Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize GADDR_n (Optional) GADDR_n[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p> |
| <p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0–TBASE7, TBASE0–TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p> |
| <p>Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0–RBASE7, RBASE0–RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p> |
| <p>Enable Transmit Queues Initialize TQUEUE</p> |
| <p>Enable Receive Queues Initialize RQUEUE</p> |
| <p>Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]</p> |

Chapter 20 I²C Interface

This chapter describes the inter-IC (IIC or I²C) bus interface implemented on this device.

20.1 I²C Introduction

The inter-IC (IIC or I²C) bus is a two-wire—serial data (SDA) and serial clock (SCL)—bidirectional serial bus that provides a simple, efficient method of data exchange between this device and other devices, such as microcontrollers, EEPROMs, real-time clock devices, A/D converters, and LCDs. Figure 20-1 shows a block diagram of the I²C interface.




Figure 20-1. I²C Block Diagram

The two-wire I²C bus minimizes interconnections between devices. The synchronous, multiple-master I²C bus allows the connection of additional devices to the bus for expansion and system development. The bus includes collision detection and arbitration that prevent data corruption if two or more masters attempt to control the bus simultaneously.

20.1.1 I²C Features

The I²C interface includes the following features:

- Two-wire interface
- Multiple-master operational
- Arbitration lost interrupt with automatic mode switching from master to slave
- Calling address identification interrupt
- START and STOP signal generation/detection
- Acknowledge bit generation/detection
- Bus busy detection
- Software-programmable clock frequency
- Software-selectable acknowledge bit
- On-chip filtering for spikes on the bus

20.1.2 I²C Modes of Operation

The I²C unit on this device can operate in one of the following modes:

- Master mode. The I²C initiates a transfer, generates clock signals, and terminates a transfer. It cannot use its own slave address as a calling address. The I²C cannot be a master and a slave simultaneously.
- Slave mode. The I²C is addressed by an I²C master. The module must be enabled before a START condition from an I²C master is detected.
- Interrupt-driven byte-to-byte data transfer. When successful slave addressing is achieved (and SCL returns to zero), the data transfer can proceed on a byte-to-byte basis in the direction specified by the R/\overline{W} bit sent by the calling master. Each byte of data must be followed by an acknowledge bit, which is signaled from the receiving device. Several bytes can be transferred during a data transfer session.
- Boot sequencer mode. I²C controller supports boot sequencer mode. This mode can be used to initialize the configuration registers in the device after the I²C module is initialized. Boot sequencer mode is selected using the BOOTSEQ field in the reset configuration word high. Note that the hard-coded reset configuration word high value is boot sequencer mode disabled.
- Reset configuration load. In this mode, the I²C interface loads the reset configuration words from an EEPROM at a specific calling address while the rest of the device is in the reset state ($\overline{\text{HRESET}}$ asserted). Once the reset configuration words are latched inside the device, I²C is reset until $\overline{\text{HRESET}}$ is negated. After $\overline{\text{HRESET}}$ is negated, the device may be initialized using boot sequence mode according to the BOOTSEQ field in the reset configuration word. See [Section 20.4.5, “Boot Sequencer Mode.”](#)

Additionally, the following three I²C-specific states are defined for the I²C interface:

- START condition. This condition denotes the beginning of a new data transfer (each data transfer contains several bytes of data) and awakens all slaves.

- Repeated START condition. A START condition that is generated without a STOP condition to terminate the previous transfer.
- STOP condition. The master can terminate the transfer by generating a STOP condition to free the bus.

20.2 I²C External Signal Descriptions

The following sections give an overview of signals and provide detailed signal descriptions.

20.2.1 I²C Signal Overview

The I²C interface uses the SDA and SCL signals, described in [Table 20-1](#), for data transfer. Note that the signal patterns driven on SDA represent address, data, or read/write information at different stages of the protocol.

Table 20-1. I²C Interface Signal Descriptions

| Signal Name | Idle State | I/O | State Meaning |
|--------------------|------------|-----|---|
| Serial Clock (SCL) | High | I | When the I ² C module is idle or acts as a slave, SCL defaults as an input. The unit uses SCL to synchronize incoming data on SDA. The bus is assumed to be busy when SCL is detected low. |
| | | O | As a master, the I ² C module drives SCL along with SDA when transmitting. As a slave, the I ² C module drives SCL negates for data pacing. |
| Serial Data (SDA) | High | I | When the I ² C module is idle or in a receiving mode, SDA defaults as an input. The unit receives data from other I ² C devices on SDA. The bus is assumed to be busy when SDA is detected low. |
| | | O | When writing as a master or slave, the I ² C module drives data on SDA synchronous to SCL. |

20.2.2 I²C Detailed Signal Descriptions

SDA and SCL, described in [Table 20-2](#), serve as a communication interconnect with other devices. All devices connected to these signals must have open-drain or open-collector outputs. The logic AND function is performed on both of these signals with external pull-up resistors. Refer to the hardware specifications for electrical characteristics.

Table 20-2. I²C Interface Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|----------------------|-----|--|
| SCL | I/O | Serial clock. Performs as an input when the device is programmed as an I ² C slave. SCL also performs as an output when the device is programmed as an I ² C master. |
| | O | As outputs for the bidirectional serial clock, these signals operate as described below. |
| | | State Meaning |
| | I | As inputs for the bi-directional serial clock, these signals operate as described below. |
| State Meaning | | Asserted/Negated—The I ² C unit uses this signal to synchronize incoming data on SDA. The bus is assumed to be busy when this signal is detected low. |

Table 20-2. I²C Interface Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|----------------------|-----|---|
| SDA | I/O | Serial data. Performs as an input when the device is in a receiving mode. SDA also performs as an output signal when the device is transmitting (as an I ² C master or a slave). |
| | O | As outputs for the bi-directional serial data, these signals operate as described below. |
| | | State Meaning |
| | I | As inputs for the bi-directional serial data, these signals operate as described below. |
| State Meaning | | Asserted/Negated—Used to receive data from other devices. The bus is assumed to be busy when SDA is detected low. |

20.3 I²C Memory Map/Register Definition

Table 20-3 lists the I²C-specific registers and their addresses.

Table 20-3. I²C Memory Map

| Address | I ² C Register | Access | Reset | Section/Page |
|--|--|--------|-------|-------------------------------|
| I²C Controller—Block Base Address 0x0_3000 | | | | |
| 0x0_3000 | I2CADR—I ² C address register | R/W | 0x00 | 20.3.1.1/20-4 |
| 0x0_3004 | I2CFDR—I ² C frequency divider register | R/W | 0x00 | 20.3.1.2/20-5 |
| 0x0_3008 | I2CCR—I ² C control register | R/W | 0x00 | 20.3.1.3/20-6 |
| 0x0_300C | I2CSR—I ² C status register | R/W | 0x81 | 20.3.1.4/20-7 |
| 0x0_3010 | I2CDR—I ² C data register | R/W | 0x00 | 20.3.1.5/20-9 |
| 0x0_3014 | I2CDFSR—I ² C digital filter sampling rate register | R/W | 0x10 | 20.3.1.6/20-9 |
| 0x0_301C– 0x0_31FF | Reserved, should be cleared | — | — | — |

20.3.1 I²C Register Descriptions

This section describes the I²C registers in detail. Note that reserved bits should always be written with the value they return when read. That is, the register should be programmed by reading the value, modifying appropriate fields, and writing back the value. The return value of the reserved fields should not be assumed, even though the reserved fields return zero. This does not apply to the I²C data register (I2CDR).

20.3.1.1 I²C Address Register (I2CADR)

Figure 20-2 shows the I2CADR register, which contains the address to which the I²C interface responds when addressed as a slave. Note that this is not the address that is sent on the bus during the address-calling cycle when the I²C module is in master mode.




Figure 20-2. I²C Address Register (I2CADR)

Table 20-4 describes the bit settings of I2CADR.

Table 20-4. I2CADR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–6 | ADDR | Slave address. Contains the specific slave address that is used by the I ² C interface. Note that the default mode of the I ² C interface is slave mode for an address match. Note that an address match is one of the conditions that can cause I2CSR[MIF] to be set, signaling an interrupt pending condition. |
| 7 | — | Reserved, should be cleared |

20.3.1.2 I²C Frequency Divider Register (I2CFDR)

Figure 20-3 shows the bits of the I²C frequency divider register.




Figure 20-3. I²C Frequency Divider Register (I2CFDR)

Table 20-5 describes the bit settings of I2CFDR. It also maps I2CFDR[FDR] to the clock divider values. Although it describes the ratio between the I²C controller internal clock and SCL, the default ratio of I²C controller clock and CSB is 1:3 (I²C controller clock frequency is three times slower than CSB clock frequency). This ratio is set in SCCR[ENCCM].

Table 20-5. I2C FDR Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|------|-------------------|--|-------------------|-------------------|-------------------|-------------------|-----|-------------------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|------|------|-------|------|------|------|------|------|-------|------|------|------|------|------|-------|------|------|------|------|------|-----|------|------|------|------|------|-----|------|------|------|------|------|-----|------|------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|------|------|-----|------|-------|------|-------|--|--|--|--|
| 0–1 | — | Reserved, should be cleared | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2–7 | FDR | <p>Frequency divider ratio. Used to prescale the clock for bit-rate selection. The serial bit clock frequency of SCL is equal to the I²C controller clock divided by the divider. The serial bit clock frequency divider selections are described as follows:</p> <table border="1"> <thead> <tr> <th>FDR</th> <th>Divider (Decimal)</th> <th>FDR</th> <th>Divider (Decimal)</th> <th>FDR</th> <th>Divider (Decimal)</th> </tr> </thead> <tbody> <tr><td>0x00</td><td>384</td><td>0x16</td><td>12288</td><td>0x2B</td><td>1024</td></tr> <tr><td>0x01</td><td>416</td><td>0x17</td><td>15360</td><td>0x2C</td><td>1280</td></tr> <tr><td>0x02</td><td>480</td><td>0x18</td><td>18432</td><td>0x2D</td><td>1536</td></tr> <tr><td>0x03</td><td>576</td><td>0x19</td><td>20480</td><td>0x2E</td><td>1792</td></tr> <tr><td>0x04</td><td>640</td><td>0x1A</td><td>24576</td><td>0x2F</td><td>2048</td></tr> <tr><td>0x05</td><td>704</td><td>0x1B</td><td>30720</td><td>0x30</td><td>2560</td></tr> <tr><td>0x06</td><td>832</td><td>0x1C</td><td>36864</td><td>0x31</td><td>3072</td></tr> <tr><td>0x07</td><td>1024</td><td>0x1D</td><td>40960</td><td>0x32</td><td>3584</td></tr> <tr><td>0x08</td><td>1152</td><td>0x1E</td><td>49152</td><td>0x33</td><td>4096</td></tr> <tr><td>0x09</td><td>1280</td><td>0x1F</td><td>61440</td><td>0x34</td><td>5120</td></tr> <tr><td>0x0A</td><td>1536</td><td>0x20</td><td>256</td><td>0x35</td><td>6144</td></tr> <tr><td>0x0B</td><td>1920</td><td>0x21</td><td>288</td><td>0x36</td><td>7168</td></tr> <tr><td>0x0C</td><td>2304</td><td>0x22</td><td>320</td><td>0x37</td><td>8192</td></tr> <tr><td>0x0D</td><td>2560</td><td>0x23</td><td>352</td><td>0x38</td><td>10240</td></tr> <tr><td>0x0E</td><td>3072</td><td>0x24</td><td>384</td><td>0x39</td><td>12288</td></tr> <tr><td>0x0F</td><td>3840</td><td>0x25</td><td>448</td><td>0x3A</td><td>14336</td></tr> <tr><td>0x10</td><td>4608</td><td>0x26</td><td>512</td><td>0x3B</td><td>16384</td></tr> <tr><td>0x11</td><td>5120</td><td>0x27</td><td>576</td><td>0x3C</td><td>20480</td></tr> <tr><td>0x12</td><td>6144</td><td>0x28</td><td>640</td><td>0x3D</td><td>24576</td></tr> <tr><td>0x13</td><td>7680</td><td>0x29</td><td>768</td><td>0x3E</td><td>28672</td></tr> <tr><td>0x14</td><td>9216</td><td>0x2A</td><td>896</td><td>0x3F</td><td>32768</td></tr> <tr><td>0x15</td><td>10240</td><td></td><td></td><td></td><td></td></tr> </tbody> </table> <p>Note: The value's shown in the table are applicable only for the default value of DFSRR. Refer to AN2919.</p> | FDR | Divider (Decimal) | FDR | Divider (Decimal) | FDR | Divider (Decimal) | 0x00 | 384 | 0x16 | 12288 | 0x2B | 1024 | 0x01 | 416 | 0x17 | 15360 | 0x2C | 1280 | 0x02 | 480 | 0x18 | 18432 | 0x2D | 1536 | 0x03 | 576 | 0x19 | 20480 | 0x2E | 1792 | 0x04 | 640 | 0x1A | 24576 | 0x2F | 2048 | 0x05 | 704 | 0x1B | 30720 | 0x30 | 2560 | 0x06 | 832 | 0x1C | 36864 | 0x31 | 3072 | 0x07 | 1024 | 0x1D | 40960 | 0x32 | 3584 | 0x08 | 1152 | 0x1E | 49152 | 0x33 | 4096 | 0x09 | 1280 | 0x1F | 61440 | 0x34 | 5120 | 0x0A | 1536 | 0x20 | 256 | 0x35 | 6144 | 0x0B | 1920 | 0x21 | 288 | 0x36 | 7168 | 0x0C | 2304 | 0x22 | 320 | 0x37 | 8192 | 0x0D | 2560 | 0x23 | 352 | 0x38 | 10240 | 0x0E | 3072 | 0x24 | 384 | 0x39 | 12288 | 0x0F | 3840 | 0x25 | 448 | 0x3A | 14336 | 0x10 | 4608 | 0x26 | 512 | 0x3B | 16384 | 0x11 | 5120 | 0x27 | 576 | 0x3C | 20480 | 0x12 | 6144 | 0x28 | 640 | 0x3D | 24576 | 0x13 | 7680 | 0x29 | 768 | 0x3E | 28672 | 0x14 | 9216 | 0x2A | 896 | 0x3F | 32768 | 0x15 | 10240 | | | | |
| FDR | Divider (Decimal) | FDR | Divider (Decimal) | FDR | Divider (Decimal) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x00 | 384 | 0x16 | 12288 | 0x2B | 1024 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x01 | 416 | 0x17 | 15360 | 0x2C | 1280 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x02 | 480 | 0x18 | 18432 | 0x2D | 1536 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x03 | 576 | 0x19 | 20480 | 0x2E | 1792 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x04 | 640 | 0x1A | 24576 | 0x2F | 2048 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x05 | 704 | 0x1B | 30720 | 0x30 | 2560 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x06 | 832 | 0x1C | 36864 | 0x31 | 3072 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x07 | 1024 | 0x1D | 40960 | 0x32 | 3584 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x08 | 1152 | 0x1E | 49152 | 0x33 | 4096 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x09 | 1280 | 0x1F | 61440 | 0x34 | 5120 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x0A | 1536 | 0x20 | 256 | 0x35 | 6144 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x0B | 1920 | 0x21 | 288 | 0x36 | 7168 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x0C | 2304 | 0x22 | 320 | 0x37 | 8192 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x0D | 2560 | 0x23 | 352 | 0x38 | 10240 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x0E | 3072 | 0x24 | 384 | 0x39 | 12288 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x0F | 3840 | 0x25 | 448 | 0x3A | 14336 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x10 | 4608 | 0x26 | 512 | 0x3B | 16384 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x11 | 5120 | 0x27 | 576 | 0x3C | 20480 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x12 | 6144 | 0x28 | 640 | 0x3D | 24576 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x13 | 7680 | 0x29 | 768 | 0x3E | 28672 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x14 | 9216 | 0x2A | 896 | 0x3F | 32768 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0x15 | 10240 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

20.3.1.3 I²C Control Register (I2CCR)

Figure 20-4 shows the I²C control register.




Figure 20-4. I²C Control Register (I2CCR)

Table 20-6 describes the I2CCR bit settings.

Table 20-6. I2CCR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0 | MEN | Module enable. Controls the software reset of the I ² C module. 0 The module is reset and disabled. The interface is held in reset, but the registers can still be accessed. 1 The I ² C module is enabled. MEN must be set before any other control register bits have any effect. All I ² C registers for slave receive or master START can be initialized before setting this bit. |
| 1 | MIEN | Module interrupt enable 0 Interrupts from the I ² C module are disabled. This does not clear any pending interrupt conditions. 1 Interrupts from the I ² C module are enabled. An interrupt occurs provided I2CSR[MIF] is also set. |
| 2 | MSTA | Master/slave mode START 0 On a transition to zero, a STOP condition is generated and the mode changes from master to slave. Cleared without generating a STOP condition when the master loses arbitration. 1 When MSTA changes from zero to one, a START condition is generated on the bus and master mode is selected. |
| 3 | MTX | Transmit/receive mode select. Selects the direction of the master and slave transfers. When configured as a slave, this bit should be set by software according to I2CSR[SRW]. In master mode, the bit should be set according to the type of transfer required. Therefore, for address cycles, this bit will always high. MTX is cleared when the master loses arbitration. 0 Receive mode 1 Transmit mode |
| 4 | TXAK | Transfer acknowledge. Specifies the value driven onto the SDA line during acknowledge cycles for both master and slave receivers. The value of this bit applies only when the I ² C module is configured as a receiver, not a transmitter. It also does not apply to address cycles; when the device is addressed as a slave, an acknowledge is always sent. 0 An acknowledge signal (low value on SDA) is sent out to the bus at the 9th clock bit after receiving one byte of data. 1 No acknowledge signal response (high value on SDA) is sent. |
| 5 | RSTA | Repeated START. Note that this bit is not readable, which means if a read is performed to RSTA, a zero value is returned. 0 No START condition is generated 1 Setting this bit always generates a repeated START condition on the bus, provides the device with the current bus master. Attempting a repeated START at the wrong time (or if the bus is owned by another master), results in loss of arbitration. |
| 6 | — | Reserved, should be cleared |
| 7 | BCST | Broadcast 0 Disables the broadcast accept capability 1 Enables the I ² C to accept broadcast messages at address zero |

20.3.1.4 I²C Status Register (I2CSR)

I2CSR is shown in Figure 20-5.




Figure 20-5. I²C Status Register (I2CSR)

Table 20-7 describes the bit settings of the I2CSR.

Table 20-7. I2CSR Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 0 | MCF | Data transfer. When one byte of data is transferred, the bit is cleared. It is set by the falling edge of the 9th clock of a byte transfer. 0 Byte transfer in progress. MCF is cleared under the following conditions: <ul style="list-style-type: none"> • When I2CDR is read in receive mode or when I2CDR is written in transmit mode. • After a start sequence is recognized by the I²C controller in slave mode. 1 Byte transfer is completed |
| 1 | MAAS | Addressed as a slave. When the value in I2CADR matches the calling address or when the calling address is the broadcast address and broadcast mode is enabled (I2CCR[BCST] is set), this bit is set. The processor is interrupted if I2CCR[MIE] is set. Next, the processor must check the SRW bit and set I2CCR[MTX] accordingly. Writing to the I2CCR automatically clears this bit. 0 Not addressed as a slave 1 Addressed as a slave |
| 2 | MBB | Bus busy. Indicates the status of the bus. When a START condition is detected, MBB is set. If a STOP condition is detected, it is cleared. 0 I ² C bus is idle 1 I ² C bus is busy |
| 3 | MAL | Arbitration lost. Automatically set when the arbitration procedure is lost. Note that the device does not automatically retry a failed transfer attempt. 0 Arbitration is not lost. Can only be cleared by software 1 Arbitration is lost |
| 4 | BCSTM | Broadcast match. Writing to the I2CCR automatically clears this bit. 0 There has not been a broadcast match. 1 The calling address matches with the broadcast address and broadcast mode is enabled. This is also set if this I ² C drives an address of all 0s. |
| 5 | SRW | Slave read/write. When MAAS is set, SRW indicates the value of the R/W command bit of the calling address, which is sent from the master. 0 Slave receive, master writing to slave 1 Slave transmit, master reading from slave. This bit is valid only when both of the following conditions are true: <ul style="list-style-type: none"> • A complete transfer occurred and no other transfers have been initiated. • The I²C interface is configured as a slave and has an address match. By checking SRW, the processor can select slave transmit/receive mode according to the command of the master. |
| 6 | MIF | Module interrupt. The MIF bit is set when an interrupt is pending, causing a processor interrupt request (provided I2CCR[MIE] is set). 0 No interrupt is pending. Can be cleared only by software. 1 Interrupt is pending. MIF is set when one of the following events occurs: <ul style="list-style-type: none"> • One byte of data is transferred (set at the falling edge of the 9th clock). • The value in I2CADR matches with the calling address in slave-receive mode. • Arbitration is lost. |
| 7 | RXAK | Received acknowledge. The value of SDA _n during the reception of acknowledge bit of a bus cycle. If the received acknowledge bit (RXAK) is low, it indicates that an acknowledge signal has been received after the completion of eight bits of data transmission on the bus. If RXAK is high, it means no acknowledge signal has been detected at the 9th clock. 0 Acknowledge received 1 No acknowledge received |

20.3.1.5 I²C Data Register (I2CDR)

The I²C data register is shown in Figure 20-6.



Figure 20-6. I²C Data Register (I2CDR)

Table 20-8 shows the bit descriptions for I2CDR.

Table 20-8. I2CDR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–7 | DATA | Transmission starts when an address and the R/W bit are written to the data register and the I ² C interface performs as the master. A data transfer is initiated when data is written to the I2CDR. The most-significant bit is sent first in both cases. In master receive mode, reading the data register allows the read to occur, but also allows the I ² C module to receive the next byte of data on the I ² C interface. In slave mode, the same function is available after it is addressed. Note that in both master receive and slave receive modes, the very first read is always a dummy read. |

20.3.1.6 Digital Filter Sampling Rate Register (I2CDFSRR)

I2CDFSRR is shown in Figure 20-7.




Figure 20-7. I²C Digital Filter Sampling Rate Register (I2CDFSRR)

Table 20-9 shows the I2CDFSRR field descriptions.

Table 20-9. I2CDFSRR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–1 | — | Reserved, should be cleared |
| 2–7 | DFSR | Digital filter sampling rate. To assist in filtering out signal noise, the sample rate is programmed. DFSR is used to prescale the frequency at which the digital filter takes samples from the I ² C bus. The resulting sampling rate is calculated by dividing the platform frequency by the non-zero value of DFSR. If I2CDFSRR is cleared, the I ² C bus sample points default to the reset divisor 0x10. |

20.4 Functional Description

The I²C unit always performs as a slave receiver as a default, unless explicitly programmed to be a master or slave transmitter. If boot sequencer mode is selected, the I²C interface performs as a slave receiver after the boot sequence has completed.

20.4.1 Transaction Protocol

A standard I²C transfer consists of the following:

- START condition
- Slave target address transmission
- Data transfer
- STOP condition

Figure 20-8 shows the interaction of these four parts with the calling address, data byte, and new calling address components of the I²C protocol. The details of the protocol are described in the following subsections.




Figure 20-8. I²C Interface Transaction Protocol

20.4.1.1 START Condition

When the I²C bus is not engaged (both SDA and SCL lines are at logic high), a master can initiate a transfer by sending a START condition. As shown in Figure 20-8, a START condition is defined as a high-to-low transition of SDA while SCL is high. This condition denotes the beginning of a new data transfer. Each data transfer can contain several bytes and awakens all slaves. The START condition is initiated by a software write that sets I2CCR[MSTA].

20.4.1.2 Slave Address Transmission

The first byte of data transferred by the master immediately after the START condition is the slave address. This is a seven-bit calling address followed by a $\overline{R/W}$ bit, which indicates the direction of the data transferred to the slave. Each slave in the system has a unique address. When the I²C module is operating as a master, it must not transmit an address that is the same as its slave address. An I²C device cannot be master and slave at the same time.

Only the slave with a calling address that matches the one transmitted by the master responds by returning an acknowledge bit (negating the SDA signal at the 9th clock) as shown in [Figure 20-8](#). If no slave acknowledges the address, the master should generate a STOP condition or a repeated START condition.

When slave addressing is successful (and SCL returns to zero), the data transfer can proceed on a byte-to-byte basis in the direction specified by the R/ \bar{W} bit sent by the calling master.

The I²C module responds to a general call (broadcast) command when I2CCR[BCST] is set. A broadcast address is always zero; however the I²C module does not check the R/W bit. The second byte of the broadcast message is the master address. Because the second byte is automatically acknowledged by hardware, the receiver device software must verify that the broadcast message is intended for itself by reading the second byte of the message. If the master address is for another receiver device and the third byte is a write command, software can ignore the third byte during the broadcast. If the master address is for another receiver device and the third byte is a read command, software must write 0xFF to I2CDR with I2CCR[TXAK] = 1, so that it does not interfere with the data written from the addressed device.

Each data byte is 8 bits long. Data bits can be changed only while SCL is low and must be held stable while SCL is high, as shown in [Figure 20-8](#). There is one clock pulse on SCL for each data bit, and the most significant bit (msb) is transmitted first. Each byte of data must be followed by an acknowledge bit, which is signaled from the receiving device by pulling SDA low at the 9th clock. Therefore, one complete data byte transfer takes 9 clock pulses. Several bytes can be transferred during a data transfer session.

If the slave receiver does not acknowledge the master, the SDA line must be left high by the slave. The master can then generate a stop condition to abort the data transfer or a START condition (repeated START) to begin a new calling.

If the master receiver does not acknowledge the slave transmitter after a byte of transmission, the slave interprets that the end-of-data has been reached. Then the slave releases the SDA line for the master to generate a STOP or a START condition.

20.4.1.3 Repeated START Condition

[Figure 20-8](#) shows a repeated START condition, which is generated without a STOP condition that can terminate the previous transfer. The master uses this method to communicate with another slave or with the same slave in a different mode (transmit/receive mode) without releasing the bus.

20.4.1.4 STOP Condition

The master can terminate the transfer by generating a STOP condition to free the bus. A STOP condition is defined as a low-to-high transition of the SDA signal while SCL is high. For more information, see [Figure 20-8](#). Note that a master can generate a STOP even if the slave has transmitted an acknowledge bit, at which point the slave must release the bus. The STOP condition is initiated by a software write that clears I2CCR[MSTA].

As described in [Section 20.4.1.3, “Repeated START Condition,”](#) the master can generate a START condition followed by a calling address without generating a STOP condition for the previous transfer. This is called a repeated START condition.

20.4.1.5 Protocol Implementation Details

The following sections give details of how aspects of the protocol are implemented in the I²C module.

20.4.1.5.1 Transaction Monitoring—Implementation Details

The different conditions of the I²C data transfers are monitored as follows (see [Figure 20-8](#)):

- START conditions are detected when an SDA fall occurs while SCL is high.
- STOP conditions are detected when an SDA rise occurs while SCL is high.
- Data transfers in progress are canceled when a STOP condition is detected or if there is a slave address mismatch. Cancellation of data transactions resets the clock module.
- The bus is detected to be busy upon the detection of a START condition and idle upon the detection of a STOP condition.

20.4.1.5.2 Control Transfer—Implementation Details

The I²C module contains logic that controls the output to the serial data (SDA) and serial clock (SCL) lines of the I²C. The SCL output is pulled low as determined by the internal clock generated in the clock module. The SDA output can change only at the midpoint of a low cycle of the SCL, unless it is performing a START, STOP, or repeated START condition. Otherwise, the SDA output is held constant.

SDA is negated when one or more of the following conditions are true:

- Master mode
 - Data bit (transmit)
 - ACK bit (receive)
 - START condition
 - STOP condition
 - Repeated START condition
- Slave mode
 - Acknowledging address match
 - Data bit (transmit)
 - ACK bit (receive)

The SCL signal corresponds to the internal SCL signal when one or more of the following conditions are true in either master or slave mode:

- Master mode
 - Bus owner
 - Lost arbitration
 - START condition
 - STOP condition
 - Repeated START condition begin
 - Repeated START condition end

- Slave mode
 - Address cycle
 - Transmit cycle
 - ACK cycle

20.4.1.6 Address Compare—Implementation Details

The address compare block determines whether a slave has been properly addressed, either by its slave address or by the general broadcast address (which addresses all slaves). The following address comparisons are performed:

- Whether a broadcast message has been received, to update I2CSR
- Whether the module has been addressed as a slave, to update I2CSR and to generate an interrupt
- Whether the address transmitted by the current master matches the general broadcast address

20.4.2 Arbitration Procedure

The I²C interface is a true multiple-master bus. If two or more masters simultaneously try to control the bus, each master's clock synchronization procedure (including the I²C module) determines the bus clock—the low period is equal to the longest clock-low period and the high is equal to the shortest one among the masters. A bus master loses arbitration if it transmits a logic 1 on SDA while another master transmits a logic 0. The losing masters immediately switch to slave-receive mode and stop driving the SDA line. In this case, the transition from master to slave mode does not generate a STOP condition. Meanwhile, the I²C unit sets I2CSR[MAL] to indicate the loss of arbitration and, as a slave, services the transaction if it is directed to itself.

If the I²C module is enabled in the middle of an ongoing byte transfer, the interface behaves as follows:

- Slave mode—the I²C module ignores the current transfer on the bus and starts operating whenever a subsequent START condition is detected.
- Master mode—the I²C module cannot tell whether the bus is busy; therefore, if a START condition is initiated, the current bus cycle can be corrupted. This ultimately causes in the current bus master to lose arbitration, after which bus operations return to normal.

20.4.2.1 Arbitration Control

The arbitration control block controls the arbitration procedure of the master mode. A loss of arbitration occurs whenever the master detects a 0 on the external SDA line while attempting to drive a 1, tries to generate a START or repeated START at an inappropriate time, or detects an unexpected STOP request on the line.

In master mode, arbitration by the master is lost (and I2CSR[MAL] is set) under the following conditions:

- SDA samples low when the master drives high during an address or data-transmit cycle (transmit).
- SDA samples low when the master drives high during a data-receive cycle of the acknowledge (ACK) bit (receive).
- A START condition is attempted when the bus is busy.

- A repeated START condition is requested in slave mode.
- A repeated START condition is attempted when the requesting device is not the bus owner
- Unexpected STOP condition detected

Note that the I²C module does not automatically retry a failed transfer attempt.

20.4.3 Handshaking

The clock synchronization mechanism can be used as a handshake in data transfer. Slave devices can hold SCL low after completion of a 1-byte transfer (9 bits). In such cases, it halts the bus clock and forces the master clock into wait states until the slave releases the SCL line.

20.4.4 Clock Control

The clock control block handles requests from the clock signal for transferring and controlling data for multiple tasks.

A 9-cycle data transfer clock is requested for the following conditions:

- Master mode
 - Transmit slave address after START condition
 - Transmit slave address after repeated START condition
 - Transmit data
 - Receive data
- Slave mode
 - Transmit data
 - Receive data
 - Receive slave address after START or repeated START condition

20.4.4.1 Clock Synchronization

Due to the wire AND logic on the SCL line, a high-to-low transition on the SCL line affects all devices connected on the bus. The devices begin counting their low period when the master negates the SCL line. After a device has negated SCL, it holds the SCL line low until the clock high state is reached. However, the change of low-to-high in a device clock may not change the state of SCL if another device is still within its low period. Therefore, SCL is held low by the device with the longest low period. Devices with shorter low periods enter a high wait state during this time. When all devices concerned have counted off their low periods, SCL is released and asserted. Then there is no difference between the devices' clocks and the state of SCL, and all the devices begin counting their high periods. The first device to complete its high period negates SCL again.

20.4.4.2 Input Synchronization and Digital Filter

The following sections describes synchronization of the input signals and the filtering of SCL and SDA in detail.

20.4.4.2.1 Input Signal Synchronization

The input synchronization block synchronizes the input SCL and SDA signals to the system clock and detects transitions of these signals.

20.4.4.2.2 Filtering of SCL and SDA Lines

The SCL and SDA inputs are filtered to eliminate noise. Three consecutive samples of the SCL and SDA lines are compared to a pre-determined sampling rate. If they are all high, the output of the filter is high. If they are all low, the output is low. If they are any combination of highs and lows, the output is whatever the value of the line was in the previous clock cycle.

The sampling rate is equal to a binary value stored in the frequency register I2CDFSRR. The duration of the sampling cycle is controlled by a down counter. This allows a software write to the I2CDFSRR to control the filtered sampling rate.

20.4.4.3 Clock Stretching

Slaves can use the clock synchronization mechanism to slow down the transfer bit rate. After the master has driven SCL low, the slave can drive SCL low for the required period and then release it. If the slave SCL low period is greater than the master SCL low period, the resulting SCL low period is extended.

20.4.5 Boot Sequencer Mode

Boot sequencer mode is selected at power-on reset by the BOOTSEQ field of the high-order reset configuration word. If boot sequencer mode is selected, the I²C module communicates with one or more EEPROMs through the I²C interface. EEPROMs can be programmed to initialize one or more configuration registers. Note that as described in [Section 4.3.2.2.3, “Boot Sequencer Configuration,”](#) the default value for BOOTSEQ is 0b00, which corresponds to the I²C boot sequencer being disabled at power-up.

Boot sequencer mode also supports an extension of the standard I²C interface that uses more address bits to allow for EEPROM devices that have more than 256 bytes. This extended addressing mode is selectable using a different encoding in the BOOTSEQ field of the high-order reset configuration word (see [Section 4.3.2.2.3, “Boot Sequencer Configuration.”](#)) In this mode, only one EEPROM device can be used and the maximum number of registers is limited by the size of the EEPROM.

If the standard I²C interface is used, the I²C module addresses the first EEPROM, and reads 256 bytes. Then it issues a repeated start and addresses the next EEPROM address. This sequence continues until the CONT bit is cleared. If the last register is not detected before wrapping back to the first address, an error condition is detected. In other words, if the CONT bit for not cleared on the final 7 bytes, an error condition is detected, causing the I²C controller to hang. The I²C module continues to read from the EEPROM as long as the continue (CONT) bit is set in the EEPROM. The CONT bit resides in the address/attributes field that is transferred from the EEPROM, as described in [Section 20.4.5.2, “EEPROM Calling Address.”](#) There should be no other I²C traffic when the boot sequencer is active.

20.4.5.1 Using the Boot Sequencer for Reset Configuration

The reset configuration word can be loaded by using the I²C boot sequencer. See [Section 4.3.2.2.3, “Boot Sequencer Configuration.”](#)

Note that this usage does not prevent using the I²C boot sequencer to initiate the device in the normal functional mode, after reset state has completed. However, an I²C serial EEPROM of extended addressing type must be used and the first two EEPROM data structures must contain dedicated reset information.

20.4.5.2 EEPROM Calling Address

The EEPROM calling address is 0b101_0000. The first EEPROM to be addressed must be programmed to respond to this address, or an error is generated. Any additional EEPROMs are addressed in sequential order.

20.4.5.3 EEPROM Data Format

The I²C module expects a particular format for data to be programmed in the EEPROM. [Figure 20-9](#) shows an example of the EEPROM contents, including the preamble, data format, and CRC.

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | |
|-------------|---------|---|---|---|-------------|---|---|--------------------------------------|
| 1 | 0 | 1 | 0 | 1 | 0 | 1 | 0 | Preamble |
| 0 | 1 | 0 | 1 | 0 | 1 | 0 | 1 | |
| 1 | 0 | 1 | 0 | 1 | 0 | 1 | 0 | |
| ACS | BYTE_EN | | | 1 | ADDR[12:13] | | | First Configuration Preload Command |
| ADDR[14:21] | | | | | | | | |
| ADDR[22:29] | | | | | | | | |
| DATA[0:7] | | | | | | | | |
| DATA[8:15] | | | | | | | | |
| DATA[16:23] | | | | | | | | |
| DATA[24:31] | | | | | | | | |
| ACS | BYTE_EN | | | 1 | ADDR[12:13] | | | Second Configuration Preload Command |
| ADDR[14:21] | | | | | | | | |
| ADDR[22:29] | | | | | | | | |
| DATA[0:7] | | | | | | | | |
| DATA[8:15] | | | | | | | | |
| DATA[16:23] | | | | | | | | |
| DATA[24:31] | | | | | | | | |
| | | | | | | | | |
| ACS | BYTE_EN | | | 1 | ADDR[12:13] | | | Last Configuration Preload Command |
| ADDR[14:21] | | | | | | | | |
| ADDR[22:29] | | | | | | | | |
| DATA[0:7] | | | | | | | | |
| DATA[8:15] | | | | | | | | |
| DATA[16:23] | | | | | | | | |
| DATA[24:31] | | | | | | | | |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | End Command |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | |
| CRC[0:7] | | | | | | | | |
| CRC[8:15] | | | | | | | | |
| CRC[16:23] | | | | | | | | |
| CRC[24:31] | | | | | | | | |

Figure 20-9. EEPROM Contents

- A preamble should be the first 3 bytes programmed into the EEPROM. It should have a value of 0xAA55AA. The I²C checks to ensure that this preamble is correctly detected before proceeding.
- Following the preamble, there should be a series of configuration registers (known as register preloads). Each configuration register should be programmed according to a particular format, as shown in [Figure 20-10](#).

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|-------------|---------|---|---|------|-------------|---|---|
| ACS | BYTE_EN | | | CONT | ADDR[12:13] | | |
| ADDR[14:21] | | | | | | | |
| ADDR[12:29] | | | | | | | |
| DATA[0:7] | | | | | | | |
| DATA[8:15] | | | | | | | |
| DATA[16:23] | | | | | | | |
| DATA[24:31] | | | | | | | |

Figure 20-10. EEPROM Data Format for One Register Preload Command

- The first byte holds alternate configuration space (ACS), byte enables, and continue (CONT) attributes.
- The 2 least-significant bits of the address are derived from the byte enables. address offset. Therefore, the address offset programmed into the EEPROM preload should be a word offset.
- The most significant 16 bits (assuming 36-bit addressing) of the address are prepended from either IMMRRBAR or alternate configuration space.
- After the first 3 bytes, 4 bytes of data should hold the desired value of the configuration register, regardless of the size of transaction.

Byte enables should be asserted for any byte that will be written, and they should be asserted contiguously, creating a 1, 2, or 4 byte write to a register. The boot sequencer assumes that a big-endian address is stored in the EEPROM. In addition, byte enable bit 0 (bit 1 of the byte) corresponds to the most-significant byte of data (data[0:7]), and byte enable bit 3 (bit 4 of the byte) corresponds to the least-significant byte of data (data[24:31]).

By asserting ACS, an alternate configuration space address is prepended to the write request from the boot sequencer according to the value in the ALTCBAR register. This will allow for external memories to be configured. Otherwise, IMMRRBAR is prepended to the EEPROM address.

If the CONT bit is cleared, the first 3 bytes, including ACS, the byte enables, and the address, should be cleared 0. Also, the data contains the final CRC. A CRC-32 algorithm is used to check the integrity of the data. The following polynomial is used:

$$1 + x^1 + x^2 + x^4 + x^5 + x^7 + x^8 + x^{10} + x^{11} + x^{12} + x^{16} + x^{22} + x^{23} + x^{26} + x^{32}$$

The CRC should cover all bytes stored in the EEPROM before the CRC. This includes the preamble, all register preloads, and the first 3 bytes of the last 7-byte preload (which should be all zeros).

20.4.5.4 Boot Sequencer Done Indication

Dedicated hardware is not provided to indicate whether the boot sequencer operation completed successfully. It is recommended to use one of the GPIO signals for that purpose. To do this, the last register preload programmed into the EEPROM should contain the address of the appropriate GPIO register and data that causes the setting of the required GPIO signal. The GPIO signal may be used for an external device or for debug purposes.

20.5 Initialization/Application Information

This section describes some programming guidelines recommended for the I²C interface. [Figure 20-11](#) is a recommended flowchart for I²C interrupt service routines.

A **sync** assembly instruction must be executed after each I²C register read/write access to guarantee that register accesses occur in order.

The I²C controller does not guarantee its recovery from all illegal I²C bus activity. In addition, a malfunctioning device may hold the bus captive. A good programming practice is for software to rely on a watchdog timer to help recover from I²C bus hangs. The recovery routine should also handle the case when the illegal I²C bus behavior causes the status bits returned after an interrupt to be inconsistent with what was expected.

20.5.1 Interrupt Service Routine Flowchart

[Figure 20-11](#) shows an example algorithm for an I²C interrupt service routine. Deviation from the flowchart may result in unpredictable I²C bus behavior. However, in the slave receive mode (not shown), the interrupt service routine may need to set I2CCR[TXAK] when the next-to-last byte is to be accepted. It is recommended that a **sync** instruction follow each I²C register read or write to guarantee that register accesses occur in order.

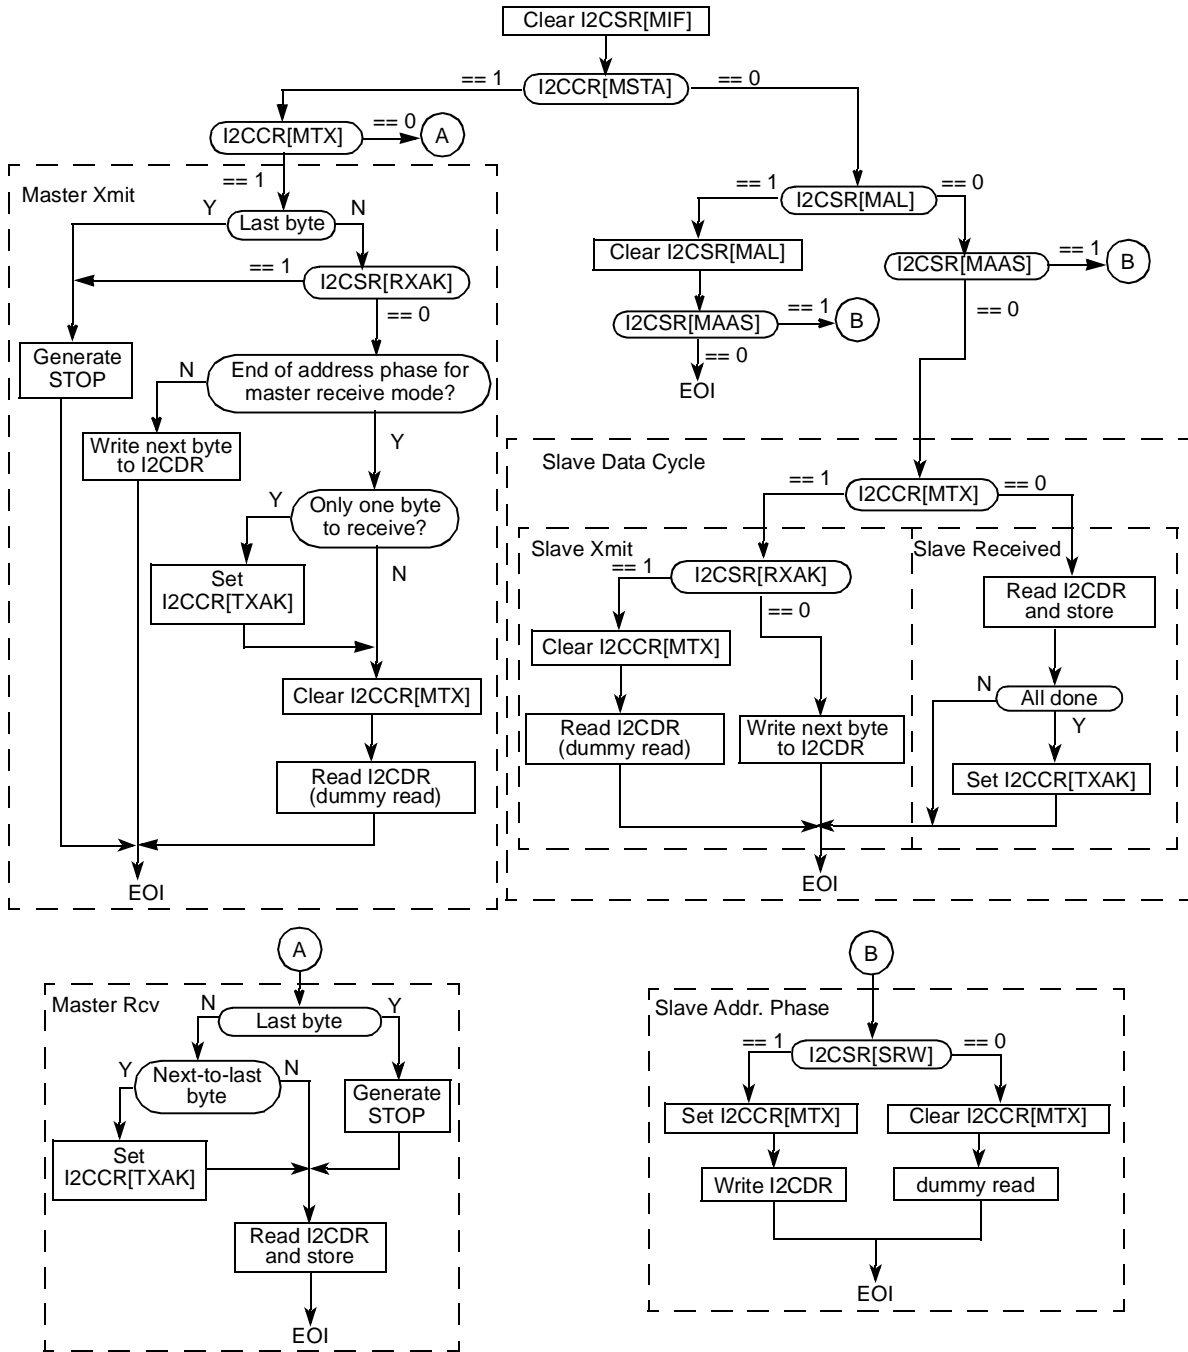


Figure 20-11. Example I²C Interrupt Service Routine Flowchart

20.5.2 Initialization Sequence

A hard reset initializes all of the I²C registers to their default states. The following initialization sequence initializes the I²C unit:

1. All I²C registers must be located in a cache-inhibited page.

2. Update I2CFDR[FDR] and select the required division ratio to obtain the SCL frequency from the CSB (platform) clock.
3. Update I2CADR to define the slave address for this device.
4. Modify I2CCR to select master/slave mode, transmit/receive mode, and interrupt-enable or disable.
5. Set the I2CCR[MEN] to enable the I²C interface.

20.5.3 Generation of START

After initialization, the following sequence can be used to generate START:

1. If the device is connected to a multimaster I²C system, check whether the serial bus is free (I2CSR[MBB] = 0) before switching to master mode.
2. Select master mode (set I2CCR[MSTA]) to transmit serial data and select transmit mode (set I2CCR[MTX]) for the address cycle.
3. Write the slave address being called into I2CDR. The data written to I2CDR[0–6] comprises the slave calling address. I2CCR[MTX] indicates the direction of transfer (transmit/receive) required from the slave.

The scenario above assumes that the I²C interrupt bit (I2CSR[MIF]) is cleared. If MIF is set at any time, an I²C interrupt is generated (provided interrupt reporting is enabled with I2CCR[MIEN] = 1).

20.5.4 Post-Transfer Software Response

Transmission or reception of a byte automatically sets the data transferring bit (I2CSR[MCF]), which indicates that one byte has been transferred. The I²C interrupt bit (I2CSR[MIF]) is also set and an interrupt is generated to the processor if the interrupt function is enabled during the initialization sequence (I2CCR[MIEN] is set). In the interrupt handler, software must take the following steps:

1. Clear I2CSR[MIF]
2. Read the I2CDR in receive mode or write to I2CDR in transmit mode. Note that this causes I2CSR[MCF] to be cleared, as shown in [Figure 20-11](#).
3. When an interrupt occurs at the end of the address cycle, the master remains in transmit mode. If master receive mode is required, I2CCR[MTX] must be toggled at this stage (see [Figure 20-11](#)).

If the interrupt function is disabled, software can service the I2CDR in the main program by monitoring I2CSR[MIF]. In this case, I2CSR[MIF] must be polled rather than I2CSR[MCF] because MCF behaves differently when arbitration is lost. Note that interrupt or other bus conditions may be detected before the I²C signals have time to settle. Thus, when polling I2CSR[MIF] (or any other I2CSR bits), software delays may be needed to give the I²C signals sufficient time to settle.

During slave-mode address cycles (I2CSR[MAAS] is set), I2CSR[SRW] should be read to determine the direction of the subsequent transfer and I2CCR[MTX] should be programmed accordingly. For slave-mode data cycles (MAAS is cleared), I2CSR[SRW] is not valid and I2CCR[MTX] must be read to determine the direction of the current transfer (see [Figure 20-11](#)).

20.5.5 Generation of STOP

A data transfer ends with a STOP condition generated by the master device. A master transmitter can generate a STOP condition after all the data has been transmitted.

If a master receiver wants to terminate a data transfer, it must inform the slave transmitter by not acknowledging the last byte of data (by setting the transmit acknowledge bit (I2CCR[TXAK])) before reading the next-to-last byte of data. At this time, the next-to-last byte of data has been transferred on the I²C interface, so the last byte does not receive the data acknowledge (because I2CCR[TXAK] is set). Before the interrupt service routine reads the last byte of data, a STOP condition must first be generated.

20.5.6 Generation of Repeated START

At the end of a data transfer, if the master still wants to communicate on the bus, it can generate another START condition followed by another slave address without first generating a STOP condition. This is accomplished by setting I2CCR[RSTA].

20.5.7 Generation of SCL When SDA is Negated

It is sometimes necessary to force the I²C module to become the I²C bus master out of reset and drive SCL_n (even though SDA may already be driven, which indicates that the bus is busy). This can occur when a system reset does not cause all I²C devices to be reset. Thus, SDA can be negated low by another I²C device while this I²C module is coming out of reset and will stay low indefinitely. The following procedure can be used to force this I²C module to generate SCL so that the device driving SDA can finish its transaction:

1. Disable the I²C module and set the master bit by setting I2CCR to 0x20.
2. Enable the I²C module by setting I2CCR to 0xA0.
3. Read I2CDR.
4. Return the I²C module to slave mode by setting I2CCR to 0x80.

20.5.8 Slave Mode Interrupt Service Routine

In the slave interrupt service routine, the module addressed as a slave should be tested to check if a calling of its own address has been received. If I2CSR[MAAS] is set, software should set the transmit/receive mode select bit (I2CCR[MTX]) according to the R/ \overline{W} command bit (I2CSR[SRW]). Writing to I2CCR clears MAAS automatically. MAAS is read as set only in the interrupt handler at the end of that address cycle where an address match occurred; interrupts resulting from subsequent data transfers clear MAAS. A data transfer can then be initiated by writing to I2CDR for slave transmits or dummy reading from I2CDR in slave-receive mode. The slave negates SCL between byte transfers. SCL is released when I2CDR is accessed in the required mode.

20.5.8.1 Slave Transmitter and Received Acknowledge

In the slave transmitter routine, the received acknowledge bit (I2CSR[RXAK]) must be tested before sending the next byte of data. The master signals an end-of-data by not acknowledging the data transfer

from the slave. When no acknowledge is received (I2CSR[RXAK] is set), the slave transmitter interrupt routine must clear I2CCR[MTX] to switch the slave from transmitter to receiver mode. A dummy read of I2CDR then releases SCL so that the master can generate a STOP condition. See [Figure 20-11](#).

20.5.8.2 Loss of Arbitration and Forcing of Slave Mode

When a master loses arbitration the following conditions all occur:

- I2CSR[MAL] is set
- I2CCR[MSTA] is cleared (changing the master to slave mode)
- An interrupt occurs (if enabled) at the falling edge of the 9th clock of this transfer

Thus, the slave interrupt service routine should first test I2CSR[MAL] and software should clear it if it is set. See [Section 20.4.2.1, “Arbitration Control.”](#)



Chapter 21

DUART

This chapter describes the two (dual) universal asynchronous receiver/transmitters (UARTs) of the device. It describes the functional operation, the DUART initialization sequence, and the programming details for the DUART registers and features.

21.1 DUART Overview

The DUART consists of two (dual) universal asynchronous receiver/transmitters (UARTs). The UARTs act independently; all references to UART refer to one of these receiver/transmitters. Each UART is clocked by the system clock. The DUART programming model is compatible with the PC16552D.

The UART interface is point-to-point, meaning that only two UART devices are attached to the connecting signals. As shown in [Figure 21-1](#), each UART module consists of the following:

- Receive and transmit buffers
- Clear to send ($\overline{\text{CTS}}$) input port and request to send ($\overline{\text{RTS}}$) output port for data-flow control.
- 16-bit counter for baud rate generation
- Interrupt control logic




Figure 21-1. UART Block Diagram

21.1.1 DUART Features

The DUART includes these features:

- Full-duplex operation
- Programming model compatible with the original PC16450 UART and the PC16550D (an improved version of the PC16450 that also operates in FIFO mode)
- PC16450 register reset values
- FIFO mode for both transmitter and receiver, providing 16-byte FIFOs
- Serial data encapsulation and decapsulation with standard asynchronous communication bits (START, STOP, and parity)
- Maskable transmit, receive, line status, and MODEM status interrupts
- Software-programmable baud generators that divide the system clock by 1 to $(2^{16}-1)$ and generate a 16x clock for the transmitter and receiver engines
- Clear-to-send (\overline{CTS}) and ready-to-send (\overline{RTS}) MODEM control functions
- Software-selectable serial interface data format (data length, parity, 1/1.5/2 STOP bit, baud rate)
- Line and MODEM status registers
- Line-break detection and generation
- Internal diagnostic support, local loopback, and break functions
- Prioritized interrupt reporting
- Overrun, parity, and framing error detection

21.1.2 DUART Modes of Operation

The communication channel provides a full-duplex asynchronous receiver and transmitter using an operating frequency derived from the system clock.

The transmitter accepts parallel data from a write to the transmitter holding register (UTHR). In FIFO mode, the data is placed directly into an internal transmitter shift register of the transmitter FIFO. The transmitter converts the data to a serial bit stream, inserting the appropriate START, STOP, and optional parity bits. Finally, it outputs a composite serial data stream on the channel transmitter serial data output signal (SOUT). The transmitter status may be polled or interrupt driven.

The receiver accepts serial data bits on the channel receiver serial data input signal (SIN), converts it to parallel format, checks for a START bit, parity (if any), STOP bits, and transfers the assembled character (with START, STOP, parity bits removed) from the receiver buffer (or FIFO) in response to a read of the UART's receiver buffer register (URBR). The receiver status may be polled or interrupt driven.

21.2 DUART External Signal Descriptions

This section contains a signal overview and detailed signal descriptions.

21.2.1 DUART Signal Overview

Table 21-1 summarizes the DUART signals. Note that although the actual device signal names are prepended with the 'UART_' prefix as shown in the table, the functional (abbreviated) signal names are often used throughout this chapter.

Table 21-1. DUART Signal Overview

| Signal Name | I/O | Pins | Reset Value | State Meaning |
|------------------------------------|-----|------|-------------|---------------------------------|
| UART_SIN[1:2] | I | 2 | 1 | Serial in data UART1 and UART2 |
| UART_SOUT[1:2] | O | 2 | 1 | Serial out data UART1 and UART2 |
| $\overline{\text{UART_CTS}}[1:2]$ | I | 2 | 1 | Clear to send UART1 and UART2 |
| $\overline{\text{UART_RTS}}[1:2]$ | O | 2 | 1 | Request to send UART1 and UART2 |

21.2.2 DUART Detailed Signal Descriptions

The DUART signals are described in detail in Table 21-2.

Table 21-2. DUART Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|----------------------------|-----|--|
| UART_SIN[1:2]/DSP_UART_SIN | I | Serial data in. Data is received on the receivers of UART1, UART2, or DSP_UART through its respective serial data input signal, with the least significant bit received first. |
| | | State Meaning Asserted/Negated—Represents the data being received on the UART interface. |
| | | Timing Assertion/Negation—An internal logic sample signal, <i>rxcnt</i> , uses the frequency of the baud-rate generator to sample the data on SIN. |

Table 21-2. DUART Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|----------------------------------|-----|---|--|
| UART_SOUT[1:2]/ DSP_UART_SOUT | O | Serial data out. The serial data output signals for the UART1, UART2, or DSP_UART are set (mark condition) when the transmitter is disabled, idle, or operating in the local loopback mode. Data is shifted out on these signals, with the least significant bit transmitted first. | |
| | | State Meaning | Asserted/Negated—Represents the data transmitted on the respective UART interface. |
| | | Timing | Assertion/Negation—An internal logic sample signal, <i>rxcnt</i> , uses the frequency of the baud-rate generator to update and drive the data on SOUT. |
| UART_CTS[1:2] | I | Clear to send. Connected to the respective $\overline{\text{RTS}}$ outputs of the other UART devices on the bus. They can be programmed to generate an interrupt on change-of-state of the signal. | |
| | | State Meaning | Asserted/Negated—Represent the clear to send condition for their respective UART. |
| | | Timing | Assertion/Negation—Sampled at the rising edge of every system clock. |
| UART_RTS[1:2] | O | Request to send. Can be programmed to be negated and asserted by either the receiver or transmitter. When connected to the CTS input of a transmitter, this signal can be used to control serial data flow. | |
| | | State Meaning | Asserted/Negated—Represents the data being transmitted on the respective UART interface. |
| | | Timing | Assertion/Negation—Updated and driven at the rising edge of every system clock. |

21.3 DUART Memory Map/Register Definition

There are two complete sets of DUART registers (one for UART1 and one for UART2). The two UARTs are identical, except that the registers for UART1 are located at offsets 0x0_4500 (local), and the registers for UART2 are located at offsets 0x0_4600 (local). Throughout this chapter, the registers are described by a singular acronym: for example, LCR represents the line control register for either UART1 or UART2.

The registers in each UART interface are used for configuration, control, and status. The divisor latch access bit, ULCR[DLAB], is used to access the divisor latch least- and most-significant bit registers and the alternate function register. Refer to [Section 21.3.1.8, “Line Control Registers \(ULCR1 and ULCR2\),”](#) for more information on ULCR[DLAB].

All DUART registers are one byte wide; reads and writes to these registers must be byte-wide operations. [Table 21-3](#) provides a register summary with references to the section and page that contain detailed

information about each register. Undefined byte address spaces within offset 0x4000–0x4FFF are reserved.

Table 21-3. DUART Register Summary

| Offset | Register | Access | Reset | Section/Page |
|--|---|--------|-------|---------------------------------|
| UART 1—Block Base Address 0x0_4000 UART 2—Block Base Address 0x0_4100 | | | | |
| 0x0_4500 | URBR—ULCR[DLAB] = 0 UART1 receiver buffer register | R | 0x00 | 21.3.1.1/21-6 |
| | UTHR—ULCR[DLAB] = 0 UART1 transmitter holding register | W | 0x00 | 21.3.1.2/21-6 |
| | UDLB—ULCR[DLAB] = 1 UART1 divisor least significant byte register | R/W | 0x00 | 21.3.1.3/21-7 |
| 0x0_4501 | UIER—ULCR[DLAB] = 0 UART1 interrupt enable register | R/W | 0x00 | 21.3.1.4/21-8 |
| | UDMB—ULCR[DLAB] = 1 UART1 divisor most significant byte register | R/W | 0x00 | 21.3.1.3/21-7 |
| 0x0_4502 | UIIR—ULCR[DLAB] = 0 UART1 interrupt ID register | R | 0x01 | 21.3.1.5/21-9 |
| | UFCR—ULCR[DLAB] = 0 UART1 FIFO control register | W | 0x00 | 21.3.1.6/21-10 |
| | UAFR—ULCR[DLAB] = 1 UART1 alternate function register | R/W | 0x00 | 21.3.1.7/21-11 |
| 0x0_4503 | ULCR—ULCR[DLAB] = x UART1 line control register | R/W | 0x00 | 21.3.1.8/21-12 |
| 0x0_4504 | UMCR—ULCR[DLAB] = x UART1 MODEM control register | R/W | 0x00 | 21.3.1.9/21-14 |
| 0x0_4505 | ULSR—ULCR[DLAB] = x UART1 line status register | R | 0x60 | 21.3.1.10/21-15 |
| 0x0_4506 | UMSR—ULCR[DLAB] = x UART1 MODEM status register | R | 0x00 | 21.3.1.11/21-16 |
| 0x0_4507 | USCR—ULCR[DLAB] = x UART1 scratch register | R/W | 0x00 | 21.3.1.12/21-17 |
| 0x0_4510 | UDSR—ULCR[DLAB] = x UART1 DMA status register | R | 0x01 | 21.3.1.13/21-17 |
| 0x0_4600 | URBR—ULCR[DLAB] = 0 UART2 receiver buffer register | R | 0x00 | 21.3.1.1/21-6 |
| | UTHR—ULCR[DLAB] = 0 UART2 transmitter holding register | W | 0x00 | 21.3.1.2/21-6 |
| | UDLB—ULCR[DLAB] = 1 UART2 divisor least significant byte register | R/W | 0x00 | 21.3.1.3/21-7 |
| 0x0_4601 | UIER—ULCR[DLAB] = 0 UART2 interrupt enable register | R/W | 0x00 | 21.3.1.4/21-8 |
| | UDMB—ULCR[DLAB] = 1 UART2 divisor most significant byte register | R/W | 0x00 | 21.3.1.3/21-7 |
| 0x0_4602 | UIIR—ULCR[DLAB] = 0 UART2 interrupt ID register | R | 0x01 | 21.3.1.5/21-9 |
| | UFCR—ULCR[DLAB] = 0 UART2 FIFO control register | W | 0x00 | 21.3.1.6/21-10 |
| | UAFR—ULCR[DLAB] = 1 UART2 alternate function register | R/W | 0x00 | 21.3.1.7/21-11 |
| 0x0_4603 | ULCR—ULCR[DLAB] = x UART2 line control register | R/W | 0x00 | 21.3.1.8/21-12 |
| 0x0_4604 | UMCR—ULCR[DLAB] = x UART2 MODEM control register | R/W | 0x00 | 21.3.1.9/21-14 |
| 0x0_4605 | ULSR—ULCR[DLAB] = x UART2 line status register | R | 0x60 | 21.3.1.10/21-15 |
| 0x0_4606 | UMSR—ULCR[DLAB] = x UART2 MODEM status register | R | 0x00 | 21.3.1.11/21-16 |
| 0x0_4607 | USCR—ULCR[DLAB] = x UART2 scratch register | R/W | 0x00 | 21.3.1.12/21-17 |
| 0x0_4610 | UDSR—ULCR[DLAB] = x UART2 DMA status register | R | 0x01 | 21.3.1.13/21-17 |

21.3.1 DUART Register Descriptions

The following sections describe the UART1 and UART2 registers.

21.3.1.1 Receiver Buffer Registers (URBR1 and URBR2)

These registers contain the data received from the transmitter on the UART buses. In FIFO mode, when read, they return the first byte received. For FIFO status information, refer to the UDSR[RXRDY] description.

Except for the case when there is an overrun, URBR returns the data in the order it was received from the transmitter. Refer to the ULSR[OE] description, [Section 21.3.1.10, “Line Status Registers \(ULSR1 and ULSR2\).”](#) [Figure 21-2](#) shows the receiver buffer registers. Note that these registers have same offset as the UTHR.




Figure 21-2. Receiver Buffer Registers (URBR1 and URBR2)

[Table 21-4](#) describes URBR.

Table 21-4. URBR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–7 | DATA | Data received from the transmitter on the UART bus [read only] |

21.3.1.2 Transmitter Holding Registers (UTHR1 and UTHR2)

A write to these 8-bit registers causes the UART devices to transfer 5 to 8 data bits on the UART bus in the format set up in the ULCR (line control register). In FIFO mode, data written to UTHR is placed into the FIFO. The data written to UTHR is the data sent onto the UART bus, and the first byte written to UTHR is the first byte onto the bus. UDSR[TXRDY] indicates when the FIFO is full. Refer to [Table 21-21](#) and [Table 21-22](#).

[Figure 21-3](#) shows the bits in the UTHR.




Figure 21-3. Transmitter Holding Registers (UTHR1 and UTHR2)

Table 21-5 describes the UTHR.

Table 21-5. UTHR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–7 | DATA | Data that is written to UTHR [Write only] |

21.3.1.3 Divisor Most and Least Significant Byte Registers (UDMB and UDLB)

UDLB is concatenated with the divisor most significant byte register (UDMB) to create the divisor used to divide the input clock into the DUART. The output frequency of the baud generator is 16 times the baud rate; therefore, the desired baud rate = platform clock frequency \div (16 \times [UDMB||UDLB]). Equivalently, [UDMB||UDLB:0b0000] = platform clock frequency/desired baud rate. Baud rates that can be generated by specific input clock frequencies are shown in Table 21-8.

Figure 21-4 shows the bits in the UDMBs.




Figure 21-4. Divisor Most Significant Byte Registers (UDMB1 and UDMB2)

Table 21-6 describes the UDMB.

Table 21-6. UDMB Field Descriptions

| Bits | Name | Description |
|------|------|-------------------------------|
| 0–7 | UDMB | Divisor most significant byte |

Figure 21-5 shows the bits in the UDLBs.




Figure 21-5. Divisor Least Significant Byte Registers (UDLB1 and UDLB2)

Table 21-7 describes the UDLB.

Table 21-7. UDLB Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–7 | UDLB | Divisor least significant byte. This is concatenated with UDMB. |

Table 21-8 shows baud rate for a variety of input clock frequencies.

Table 21-8. Baud Rate Examples

| Baud Rate (Decimal) | Divisor | | Input Clock (System Clock) Frequency (MHz) | Percent Error (Decimal) |
|------------------------|---------|-----|---|----------------------------|
| | Decimal | Hex | | |
| 9,600 | 866 | 362 | 133 | 0.013 |
| 19,200 | 433 | 1B1 | 133 | 0.013 |
| 38,400 | 216 | D8 | 133 | 0.218 |
| 56,000 | 148 | 94 | 133 | 0.300 |
| 128,000 | 65 | 41 | 133 | 0.090 |
| 256,000 | 32 | 20 | 133 | 1.471 |

To get the percent error value, the following three steps are taken:

1. The input clock frequency (ICF) is divided by the actual frequency input (AFI) to get the correct divisor value (ICF/AFI, where AFI = baud rate × 16 × divisor).
2. The divisor value is subtracted from 1.
3. The result from the step two is multiplied by 100 to calculate the final percent error. The result is calculated in absolute value (no negative numbers).

These steps can be described with the following equation:

$$\text{Percent error value} = (1 - \text{AFI/ICF}) \times 100$$

21.3.1.4 Interrupt Enable Registers (UIER1 and UIER2)

The UIER gives the user the ability to mask specific UART interrupts to the programmable interrupt controller (PIC).

Figure 21-6 shows the bits in the UIER.




Figure 21-6. Interrupt Enable Registers (UIER1 and UIER2)

Table 21-9 describes the UIER fields.

Table 21-9. UIER Field Descriptions

| Bits | Name | Description |
|------|--------|--|
| 0–3 | — | Reserved |
| 4 | EMSI | Enable MODEM status interrupt 0 Mask interrupts caused by UMSR[DCTS] being set. 1 Enable and assert interrupts when UMSR[CTS] changes state. |
| 5 | ERLSI | Enable receiver line status interrupt 0 Mask interrupts when ULSR's overrun, parity error, framing error, or break interrupt bits are set. 1 Enable and assert interrupts when ULSR's overrun, parity error, framing error or break interrupt bits are set. |
| 6 | ETHREI | Enable transmitter holding register empty interrupt 0 Mask interrupt when ULSR[THRE] is set. 1 Enable and assert interrupts when ULSR[THRE] is set. |
| 7 | ERDAI | Enable received data available interrupt 0 Mask interrupt when new receive data is available or receive data time-out has occurred. 1 Enable and assert interrupts when a new data character is received from the external device and/or a time-out interrupt occurs in FIFO mode. |

21.3.1.5 Interrupt ID Registers (UIIR1 and UIIR2)

The UIIRs indicate when an interrupt is pending from the corresponding UART and what type of interrupt is active. They also indicate if the FIFOs are enabled.

The DUART prioritizes interrupts into four levels and records these in the corresponding UIIR. The four levels of interrupt conditions in order of priority are as follows:

1. Receiver line status
2. Received data ready/character time-out
3. Transmitter holding register empty
4. MODEM status

See Table 21-11 for more details.

When the UIIR is read, the associated DUART serial channel freezes all interrupts and indicates the highest priority pending interrupt. While this read transaction is occurring, the associated DUART serial channel records new interrupts, but does not change the contents of UIIR until the read access is complete.

Figure 21-7 shows the bits in the UIIR.



Figure 21-7. Interrupt ID Registers (UIIR1 and UIIR2)

Table 21-10 describes the fields of the UIIR.

Table 21-10. UIIR Field Descriptions

| Bits | Name | Description |
|------|-----------|---|
| 0–1 | FE | FIFOs enabled. Reflects the setting of UFCR[FEN]. |
| 2–3 | — | Reserved |
| 4 | IID3 | Interrupt ID bits identify the highest priority interrupt that is pending as indicated in Table 21-11. IID3 is set along with IID2 only when a time out interrupt is pending for FIFO mode. |
| 5–6 | IID2–IID1 | Interrupt ID bits identify the highest priority pending interrupt as indicated in Table 21-11. |
| 7 | IID0 | IID0 indicates when an interrupt is pending. 0 The UART has an active interrupt ready to be serviced. 1 No interrupt is pending. |

The bits contained in the UIIR registers are described in Table 21-11.

Table 21-11. UIIR IID Bits Summary

| IID3–IID0 | Priority Level | Interrupt Type | Interrupt Description | How To Reset Interrupt |
|-----------|----------------|-------------------------|--|--|
| 0001 | — | — | — | — |
| 0110 | Highest | Receiver line status | Overrun error, parity error, framing error, or break interrupt | Reading the line status register |
| 0100 | Second | Received data available | Receiver data available or trigger level reached in FIFO mode. | Reading the receiver buffer register or if the number of bytes in the receiver FIFO drops below the trigger level. |
| 1100 | Second | Character time-out | No characters were removed from or input to the receiver FIFO during the last four character times and at least one character is in the receiver FIFO. | Reading the receiver buffer register |
| 0010 | Third | UTHR empty | Transmitter holding register is empty. | Reading UIIR or writing to UTHR |
| 0000 | Fourth | MODEM status | $\overline{\text{CTS}}$ input value changed since last read of UMSR. | Reading UMSR |

21.3.1.6 FIFO Control Registers (UFCR1 and UFCR2)

UFCR is used to enable and clear the receiver and transmitter FIFOs, set a receiver FIFO trigger level to control the received data available interrupt, and select the type of DMA signaling.

UFCR bits cannot be programmed unless FIFO enable bits are set. When changing from FIFO mode to 16450 mode (non-FIFO mode) and vice versa, data is automatically cleared from the FIFOs.

After all of the bytes in the receiver FIFO are cleared, the receiver internal shift register is not cleared. Similarly, the bytes are cleared in the transmitter FIFO, but the transmitter internal shift register is not cleared. Both TFR and RFR are self clearing.

Figure 21-8 shows the bits in the UFCRs.




Figure 21-8. FIFO Control Registers (UFCR1 and UFCR2)

Table 21-12 describes the fields of the UFCRs.

Table 21-12. UFCR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–1 | RTL | Receiver trigger level. A received data available interrupt occurs when UIER[ERDAI] is set and the number of bytes in the receiver FIFO equals RTL value. 00 1 byte 01 4 bytes 10 8 bytes 11 14 bytes |
| 2–3 | — | Reserved |
| 4 | DMS | DMA mode select. See Section 21.4.5.2, “DMA Mode Select” 0 UDSR[RXRDY] and UDSR[TXRDY] bits are in mode 0. 1 UDSR[RXRDY] and UDSR[TXRDY] bits are in mode 1 if UFCR[FEN] = 1. |
| 5 | TFR | Transmitter FIFO reset 0 No action 1 Clears all bytes in the transmitter FIFO and resets the FIFO counter/pointer to 0 |
| 6 | RFR | Receiver FIFO reset 0 No action 1 Clears all bytes in the receiver FIFO and resets the FIFO counter/pointer to 0 |
| 7 | FEN | FIFO enable 0 FIFOs are disabled and cleared 1 Transmitter and receiver FIFOs are enabled. |

21.3.1.7 Alternate Function Registers (UAFR1 and UAFR2)

The UAFRs, shown in [Figure 21-9](#), allow software to write to both UART1 and UART2 registers simultaneously with the same write operation. The UAFRs also provide a means for the device's performance monitor to track the baud clock.



Figure 21-9. Alternate Function Register (UAFR)

Table 21-13 describes UAFR fields.

Table 21-13. UAFR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–5 | — | Reserved |
| 6 | BO | Baud clock select 0 The baud clock is not gated off. 1 The baud clock is gated off. |
| 7 | CW | Concurrent write enable 0 Disables writing to both UART1 and UART2. 1 Enables concurrent writes to corresponding UART registers. A write to a register in UART1 is also a write to the corresponding register in UART2 and vice versa. |

21.3.1.8 Line Control Registers (ULCR1 and ULCR2)

The ULCRs specify the data format for the UART bus and set the divisor latch access bit ULCR[DLAB], which controls the ability to access the divisor latch least and most significant bit registers and the alternate function register.

After initializing ULCR, the software should not rewrite the ULCR while valid transfers on the UART bus are active. The software should not rewrite the ULCR until the last STOP bit is received and no new characters are being transferred on the bus.

The stick parity bit, ULCR[SP], assigns a set parity value for the parity bit time slot sent on the UART bus. The set value is defined as mark parity (logic 1) or space parity (logic 0). ULCR[PEN] and ULCR[EPS] help determine the set parity value. See Table 21-15. ULCR[NSTB] defines the number of STOP bits to be sent at the end of the data transfer. The receiver checks only the first STOP bit, regardless of the number of STOP bits selected. The word length select bits (1 and 0) define the number of data bits transmitted or received as a serial character. The word length does not include START, parity, and STOP bits.

Figure 21-10 shows the bits in the ULCRs.




Figure 21-10. Line Control Register (ULCR1 and ULCR2)

Table 21-14 describes the ULCR fields.

Table 21-14. ULCR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0 | DLAB | Divisor latch access bit 0 Access to all registers except UDLB, UAFR, and UDMB. 1 Ability to access UDMB, UDLB, and UAFR. |
| 1 | SB | Set break 0 Send normal UTHR data onto the SOUT signal. 1 Force logic 0 to be on SOUT. Data in the UTHR is not affected. |
| 2 | SP | Stick parity 0 Stick parity is disabled. 1 If PEN = 1 and EPS = 1, space parity is selected; if PEN = 1 and EPS = 0, mark parity is selected. |
| 3 | EPS | Even parity select. See Table 21-15 . 0 If PEN = 1 and SP = 0 then odd parity is selected. 1 If PEN = 1 and SP = 0 then even parity is selected. |
| 4 | PEN | Parity enable 0 No parity generation and checking. 1 Generate parity bit as a transmitter, and check parity as a receiver. |
| 5 | NTSB | Number of STOP bits 0 One STOP bit is generated in the transmitted data. 1 When a 5-bit data length is selected, 1 1/2 STOP bits are generated. When either a 6-, 7-, or 8-bit word length is selected, two STOP bits are generated. |
| 6–7 | WLS | Word length select. Number of bits that comprise the character length. 00 5 bits 01 6 bits 10 7 bits 11 8 bits |

Table 21-15. Parity Selection Using ULCR[PEN], ULCR[SP], and ULCR[EPS]

| PEN | SP | EPS | Parity Selected |
|-----|----|-----|-----------------|
| 0 | 0 | 0 | No parity |
| 0 | 0 | 1 | No parity |
| 0 | 1 | 0 | No parity |
| 0 | 1 | 1 | No parity |
| 1 | 0 | 0 | Odd parity |
| 1 | 0 | 1 | Even parity |
| 1 | 1 | 0 | Mark parity |
| 1 | 1 | 1 | Space parity |

21.3.1.9 MODEM Control Registers (UMCR1 and UMCR2)

The UMCRs, shown in [Figure 21-11](#), control the interface with the external peripheral device on the UART bus.




Figure 21-11. Modem Control Register (UMCR1 and UMCR2)

[Table 21-16](#) describes the UMCR fields.

Table 21-16. UMCR Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–2 | — | Reserved, should be cleared |
| 3 | LOOP | Local loopback mode 0 Normal operation. 1 Functionally, the data written to UTHR can be read from URBR of the same UART, and UMCR[RTS] is tied to UMSR[CTS]. |
| 4–5 | — | Reserved |
| 6 | RTS | Ready to send 0 Negates corresponding $\overline{\text{UART_RTS}}$ output. 1 Assert corresponding $\overline{\text{UART_RTS}}$ output. Informs external MODEM or peripheral that the UART is ready for sending/receiving data. |
| 7 | — | Reserved |

21.3.1.10 Line Status Registers (ULSR1 and ULSR2)

The ULSRs, shown in [Figure 21-12](#), monitor the status of the data transfer on the UART buses. To isolate the status bits from the proper character received through the UART bus, software should read the ULSR and then the URBR.




Figure 21-12. Line Status Register (ULSR1 and ULSR2)

[Table 21-17](#) describes the ULSR fields.

Table 21-17. ULSR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0 | RFE | Receiver FIFO error. 0 Cleared when there are no errors in the receiver FIFO or on a read of the ULSR with no remaining receiver FIFO errors. 1 Set when one of the characters in the receiver FIFO encounters an error (framing, parity, or break interrupt). |
| 1 | TEMT | Transmitter empty 0 Either or both the UTHR or the internal transmitter shift register has a data character. In FIFO mode, a data character is in the transmitter FIFO or the internal transmitter shift register. 1 Both the UTHR and the internal transmitter shift register are empty. In FIFO mode, both the transmitter FIFO and the internal transmitter shift register are empty. |
| 2 | THRE | Transmitter holding register empty 0 UTHR is not empty. 1 A data character has transferred from the UTHR into the internal transmitter shift register. In FIFO mode, the transmitter FIFO contains no data character. |
| 3 | BI | Break interrupt 0 Cleared when the ULSR is read or when a valid data transfer is detected (that is, STOP bit is received). 1 Received data of logic 0 for more than START bit + Data bits + Parity bit + one STOP bits length of time. A new character is not loaded until SIN returns to the mark state (logic 1) and a valid START is detected. In FIFO mode, a zero character is encountered in the FIFO (the zero character is at the top of the FIFO). In FIFO mode, only one zero character is stored. |
| 4 | FE | Framing error 0 Cleared when ULSR is read or when a new character is loaded into the URBR from the receiver shift register. 1 Invalid STOP bit for receive data (only the first STOP bit is checked). In FIFO mode, FE is set when the character that detected a framing error is encountered in the FIFO (that is the character at the top of the FIFO). An attempt to resynchronize occurs after a framing error. The UART assumes that the framing error (due to a logic 0 being read when a logic 1 (STOP) was expected) was due to a STOP bit overlapping with the next START bit, so it assumes this logic 0 sample is a true START bit and then will receive the following new data. |
| 5 | PE | Parity error 0 Cleared when ULSR is read or when a new character is loaded into URBR. 1 Unexpected parity value encountered when receiving data. In FIFO mode, the character with the error is at the top of the FIFO. |

Table 21-17. ULSR Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 6 | OE | Overrun error 0 Cleared when ULSR is read 1 Before URBR was read, it was overwritten with a new character. The old character is lost. In FIFO mode, the receiver FIFO is full (regardless of the receiver FIFO trigger level setting) and a new character has been received into the internal receiver shift register. The old character was overwritten by the new character. Data in the receiver FIFO was not overwritten. |
| 7 | DR | Data ready 0 Cleared when URBR is read or when all of the data in the receiver FIFO is read. 1 A character was received in the URBR or the receiver FIFO. |

21.3.1.11 MODEM Status Registers (UMSR1 and UMSR2)

The UMSRs, shown in [Figure 21-13](#), track the status of the MODEM (or external peripheral device) $\overline{\text{CTS}}$, set for the corresponding UART.




Figure 21-13. Modem Status Register (UMSR1 and UMSR2)

[Table 21-18](#) describes UMSR fields.

Table 21-18. UMSR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0-2 | — | Reserved, should be cleared |
| 3 | CTS | Clear to send. Represents the inverted value of the $\overline{\text{CTS}}$ input pin from the external peripheral device. 0 Corresponding $\overline{\text{CTS}}_n$ is negated. 1 Corresponding $\overline{\text{CTS}}_n$ is asserted. The MODEM or peripheral device is ready for data transfers. |
| 4-6 | — | Reserved, should be cleared |
| 7 | DCTS | Delta clear to send 0 No change on the corresponding $\overline{\text{CTS}}_n$ signal since the last read of UMSR[CTS]. 1 $\overline{\text{CTS}}_n$ changed since the last read of UMSR[CTS]. Causes an interrupt if UIER[EMSI] is set to detect this condition. |

21.3.1.12 Scratch Registers (USCR1 and USCR2)

USCR, shown in [Figure 21-14](#), are for debugging software or the DUART hardware. The USCRs do not affect the operation of the DUART.



Figure 21-14. Scratch Register (USCR)

[Table 21-19](#) describes USCR fields.

Table 21-19. USCR Field Descriptions

| Bits | Name | Description |
|------|------|-------------|
| 0–7 | DATA | Data |

21.3.1.13 DMA Status Registers (UDSR1 and UDSR2)

The DMA status registers (UDSRs), shown in [Figure 21-15](#), return transmitter and receiver FIFO status and provide the ability to assist DMA data operations to and from the FIFOs.



Figure 21-15. DMA Status Register (UDSR)

[Table 21-20](#) describes the fields of the UDSRs.

Table 21-20. UDSR Field Descriptions

| Bits | Name | Description |
|------|-------|---|
| 0–5 | — | Reserved |
| 6 | TXRDY | Transmitter ready. Reflects the status of the transmitter FIFO or the UTHR. The status depends on the DMA mode selected, which is determined by UFCR[DMS] and UFCR [FEN]. 0 The bit is cleared, as shown in Table 21-22 . 1 This bit is set, as shown in Table 21-21 . |
| 7 | RXRDY | Receiver ready. This read-only bit reflects the status of the receiver FIFO or URBR. The status depends on the DMA mode selected, which is determined by UFCR[DMS] and UFCR [FEN]. 0 The bit is cleared, as shown in Table 21-24 . 1 This bit is set, as shown in Table 21-23 . |

Table 21-21. UDSR[TXRDY] Set Conditions

| DMS | FEN | DMA Mode | Meaning |
|-----|-----|----------|---|
| 0 | 0 | 0 | TXRDY is set after the first character is loaded into the transmitter FIFO or UTHR. |
| 0 | 1 | 0 | |
| 1 | 0 | 0 | |
| 1 | 1 | 1 | TXRDY is set when the transmitter FIFO is full. |

Table 21-22. UDSR[TXRDY] Cleared Conditions

| DMS | FEN | DMA Mode | Meaning |
|-----|-----|----------|--|
| 0 | 0 | 0 | TXRDY is cleared when there are no characters in the transmitter FIFO or UTHR. |
| 0 | 1 | 0 | |
| 1 | 0 | 0 | |
| 1 | 1 | 1 | TXRDY is cleared when there are no characters in the transmitter FIFO or UTHR. TXRDY remains clear while the transmitter FIFO is not yet full. |

Table 21-23. UDSR[RXRDY] Set Conditions

| DMS | FEN | DMA Mode | Meaning |
|-----|-----|----------|--|
| 0 | 0 | 0 | RXRDY is set when there are no characters in the receiver FIFO or URBR. |
| 0 | 1 | 0 | |
| 1 | 0 | 0 | |
| 1 | 1 | 1 | RXRDY is set when the trigger level has not been reached and there has been no time out. |

Table 21-24. UDSR[RXRDY] Cleared

| DMS | FEN | DMA Mode | Meaning |
|-----|-----|----------|---|
| 0 | 0 | 0 | RXRDY is cleared when there is at least one character in the receiver FIFO or URBR. |
| 0 | 1 | 0 | |
| 1 | 0 | 0 | |
| 1 | 1 | 1 | RXRDY is cleared when the trigger level or a time-out has been reached. RXRDY remains cleared until the receiver FIFO is empty. |

21.4 Functional Description

The communication channel provides a full-duplex asynchronous receiver and transmitter using an operating frequency derived from the system clock signal.

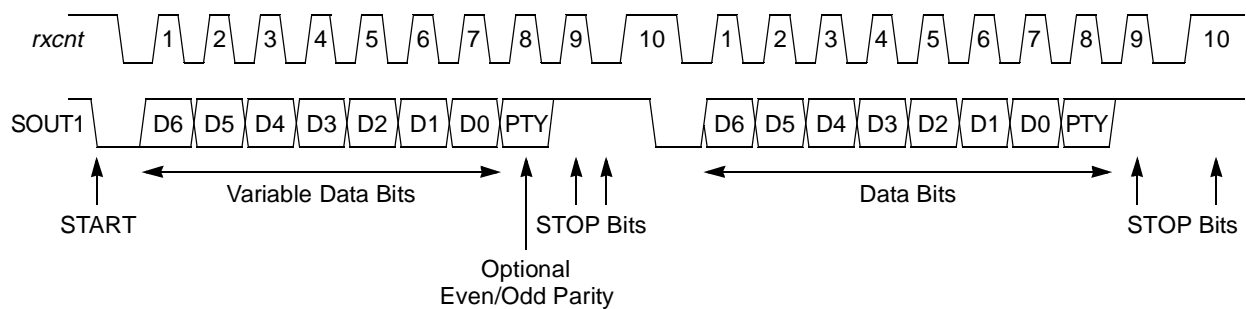
The transmitter accepts parallel data with a write access to UTHR. In FIFO mode, the data is placed directly into an internal transmitter shift register, or into the transmitter FIFO—see [Section 21.4.5, “FIFO Mode.”](#) The transmitting registers convert the data to a serial bit stream by inserting the appropriate

START, STOP, and optional parity bits. Finally, the registers output a composite serial data stream on the channel transmitter serial data output (SOUT). The transmitter status may be polled or interrupt driven.

The receiver accepts serial data on the channel receiver serial data input (SIN), converts the data into parallel format, and checks for START, STOP, and parity bits. In FIFO mode, the receiver removes the START, STOP, and parity bits and then transfers the assembled character from the receiver buffer, or receiver FIFO. This transfer occurs in response to a read of the UART receiver buffer register (URBR). The receiver status may be polled or interrupt-driven.

21.4.1 Serial Interface

The UART bus is a serial, full-duplex, point-to-point bus as shown in [Figure 21-16](#). Therefore, only two devices are attached to the same signals and there is no need for address or arbitration bus cycles.



Two 7-Bit Data Transmissions with Parity and 2-Bit STOP Transactions

Figure 21-16. UART Bus Interface Transaction Protocol Example

A standard UART bus transfer is composed of either three or four parts:

- START bit
- Data transfer (least significant bit is first data bit on the bus)
- Parity bit (optional)
- STOP bits

An internal logic sample signal, *rxcnt*, uses the frequency of the baud-rate generator to drive the bits on SOUT.

The following sections describe the four components of the serial interface, the baud-rate generator, local loopback mode, different errors, and FIFO mode.

21.4.1.1 START Bit

A write to UTHR generates a START bit on the SOUT signal. [Figure 21-16](#) shows that the START bit is defined as a logic 0. The START bit denotes the beginning of a new data transfer which is limited to the bit length programmed in ULCR. When the bus is idle, SOUT is high.

21.4.1.2 Data Transfer

Each data transfer contains 5, 6, 7, or 8 bits of data. The ULCR data bit length for the transmitter and receiver UART devices must agree before a transfer begins; otherwise, a parity or framing error may occur. A transfer begins when UTHR is written. At that time, a START bit is generated followed by 5 to 8 of the data bits previously written to the UTHR. The data bits are driven from the least- to the most-significant bits. After the parity and STOP bits, a new data transfer can begin if new data is written to UTHR.

21.4.1.3 Parity Bit

The user has the option of using even, odd, no parity, or stick parity (see [Section 21.3.1.8, “Line Control Registers \(ULCR1 and ULCR2\)”](#)). Both the receiver and transmitter parity definitions must agree before transferring data. When receiving data, a parity error can occur if an unexpected parity value is detected (see [Section 21.3.1.10, “Line Status Registers \(ULSR1 and ULSR2\)”](#)).

21.4.1.4 STOP Bit

The transmitter device ends the write transfer by generating a STOP bit. The STOP bit is always high. The user can program the length of the STOP bit(s) in the ULCR. Both the receiver and transmitter STOP bit length must agree before attempting to transfer data. A framing error can occur if an invalid STOP bit is detected.

21.4.2 Baud-Rate Generator Logic

Each UART contains an independent programmable baud-rate generator, that is capable of taking the system clock input and dividing the input by any divisor from 1 to $2^{16} - 1$.

The baud rate is defined as the number of bits per second that can be sent over the UART bus. The formula for calculating baud rate is as follows:

$$\text{Baud rate} = (1/16) \times (\text{system clock frequency/divisor value})$$

Therefore, the output frequency of the baud-rate generator is 16 times the baud rate.

- The divisor value is determined by the following two 8-bit registers to form a 16-bit binary number:
- UART divisor most significant byte register (UDMB)
- UART divisor least significant byte register (UDLB)

Upon loading either of the divisor latches, a 16-bit baud-rate counter is loaded.

The divisor latches must be loaded during initialization to ensure proper operation of the baud-rate generator. Both UART devices on the same bus must be programmed for the same baud rate before starting a transfer.

The baud clock can be passed to the performance monitor by enabling UAFR[BO]. This can be used to determine baud-rate errors.

21.4.3 Local Loopback Mode

Local loopback mode is provided for diagnostic testing. The data written to UTHR can be read from the receiver buffer register (URBR) of the same UART. In this mode, the MODEM control register UMCR[RTS] is internally tied to the MODEM status register UMSR[CTS]. The transmitter SOUT is set to a logic 1 and the receiver SIN is disconnected. The output of the transmitter shift register is looped back into the receiver shift register input. The $\overline{\text{CTS}}$ (input signal) is disconnected, $\overline{\text{RTS}}$ is internally connected to $\overline{\text{CTS}}$, and the $\overline{\text{RTS}}$ (output signal) becomes inactive. In this diagnostic mode, data that is transmitted is immediately received. In local loopback mode the transmit and receive data paths of the DUART can be verified. Note that in local loopback mode, the transmit/receive interrupts are fully operational and can be controlled by the interrupt enable register (UIER).

21.4.4 Errors

The following sections describe framing, parity, and overrun errors which may occur while data is transferred on the UART bus. Each of the error bits are usually cleared, as described below, when the line status register (ULSR) is read.

21.4.4.1 Framing Error

When an invalid STOP bit is detected, a framing error occurs and ULSR[FE] is set. Note that only the first STOP bit is checked. In FIFO mode, ULSR[FE] is set when the character at the top of the FIFO detects a framing error. An attempt to re-synchronize occurs after a framing error. The UART assumes that the framing error (due to a logic 0 being read when a logic 1 (STOP) was expected) was due to a STOP bit overlapping with the next START bit. ULSR[FE] is cleared when ULSR is read or when a new character is loaded into the URBR from the receiver shift register.

21.4.4.2 Parity Error

When unexpected parity values are encountered while receiving data, a parity error occurs and ULSR[PE] is set. In FIFO mode, ULSR[PE] is set when the character with the error is at the top of the FIFO. ULSR[PE] is cleared when ULSR is read or when a new character is loaded into the URBR.

21.4.4.3 Overrun Error

When a new (overwriting character) STOP bit is detected and the old character is lost, an overrun error occurs and ULSR[OE] is set. In FIFO mode, ULSR[OE] is set after the receiver FIFO is full (despite the receiver FIFO trigger level setting) and a new character has been received into the internal receiver shift register. Data in the FIFO is not overwritten; only the shift register data is overwritten. Therefore, the interrupt occurs immediately. ULSR[OE] is cleared when ULSR is read.

21.4.5 FIFO Mode

The UARTs use an alternate mode (FIFO mode) to relieve the processor core from excessive software overhead. The FIFO control register (UFCCR) is used to enable and clear the receiver and transmitter FIFOs

and set the FIFO receiver trigger level UFCR[RTL] to control the received data available interrupt UIER[ERDAI].

The UFCR also selects the type of DMA signaling. The UDSR[RXRDY] indicates the status of the receiver FIFO. UDSR[TXRDY] indicate when the transmitter FIFO is full. When in FIFO mode, data written to UTHR is placed into the transmitter FIFO. The first byte written to UTHR is the first byte onto the UART bus.

21.4.5.1 FIFO Interrupts

In FIFO mode, the UIER[ERDAI] is set when a time-out interrupt occurs. A receive data time-out generates a maskable interrupt condition (through UIER[ERDAI]). See [Section 21.3.1.4, “Interrupt Enable Registers \(UIER1 and UIER2\)”](#).

UIIR indicates whether the FIFOs are enabled. UIIR[IID3] is set only for FIFO mode interrupts. The character time-out interrupt occurs when no characters have been removed from or input to the receiver FIFO during the last four character times and at least one character is in the receiver FIFO. The character time-out interrupt (controlled by UIIR[IID n]) is cleared when URBR is read. See [Section 21.3.1.5, “Interrupt ID Registers \(UIIR1 and UIIR2\)”](#).

UIIR[FE] indicates whether FIFO mode is enabled.

21.4.5.2 DMA Mode Select

UDSR[RXRDY] reflects the status of the receiver FIFO or URBR. In mode 0 (UFCR[DMS] is cleared), UDSR[RXRDY] is cleared when at least one character is in the receiver FIFO or URBR; it is set when there are no more characters in the receiver FIFO or URBR. This occurs regardless of the UFCR[FEN] setting. In mode 1 (UFCR[DMS] and UFCR[FEN] are set), UDSR[RXRDY] is cleared when the trigger level or a time-out has been reached; it is set when there are no more characters in the receiver FIFO.

UDSR[TXRDY] reflects the status of the transmitter FIFO or UTHR. In mode 0 (UFCR[DMS] is cleared), UDSR[TXRDY] is cleared when there are no characters in the transmitter FIFO or UTHR; it is set after the first character is loaded into the transmitter FIFO or UTHR. This occurs regardless of the UFCR[FEN] setting. In mode 1 (UFCR[DMS] and UFCR[FEN] are set), UDSR[TXRDY] is cleared when there are no characters in the transmitter FIFO or UTHR; it is set when the transmitter FIFO is full.

See [Section 21.3.1.13, “DMA Status Registers \(UDSR1 and UDSR2\)”](#) for a complete description of the UDSR[RXRDY] and UDSR[TXRDY] bits.

21.4.5.3 Interrupt Control Logic

An interrupt is active when DUART interrupt ID register bit 7 (UIIR[IID0]), is cleared. UIER is used to mask specific interrupt types. See [Section 21.3.1.4, “Interrupt Enable Registers \(UIER1 and UIER2\)”](#).

When the interrupts are disabled in UIER, polling software can not use UIIR[IID0] to determine whether the UART is ready for service. Software must monitor the appropriate ULSR and UMSR bits. UIIR[IID0] can be used for polling if the interrupts are enabled in UIER.

21.5 DUART Initialization/Application Information

The following requirements must be met for DUART accesses:

- All DUART registers must be mapped to a cache-inhibited and guarded area. (That is, the WIMG setting in the MMU needs to be 0b01x1.)
- All DUART registers are 1 byte wide. Reads and writes to these registers must be byte-length operations.

A system reset puts the DUART registers to a default state. Before the interface can transfer serial data, the following initialization steps are recommended:

1. Update the programmable interrupt controller (PIC) DUART channel interrupt vector source registers.
2. Set data attributes and control bits in the ULCR, UFCR, UAFR, UMCR, UDLB, and UDMB.
3. Set the data attributes and control bits of the external MODEM or peripheral device.
4. Set the interrupt enable register (UIER).
5. To start a write transfer, write to the UTHR.
6. Poll UIIR if the interrupts generated by the DUART are masked.



Chapter 22

Serial Peripheral Interface

22.1 Overview

The serial peripheral interface (SPI) allows the device to exchange data between other PowerQUICC® family chips, the MC68360, MC68302, M68HC11, and M68HC05 microcontroller families, and other family devices. The SPI can be used to communicate with peripheral devices such as EEPROMs, real-time clocks, A/D converters, and ISDN devices.

The SPI is a full-duplex, synchronous, character-oriented channel that supports a four-wire interface (receive, transmit, clock, and slave select). The SPI block consists of transmitter and receiver sections, an independent baud-rate generator, and a control unit. The transmitter and receiver sections use the same clock, which is derived from the SPI baud rate generator in master mode or externally in slave mode. During an SPI transfer, data is sent and received simultaneously.

The SPI receiver and transmitter are double-buffered, as shown in [Figure 22-1](#), giving an effective FIFO size (latency) of 2 characters. The SPI's MSB/LSB is shifted out first. When the SPI is disabled in the SPI mode register (SPMODE[EN] = 0), it consumes little power.

22.2 Introduction

The SPI block diagram is shown in [Figure 22-1](#).




Figure 22-1. SPI Block Diagram

22.2.1 Features

The major features of the SPI are listed as follows:

- Four-signal interface (SPIMOSI, SPIMISO, SPICLK, and $\overline{\text{SPISSEL}}$)
- Full-duplex operation
- Works with 32-bit data characters or with a range from 4-bit to 16-bit data characters
- Supports back-to-back character transmission and reception
- Supports reverse data mode for 8/16/32 character length
- Supports master SPI mode
- Supports multiple-master environment
- Maximum clock rate is (input clock rate/4) in master mode; (input clock rate/2) in slave mode
- Independent programmable baud rate generator
- Programmable clock phase and polarity
- Local loopback capability for testing
- Open-drain outputs support multiple-master configuration

22.2.2 SPI Transmission and Reception Process

Because the SPI is a character-oriented communication unit, the core is responsible for packing and unpacking the receive and transmit frames. A frame consists of all of the characters transmitted or received during a completed SPI transmission session, from the first character written to the SPITD register to the last character transmitted following the setting of SPCOM[LST]. See [Section 22.4.1.4, “SPI Command Register \(SPCOM\),”](#) for more information.

The core receives data by reading the SPI receive data hold register (SPIRD). The SPI then clears the not empty SPIE[NE] to free up the SPIRD register for the next receive operation. The core transmits data by writing it into the SPI transmit data hold register (SPITD). The SPI then clears the not full (NF) bit in the SPI event register (SPIE) to indicate that the SPITD register contains a character for transmission. When the next character to be transmitted is going to be the final one in the current frame, the core sets SPCOM[LST], and then writes the final character to SPITD.

The SPI core handshake protocol can be implemented by either using polling or interrupts. When using a polling, the core reads the SPIE in a predefined frequency and acts according to the value of the SPIE bits. The polling frequency depends on the SPI serial channel frequency. When using the interrupt mechanism, setting either the not full (NF) or not empty (NE) bits of SPIE causes an interrupt to the processor core. The core then reads SPIE and acts accordingly. The three basic modes of operation for transmitting and receiving are master, slave and multiple-master.

NOTE

When both NE and NF bits are set, the processor core should read the received data before transmitting new data.

The SPMODE[LEN] determines the character length sent by the hardware. The core is responsible for any bit manipulation to pack/unpack data into the appropriate character length. See the SPMODE[LEN] description in [Table 22-4](#) for more information.

22.2.3 Modes of Operation

The SPI can be programmed to work in a single- or multiple-master environment. This section describes SPI master and slave operations in a single-master configuration. It also discusses the multiple master environment.

The following sections summarize the main modes of operation that the SPI supports.

22.2.3.1 SPI as a Master Device

In master mode, the SPI sends a message to the slave peripheral, which sends back a simultaneous reply. A single master device with multiple slaves can use general-purpose parallel I/O signals to selectively enable slaves, as shown in [Figure 22-2](#). To eliminate the multi-master error in a single-master environment, the master's $\overline{\text{SPISEL}}$ input should be forced inactive by an external pull up.

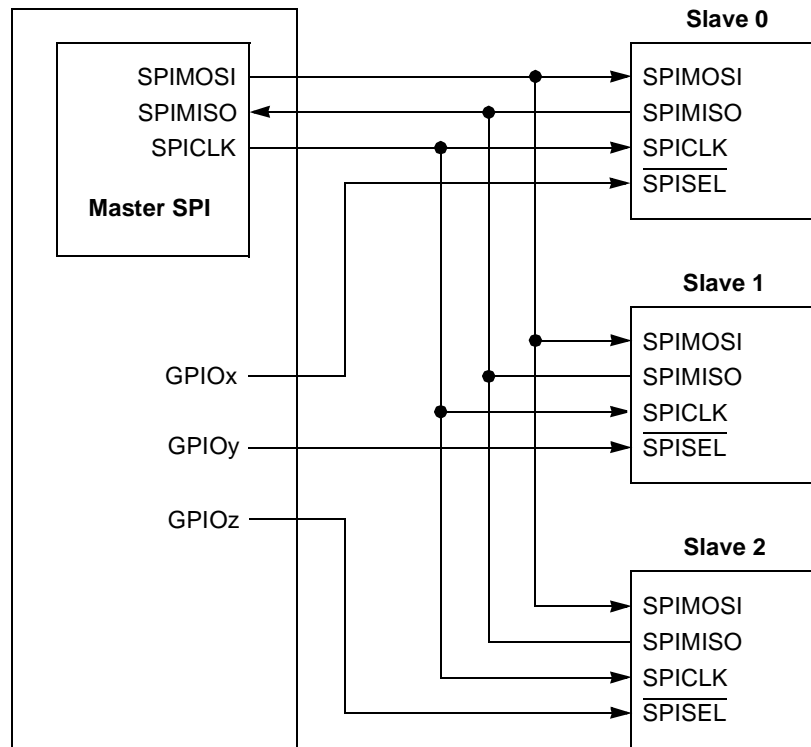


Figure 22-2. Single-Master/Multi-Slave Configuration

To start exchanging data, the processor core writes the data to be sent into the SPITD register. The SPI then generates programmable clock pulses on SPICLK for each character. It shifts Tx data out on the SPI master-out slave-in (SPIMOSI) and Rx data in on the SPI master-in slave-out (SPIMISO) simultaneously. During transmission, the core is responsible for supplying the data whenever the SPI requests it to ensure smooth operation. After the last data (LST command and data afterwards), the first character written to SPITD acts as a start command for the SPI.

The SPI continues transmitting and receiving characters until SPCOM[LST] is set or an error occurs.

The SPI sets SPIE[NF] to issue a maskable interrupt to the interrupt controller whenever its transmit buffer is not full. It also sets the NF bit after sending the last word. In response, the core should read the exception flags that relate to the last word. The SPI sets SPIE[NE] to issue a maskable interrupt to the interrupt controller whenever the receiver buffer has been filled with data.

22.2.3.2 SPI as a Slave Device

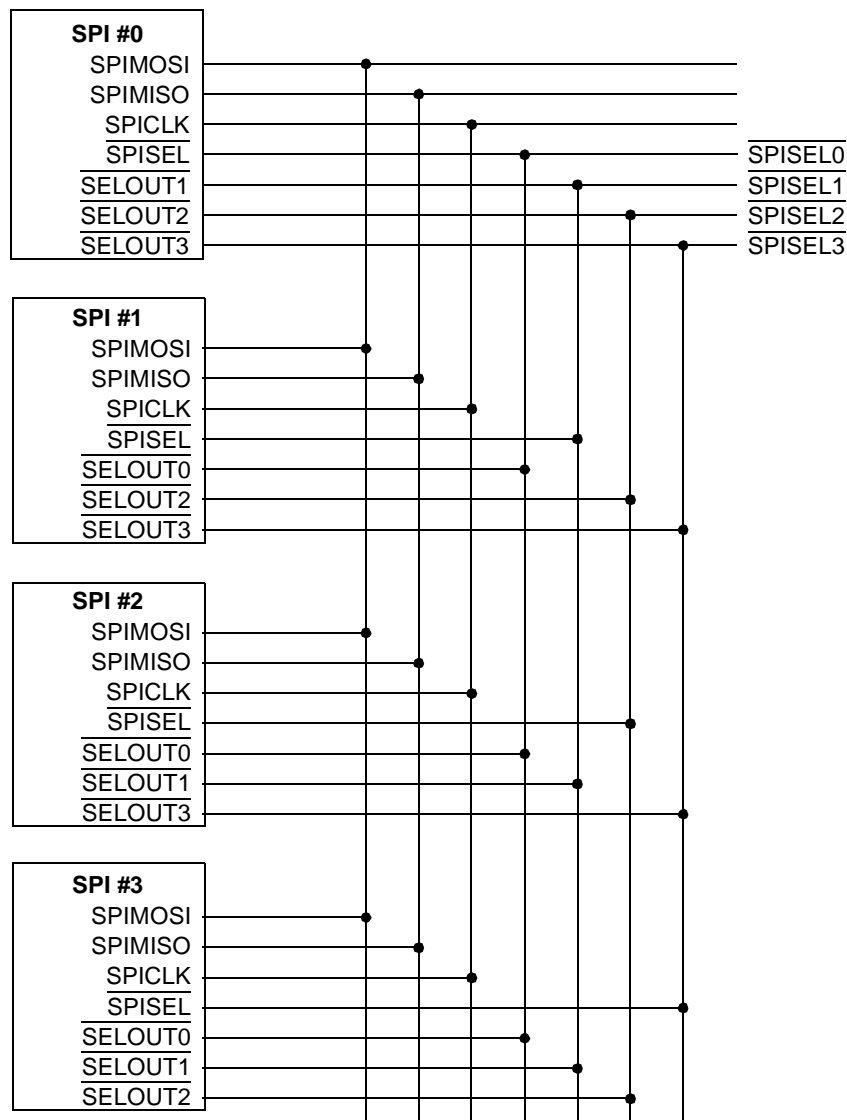
In slave mode, the SPI receives messages from an SPI master and sends a simultaneous reply. The slave's $\overline{\text{SPISEL}}$ must be asserted before Rx clocks are recognized. Once $\overline{\text{SPISEL}}$ is asserted, SPICLK becomes an input from the master to the slave. SPICLK can be any frequency from DC to input clock/2.

To prepare for data transfers, the core writes data to be sent into the SPITD register. Once $\overline{\text{SPISEL}}$ is asserted, the slave shifts data out from SPIMISO and in through SPIMOSI. The SPI sets the NF bit of the SPIE register and a maskable interrupt is issued when a full buffer finishes receiving and sending or after an error. The SPI continues reception until $\overline{\text{SPISEL}}$ is negated.

Transmission continues until no more data is available or $\overline{\text{SPISEL}}$ is negated. Transmission continues once $\overline{\text{SPISEL}}$ is reasserted and SPICLK begins toggling. After the characters in the buffer are sent, the SPI sends one as long as $\overline{\text{SPISEL}}$ remains asserted.

22.2.3.3 SPI in Multiple-Master Operation

The SPI can operate in a multiple-master environment in which all SPI devices are connected to the same bus. In this configuration, the SPIMOSI , SPIMISO , and SPICLK signals of all SPIs are shared; but the $\overline{\text{SPISEL}}$ inputs are connected separately, as shown in [Figure 22-13](#). Only one SPI device can act as master at a time—all others must be slaves. When a SPI is configured as a master, if its $\overline{\text{SPISEL}}$ input is asserted, a multiple-master error occurs because more than one SPI device is a bus master. The SPI sets $\text{SPIE}[\text{MME}]$ in the SPI event register and a maskable interrupt is issued to the core. It also disables SPI operation and the output drivers of the SPI signals. The core must clear $\text{SPMODE}[\text{EN}]$, correct the problems, and clear $\text{SPIE}[\text{MME}]$ before the SPI can be used again.



Notes:

1. All signals are open-drain.
2. For a multiple-master configuration with more than two masters, $\overline{\text{SPISEL}}$ and SPIE[MME] do not detect all possible conflicts.
3. It is the responsibility of software to arbitrate for the SPI bus (with token passing, for example).
4. SELOUTx signals are implemented in software with general-purpose I/O signals.

Figure 22-3. Multiple-Master Configuration

The SPI can transfer a single character at much higher rates—input clock/4 in master mode and input clock/2 in slave mode, and subjected to the timing parameters of the interconnected devices, and board trace delays. Gaps should be inserted between multiple characters to keep from exceeding the maximum sustained data rate.

22.3 External Signal Descriptions

The SPI's four wire interface consists of transmit, receive, clock, and slave select.

22.3.1 Overview

Table 22-1 lists signal properties.

Table 22-1. Signal Properties

| Name | Function | Reset | Pull Up |
|----------------------------|---|-------|-----------------------------|
| SPIMISO | Master input slave output | — | Required in open drain mode |
| SPIMOSI | Master output slave input | — | Required in open drain mode |
| SPICLK | Input/output serial clock connected to the other SPICLK | — | Required in open drain mode |
| $\overline{\text{SPISEL}}$ | SPI slave select | — | Required in open drain mode |

22.3.2 Detailed Signal Descriptions

Table 22-2 describes the signals in detail.

Table 22-2. Detailed Signal Descriptions

| Signal | I/O | Description | |
|---------|-----|---------------------------|---|
| SPIMISO | I/O | Master input slave output | |
| | | State Meaning | Asserted—The data that has been transmitted/received from/to the SPI (depends if master or slave mode) is high Negated—The data that has ben transmitted/received from/to the SPI (depends if master or slave mode) is low |
| | | Timing | Assertion—According to the SPICLK assertion/negation/in the middle of phase (depends on SPMODE) Negation—According to the SPICLK assertion/negation/in the middle of phase (depends on SPMODE) |
| SPIMOSI | I/O | Master output slave input | |
| | | State Meaning | Asserted—The data that has been transmitted/received from/to the SPI (depends if master or slave mode) is high Negated—The data that has ben transmitted/received from/to the SPI (depends if master or slave mode) is low |
| | | Timing | Assertion—According to the SPICLK assertion/negation/in the middle of phase (depends on SPMODE) Negation—According to the SPICLK assertion/negation/in the middle of phase (depends on SPMODE) |

Table 22-2. Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|----------------------------|-----|---|---|
| SPICLK | I/O | Serial clock in or serial clock out for slave or master mode respectively | |
| | | State Meaning | Assertion/Negation according to SPMODE[PM, DIV16] register rate configuration |
| | | Timing | Assertion/Negation—during frame reception/transmission |
| $\overline{\text{SPISEL}}$ | I | SPI slave select | |
| | | State Meaning | Asserted—In slave mode declares the slave has been selected for the coming frame. In master mode assertion causes MME multiple-master error. Negated—In slave mode means the specific SPI has not been selected. In master mode needs to be negated for regular operation. |
| | | Timing | Assertion—In slave mode along with the data from the slave Negation—In slave mode with the end of the frame (according to SPMODE[LEN]). In master mode before data is first written to SPITD and remains constant. |

The SPI can be configured as a slave or a master in single- or multiple-master environments mode. The master SPI generates the transfer clock SPICLK using the SPI baud rate generator (BRG). The SPI BRG takes its input from input clock, which is generated in the device clock synthesizer.

SPICLK is a gated clock, active only during data transfers. Four combinations of SPICLK phase and polarity can be configured with the clock invert (SPMODE[CI]) and clock phase (SPMODE[CP]) register bits. SPI signals can also be configured as open-drain to support a multiple-master configuration in which a shared SPI signal is driven by the device or an external SPI device.

The SPI master-in slave-out SPIMISO signal acts as an input for master devices and as an output for slave devices. Conversely, the master-out slave-in SPIMOSI signal is an output for master devices and an input for slave devices. The dual functionality of these signals allows the SPIs in a multiple-master environment to communicate with one another using a common hardware configuration.

- When the SPI is a master, SPICLK is the clock output signal that shifts received data in from SPIMISO and transmitted data out to SPIMOSI. SPI masters must output a slave select signal to enable SPI slave devices by using a separate general-purpose I/O signal. Assertion of the $\overline{\text{SPISEL}}$ while the SPI is configured as a master causes an error.
- When the SPI is a slave, SPICLK is the clock input that shifts received data in from SPIMOSI and transmitted data out through SPIMISO. $\overline{\text{SPISEL}}$ is the enable input to the SPI slave. In a multiple-master environment, $\overline{\text{SPISEL}}$ (always an input) is also used to detect an error when more than one master is operating.

22.4 Memory Map/Register Definition

Table 22-3 shows the memory-mapped registers of the SPI and their offsets. It lists the offset, name, and a cross-reference to the complete description of each register. Note that the full register address is comprised of IMMRBAR together with the SPI block base address and offset listed in Table 22-3. Undefined 4-byte address spaces within offset 0x000–0xFFF are reserved.

Table 22-3. SPI Register Summary

| Offset | Register | Access | Reset Value | Section/Page |
|--|-------------------------------|--------|-------------|----------------------------------|
| Serial Peripheral Interface (SPI)—Block Base Address 0x0_7000 | | | | |
| 0x000–0x01F | Reserved | — | — | — |
| 0x020 | SPI mode register (SPMODE) | R/W | All zeros | 22.4.1.1/2222-9 |
| 0x024 | SPI event register (SPIE) | Mixed | All zeros | 22.4.1.2/2222-12 |
| 0x028 | SPI mask register (SPIM) | R/W | All zeros | 22.4.1.3/2222-13 |
| 0x02C | SPI command register (SPCOM) | W | All zeros | 22.4.1.4/2222-14 |
| 0x030 | SPI transmit register (SPITD) | W | All zeros | 22.4.1.5/2222-14 |
| 0x034 | SPI receive register (SPIRD) | R | 0xFFFF_FFFF | 22.4.1.6/2222-15 |
| 0x038–0xFFFF | Reserved | — | — | — |

22.4.1 Register Descriptions

22.4.1.1 SPI Mode Register (SPMODE)

SPMODE, shown in [Figure 22-4](#), controls both the SPI operation mode and clock source.

Offset 0x020

Access: Read/write

| | | | | | | | | | | | | | | | | | | | |
|-------|---|------|----|----|-------|-----|-----|----|-----|----|----|----|----|----|----|----|--|--|-----------|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 11 | 12 | 15 | 16 | 18 | 19 | 20 | | | 31 |
| R | — | LOOP | CI | CP | DIV16 | REV | M/S | EN | LEN | | PM | | — | OD | | | | | |
| W | — | | | | | | | | | | | | | | | | | | |
| Reset | | | | | | | | | | | | | | | | | | | All zeros |

Figure 22-4. SPMODE-SPI Mode Register Definition

[Table 22-4](#) describes the SPMODE fields.


Table 22-4. SPMODE Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0 | — | Reserved |
| 1 | LOOP | Loop mode. Enables local loopback operation. 0 Normal operation. 1 Loopback mode. Used to test the SPI controller internal functionality, the transmitter output is internally connected to the receiver input. The receiver and transmitter operate normally, except that received data is ignored. The SPI acts normally in loop back mode; therefore, negating SPISEL in slave mode stops transmission, negating it in master mode and causing an MME error. |
| 2 | CI | Clock invert. Inverts SPI clock polarity. See Figure 22-5 and Figure 22-6 for more information. 0 The inactive state of SPICLK is low. 1 The inactive state of SPICLK is high. |

Table 22-4. SPMODE Field Descriptions (continued)

| Bits | Name | Description |
|-------|-------|--|
| 3 | CP | Clock phase. Selects the transfer format. See Figure 22-5 and Figure 22-6 for more information. 0 SPICLK starts toggling at the middle of the data transfer. 1 SPICLK starts toggling at the beginning of the data transfer. |
| 4 | DIV16 | Divide by 16. Selects the clock source for the SPI baud rate generator (SPI BRG) when configured as an SPI master. In slave mode, SPICLK is the clock source. 0 The SPI block input clock is the input to the SPI BRG. 1 The SPI block input clock/16 is the input to the SPI BRG. In slave mode this bit must be cleared. |
| 5 | REV | Reverse data mode for 8-/16-/32-bit character length only (see Section 22.4.1.6.1, “Reverse Mode SPMODE[REV] Examples.”) 0 LSB sent/received first (for data LEN < 32 the data is located at the lower half-word LSB) 1 MSB sent/received first |
| 6 | M/S | Master/slave. Selects master or slave mode. 0 The SPI is a slave. 1 The SPI is a master. |
| 7 | EN | Enable SPI. Any other bits in SPMODE must not change when EN is set. 0 The SPI is disabled. The SPI is in a idle state and consumes minimal power. The SPI BRG is not functioning and the input clock is disabled. 1 The SPI is enabled. Note: The SPI controller requires a minimal gap of at least 10 input clocks between disabling the SPI and re-enabling. This minimal gap is sufficient provided that SPMODE[PM] and SPMODE[DIV16] are cleared during the time in which SPMODE[EN] is cleared. |
| 8–11 | LEN | Character length in bits per character. LEN can be either 32-bits, or 4- to 16-bits that are shown as follows: 0000 32-bit characters 0001–0010 Reserved, causes erratic behavior. 0011 4-bit characters ... 1111 16-bit characters The TX and RX registers (SPITD, SPIRD) hold 32 bits at a time. A character length of 32 bits fills the TX and RX registers; therefore, all of the bits in these registers are valid. However, if the character length selected by LEN is equal or less than 16 bits, then the valid bits will reside in the lower half-word of the transmit and receive registers. For example, if the character length is set to 16 bits than the valid bits will be 16–31, if the character length is set to 5 bits that the valid bits will be 16–20. Note that the transmit and receive registers each can hold only one character regardless of the character length. |
| 12–15 | PM | Prescale modulus select. Specifies the divide ratio of the prescale divider in the SPI clock generator. The SPI baud rate generator clock source (either input clock or input clock divided by 16, depending on DIV16 bit) is divided by $4 \times ([PM] + 1)$, a range from 4 to 64. The clock has a 50% duty cycle. For example, if the prescale modulus is set to PM = 0011 and DIV16 is set, the system/SPICLK clock ratio will be $16 \times (4 \times (0011 + 1)) = 256$. In slave mode this field must be cleared. |
| 16–18 | — | Reserved |
| 19 | OD | Open drain mode. 0 All output pins are configured to normal mode. 1 All output pins are configured to open drain mode. |
| 20–31 | — | Reserved |

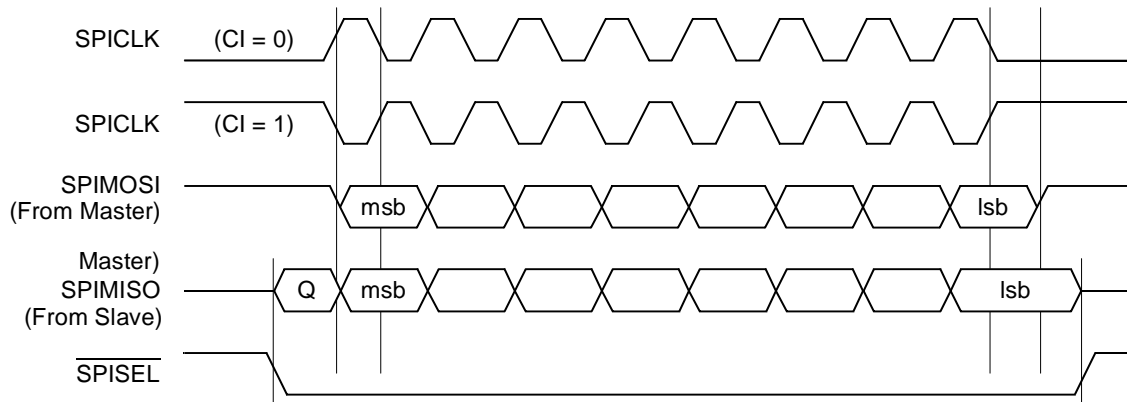
Figure 22-5 shows the SPI transfer format in which SPICLK starts toggling in the middle of the transfer (SPMODE[CP] = 0).



NOTE: Q = Undefined signal.

Figure 22-5. SPI Transfer Format with SPMODE[CP] = 0

Figure 22-6 shows the SPI transfer format in which SPICLK starts toggling at the beginning of the transfer (SPMODE[CP] = 1).



NOTE: Q = Undefined signal.

Figure 22-6. SPI Transfer Format with SPMODE[CP] = 1

22.4.1.2 SPI Event Register (SPIE)

The SPI event register (SPIE) generates interrupts and reports events recognized by the SPI. When an event is recognized, the SPI sets the corresponding SPIE bit. Most SPIE bits can be cleared by writing a '1'. Writing '0' has no effect. Setting a bit in the SPI mask register (SPIM) enables, and clearing a bit

SPIM bit enables and clearing a SPIM bit masks the corresponding interrupt. Unmasked SPIE bits must be cleared before the core clears its internal interrupt requests.

Offset 0x028

Access: Read/write




Figure 22-8. SPIM—SPI Mask Register Definition

Table 22-6 describes the SPIM fields.

Table 22-6. SPIM Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–16 | — | Reserved |
| 17 | LT | Last character transmitted 0 LT event will not cause an SPI interrupt 1 LT event causes an SPI interrupt |
| 18 | DNR | In slave mode data not ready 0 Slave DNR event will not cause an SPI interrupt Note: 1 Slave DNR event causes an SPI interrupt |
| 19 | OV | Slave/Master Overrun interrupt mask 0 Slave/Master Overrun event will not cause an SPI interrupt 1 Slave/Master Overrun event causes an SPI interrupt |
| 20 | UN | Slave Underrun interrupt mask 0 Slave Underrun event will not cause an SPI interrupt 1 Slave Underrun event causes an SPI interrupt |
| 21 | MME | Multimaster error interrupt mask 0 Multimaster error event will not cause an SPI interrupt 1 Multimaster error event causes an SPI interrupt |
| 22 | NE | Not Empty interrupt mask 0 Not Empty event will not cause an SPI interrupt 1 Not Empty event causes an SPI interrupt |
| 23 | NF | Not Full interrupt mask 0 Not Full event will not cause an SPI interrupt 1 Not Full event causes an SPI interrupt |
| 24–31 | — | Reserved |

22.4.1.4 SPI Command Register (SPCOM)

The SPI command register (SPCOM), shown in [Figure 22-9](#), is used to end SPI operation.

Offset 0x02C

Access: Write only




Figure 22-9. SPI Command Register Definition

[Table 22-7](#) describes the SPCOM fields.

Table 22-7. SPCOM Field Descriptions

| Bits | Name | Description |
|-------|------|---|
| 0–8 | — | Reserved |
| 9 | LST | This bit represents the last character. Should be set before the last character is written to the SPITD. This results in SPIE[LT] being set when the character is fully transmitted and by that gives indication about the frame being fully transmitted. 0 This character is not the last character of the frame 1 This character is the last character of the frame |
| 10–31 | — | Reserved |

22.4.1.5 SPI Transmit Data Hold Register (SPITD)

SPITD holds the character to be transmitted. The number of bits in each character is specified by SPMODE[LEN]. Each time SPIE[NF] is set, the core can write another character of data to SPITD, if there is no error indication in the SPIE. At the end of the frame the core should set SPCOM[LST] and prepare the last character of data. [Figure 22-10](#) shows the SPI transmit data hold register.

Offset 0x030

Access: Write only




Figure 22-10. SPI Transmit Data Hold Register Definition

[Table 22-8](#) shows the field descriptions of the SPI transmit data hold register.

Table 22-8. SPI Transmit Data Hold Field Descriptions

| Bits | Name | Description |
|------|------|-------------------------------------|
| 0–31 | DATA | These bits are the data to be sent. |

22.4.1.6 SPI Receive Data Hold Register (SPIRD)

SPIRD, shown in [Figure 22-11](#), is used to receive a character of data from the SPI channel. Each time SPIE[NE] is set, the core can read SPIRD.



Figure 22-11. SPI Receive Data Hold Register Definition

[Table 22-9](#) shows the field descriptions of the SPI receive data hold register.

Table 22-9. SPI Receive Data Hold Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–31 | DATA | Received data. These bits are the received data from the SPI bus. |

22.4.1.6.1 Reverse Mode SPMODE[REV] Examples

In reverse data mode (SPMODE[REV] = 1) and regular data mode (SPMODE[REV] = 0) the data is placed in the SPIRD after reception is completed as described below for character length of 8 bits (SPMODE[LEN] = 7).




Figure 22-12. Example SPMODE[REV] = 0 SPMODE[LEN] = 7 LSB Sent First




Figure 22-13. Example SPMODE[REV] = 1 SPMODE[LEN] = 7 MSB Sent First




Figure 22-14. Example SPMODE[REV] = 1 SPMODE[LEN] = 15 MSB Sent First




Figure 22-15. Example SPMODE[REV] = 0 SPMODE[LEN] = 15 LSB Sent First

22.5 Initialization/Application Information

The following sections describe programming examples of the SPI master and slave.

22.5.1 SPI Master Programming Example

The following sequence initialize the SPI to run at a high speed in master mode:

1. Configure a parallel I/O signal to operate as the SPI select output signal if needed.
2. Write 0xFFFFFFFF to SPIE to clear any previous events. Configure SPIM to enable all desired SPI interrupts.
3. Configure SPMODE to enable normal operation (not loopback), master mode, SPI enabled, character length, and the fastest speed possible.
4. Write the first character to be sent to SPITD.

22.5.2 SPI Slave Programming Example

The following is an example initialization sequence to follow when the SPI is in slave mode. It is very similar to the SPI master example, except that $\overline{\text{SPISEL}}$ is used instead of a general-purpose I/O signal.

1. Write 0xFFFFFFFF to SPIE to clear any previous events.
2. Configure SPIM to enable all desired SPI interrupts.
3. Configure SPMODE to enable normal operation (not loopback), slave mode, SPI enabled, and characters length.

Chapter 23

JTAG/Testing Support

23.1 JTAG Overview

The device provides a JTAG (Joint Test Action Group) interface to facilitate boundary-scan testing. The JTAG interface complies to the IEEE 1149.1 boundary-scan specification. For additional information about JTAG operations, refer to the IEEE 1149.1 specification.

The JTAG interface consists of a set of five signals, three JTAG registers (see [Section 23.3, “JTAG Registers and Scan Chains,”](#)) and a test access port (TAP) controller, described in the following sections. A block diagram of the JTAG interface is shown in [Figure 23-1](#).




Figure 23-1. JTAG Interface Block Diagram

23.2 JTAG Signals

The device provides the following five dedicated JTAG signals:

- Test data input (TDI)
- Test data output (TDO)
- Test mode select (TMS)

- Test reset ($\overline{\text{TRST}}$)
- Test clock (TCK)

The TDI and TDO signals input and output all instructions and data to the JTAG scan registers. JTAG operations are controlled by the TAP controller through the TMS and TCK signals. Boundary-scan data is latched by the TAP controller on the rising edge of the TCK signal. The $\overline{\text{TRST}}$ signal is specified as optional by the IEEE 1149.1 specification, and is used to reset the TAP controller asynchronously. The assertion of the $\overline{\text{TRST}}$ signal at power-on reset ensures that the JTAG logic does not interfere with the normal operation of the device.

23.2.1 JTAG External Signal Descriptions

The JTAG signals are summarized in [Table 23-1](#).

Table 23-1. JTAG Test Signals Summary

| Name | Description | Functional Block | Function | Reset Value | I/O |
|--------------------------|------------------|------------------|---|----------------|-----|
| TCK | Test clock | Debug | Clock for JTAG testing. | — | I |
| TDI | Test data input | | Serial input for instructions and data to the JTAG test subsystem. Internally pulled up. | — | I |
| TDO | Test data output | | Serial data output for the JTAG test subsystem. High impedance except when scanning out data. | High impedance | O |
| TMS | Test mode select | | Carries commands to the TAP controller for boundary scan operations. Internally pulled up. | — | I |
| $\overline{\text{TRST}}$ | Test reset | | Resets the TAP controller asynchronously. Internally pulled up. | — | I |

[Table 23-2](#) shows detailed descriptions of the JTAG test signals.

Table 23-2. JTAG Test—Detailed Signal Descriptions

| Signal | I/O | Description | |
|--------|-----|-----------------------|--|
| TCK | I | JTAG test clock. | |
| | | State Meaning | Asserted/Negated—Should be driven by a free-running clock signal with a 30–70% duty cycle. Input signals to the TAP are clocked in on the rising edge. Changes to the TAP output signals occur on the falling edge. The test logic allows TCK to be stopped. |
| | | Timing | See IEEE 1149.1 specification for more details. |
| TDI | I | JTAG test data input. | |
| | | State Meaning | Asserted/Negated—The value present on the rising edge of TCK is clocked into the selected JTAG test instruction or data register. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor. |
| | | Timing | See IEEE 1149.1 specification for more details. |

Table 23-2. JTAG Test—Detailed Signal Descriptions (continued)

| Signal | I/O | Description | |
|--------------------------|-----|------------------------|--|
| TDO | O | JTAG test data output. | |
| | | State Meaning | Asserted/Negated—The contents of the selected internal instruction or data register are shifted out on this signal on the falling edge of TCK. Remains in a high-impedance state except when scanning data. |
| | | Timing | See IEEE 1149.1 specification for more details. |
| TMS | I | JTAG test mode select. | |
| | | State Meaning | Asserted/Negated—Decoded by the internal JTAG TAP controller to distinguish the primary operation of the test support circuitry. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor. |
| | | Timing | See IEEE 1149.1 specification for more details. |
| $\overline{\text{TRST}}$ | I | JTAG test reset. | |
| | | State Meaning | Asserted—Causes asynchronous initialization of the internal JTAG TAP controller. Must be asserted during power-on reset in order to properly initialize the JTAG TAP and for normal operation of the device. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor. Negated— Normal operation. |
| | | Timing | See IEEE 1149.1 specification for more details. |

23.3 JTAG Registers and Scan Chains

The bypass, boundary-scan, and instruction JTAG registers and their associated scan chains are mandatory for compliance with the IEEE 1149.1 specification.

- Bypass register. The bypass register is a single-stage register used to bypass the boundary-scan latches of the device during board-level boundary-scan operations involving components other than the device. The use of the bypass register reduces the total scan string size of the boundary-scan test.
- Boundary-scan registers. The JTAG interface provides a chain of registers dedicated to boundary-scan operations. To be JTAG-compliant, these registers cannot be shared with any functional registers of the device. The boundary-scan register chain includes registers controlling the direction of the input/output drivers, in addition to the registers reflecting the signal value received or driven.

The boundary-scan registers capture the input or output state of the device's signals during a Capture_DR TAP controller state. When a data scan is initiated following the Capture_DR state, the sampled values are shifted out through the TDO output while new boundary-scan register values are shifted in through the TDI input. At the end of the data scan operation, the boundary-scan registers are updated with the new values during an update_DR TAP controller state.

- Instruction register. The 8-bit JTAG instruction register serves as an instruction and status register. As TAP controller instructions are scanned in through the TDI input, the TAP controller status bits are scanned out through the TDO output.

- TAP controller. The device provides a standard JTAG TAP controller that controls instruction and data scan operations. The TMS signal controls the state transitions of the TAP controller.

Chapter 24

General Purpose I/O (GPIO)

24.1 Introduction

This chapter describes the general-purpose I/O module, including pin descriptions, register settings, and interrupt capabilities. [Figure 24-1](#) shows the block diagram of the GPIO module.




Figure 24-1. GPIO Module Block Diagram

24.1.1 Overview

The GPIO module supports 32 general-purpose I/O ports. Each port can be configured as an input or as an output. If a port is configured as an input, it can optionally generate an interrupt on detection of a change. If a port is configured as an output, it can be individually configured as an open-drain or a fully active output.

24.1.2 Features

The GPIO unit implements the following features:

- 32 input/output ports
- Some ports have dedicated processor signals. Others are multiplexed together with other functional signals. See [Chapter 2, “Signal Descriptions.”](#)
- All signals are configured as inputs when the device comes out of reset and also when $\overline{\text{HRESET}}$ is asserted.

- Open-drain capability on all ports
- All ports can optionally generate an interrupt upon changing their state.

24.2 External Signal Description

The following section provides information about GPIO signals.

24.2.1 Signals Overview

Table 24-1 provides detailed descriptions of the external GPIO signals.

Table 24-1. GPIO External Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|------------|-----|---|
| GPIO[0:31] | I/O | General purpose I/O. Each signal can be set individually to act as input or output, according to application needs. |
| | | State Meaning Asserted/Negated—Defined per application. |
| | | Timing Assertion/Negation—Inputs can be asserted completely asynchronously. Outputs are asynchronous to any externally visible clock |

24.3 Memory Map/Register Definition

The GPIO has programmable registers that occupy 24 bytes of memory-mapped space. Note that reading undefined portions of the memory map returns all zeros and writing has no effect.

All GPIO registers are 32 bits wide and are located on 32-bit address boundaries. All addresses used in this chapter are offsets from the address held in IMMRBAR as defined in Chapter 2, “Memory Map.”

Table 24-2 shows the memory map of GPIO.

Table 24-2. GPIO Register Address Map

| Offset | Register | Access | Reset Value | Section/Page |
|---|--|--------|-------------|-----------------------------|
| General Purpose I/O (GPIO)—Block Base Address 0x0_0C00 | | | | |
| 0xC00 | GPIO direction register (GPDIR) | R/W | 0x0000_0000 | 24.3.1/24-3 |
| 0xC04 | GPIO open drain register (GPODR) | R/W | 0x0000_0000 | 24.3.2/24-3 |
| 0xC08 | GPIO data register (GPDAT) | R/W | 0x0000_0000 | 24.3.3/24-4 |
| 0xC0C | GPIO interrupt event register (GPIER) | w1c | Undefined | 24.3.4/24-4 |
| 0xC10 | GPIO interrupt mask register (GPIMR) | R/W | 0x0000_0000 | 24.3.5/24-4 |
| 0xC14 | GPIO external interrupt control register (GPICR) | R/W | 0x0000_0000 | 24.3.6/24-5 |

24.3.1 GPIO Direction Register (GPDIR)

The GPIO direction registers (GPDIR), shown in [Figure 24-2](#), defines the direction of the individual ports.




Figure 24-2. GPIO Direction Register (GPDIR)

[Table 24-3](#) defines the bit fields of GPDIR.

Table 24-3. GPDIR Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–31 | Dn | Direction. Indicates whether a signal is used as an input or an output. 0 The corresponding signal is an input. 1 The corresponding signal is an output. |

24.3.2 GPIO Open Drain Register (GPODR)

The GPIO open drain register (GPODR), shown in [Figure 24-3](#), defines the way individual ports drive their output.



Figure 24-3. GPIO Open Drain Register (GPODR)

[Table 24-4](#) defines the bit fields of GPODR.

Table 24-4. GPODR Bit Settings

| Bits | Name | Description |
|------|------|---|
| 0–31 | Dn | Open-drain configuration. Indicates whether a signal is actively driven as an output or an open-drain driver. This register has no effect on signals programmed as inputs in the corresponding GPDIR. 0 The I/O signal is actively driven as an output. Note: The I/O signal is an open-drain driver. As an output, the signal is driven active-low, otherwise it is three-stated. |

24.3.3 GPIO Data Register (GPDAT)

The GPIO data register (GPDAT), shown in [Figure 24-4](#), carries the data in/out for the individual ports.




Figure 24-4. GPIO Data Register (GPDAT)

[Table 24-5](#) defines the bit fields of GPDAT.

Table 24-5. GPnDAT Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–31 | Dn | Data. Write data is latched and presented on external signals if GPDIR has configured the port as an output. Read operation always returns the data at the signal. |

24.3.4 GPIO Interrupt Event Register (GPIER)

The GPIO interrupt event register (GPIER), shown in [Figure 24-5](#), carries information of the events that caused an interrupt. Each bit in GPIER, corresponds to an interrupt source. GPIER bits are cleared by writing ones. However, writing zero has no effect.




Figure 24-5. GPIO Interrupt Event Register (GPIER)

[Table 24-6](#) defines the bit fields of GPIER.

Table 24-6. GPIER Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–31 | Dn | Interrupt events. Indicates whether an interrupt event occurred on the corresponding GPIO signal. 0 No interrupt event occurred on the corresponding GPIO signal. 1 Interrupt event occurred on the corresponding GPIO signal. |

24.3.5 GPIO Interrupt Mask Register (GPIMR)

The GPIO interrupt mask register (GPIMR), shown in [Figure 24-6](#), defines the interrupt masking for the individual ports. When a masked interrupt request occurs, the corresponding GPIER bit is set, regardless

of the GPIMR state. When one or more non-masked interrupt events occur, the GPIO module issues an interrupt to the on chip interrupt controller.




Figure 24-6. GPIO Interrupt Mask Register (GPIMR)

Table 24-7 defines the bit fields of GPIMR.

Table 24-7. GPIMR Bit Settings

| Bits | Name | Description |
|------|-----------|--|
| 0–31 | <i>Dn</i> | Interrupt mask. Indicates whether an interrupt event is masked or not masked. 0 The input interrupt signal is masked (disabled). 1 The input interrupt signal is not masked (enabled). |

24.3.6 GPIO Interrupt Control Register (GPICR)

The GPIO interrupt control register (GPICR), shown in Figure 24-7, determines whether the corresponding port line asserts an interrupt request on either a high-to-low change or any change on the state of the signal.



Figure 24-7. GPIO Interrupt Control Register (GPICR)

Table 24-8 defines the bit fields of GPICR.

Table 24-8. GPICR Bit Settings

| Bits | Name | Description |
|------|-----------|--|
| 0–31 | <i>Dn</i> | Edge detection mode. The corresponding port line asserts an interrupt request according to the following: 0 Any change on the state of the port generates an interrupt request. 1 High-to-low change on the port generates an interrupt request. |



Chapter 25

Time Division Multiplexing (TDM) Interface

25.1 Introduction

This chapter discusses time division multiplexing (TDM) interface, discusses the architecture, the programming model, the operating modes, and the initialization of the TDM. The TDM is a full-duplex, serial port designed to allow digital signal processors (DSPs) to communicate with a variety of serial devices, including industry-standard framers, codecs, other DSPs, and microprocessors. It is typically used to transfer samples in a periodic manner. The TDM consists of independent transmitter and receiver sections with independent clock generation and frame synchronization.

The TDM interface consists of a TDM module supporting 128 channels running at up to 50 Mbps with 8- and 16-bit word size. The TDM bus connects gluelessly to most T1/E1 frames as well as to common buses such as the H.110, SCAS, and MVIP. The TDM also supports an I²S mode. The TDM module operates in independent or shared mode when receiving or transmitting data:

- In independent mode, there are different sync, clock, and data for receive and transmit.
- In shared sync and clock mode, the clock and the sync are shared between the transmit and receive with different receive and transmit data.

25.1.1 Features

TDM capabilities include the following:

- Independent (asynchronous) or shared (synchronous) transmit and receive sections with separate or shared internal/external clocks and frame syncs
- TDM (network) mode operation allowing multiple devices to share the port with as many as 128 time slots
- Single-channel mode operation using frame sync
- Time slot enable registers (receive and transmit)
- End-of-frame interrupt
- Programmable internal clock divider
- Programmable word length (8 or 16 bits)
- Program options for frame sync and clock generation
- Programmable to MSB or LSB first
- TDM power-down feature
- A-law/ μ -law hardware conversion is supported for 8-bit channels
- Loopback mode

25.2 Signal Descriptions

The following sections provide detailed descriptions of the TDM signals.

25.2.1 Overview

Table 25-1 gives a summary of the different TDM signals.

Table 25-1. TDM Signal Properties

| Name | Function | Type | Reset State |
|---------|-------------------------|------|-------------|
| TDM_TCK | TDM transmit clock | I/O | Input |
| TDM_TFS | TDM transmit frame sync | I/O | Input |
| TDM_RCK | TDM receive clock | I/O | Input |
| TDM_RFS | TDM receive frame sync | I/O | Input |
| TDM_TD | TDM transmit data | O | High Z |
| TDM_RD | TDM receive data | I | Input |

25.2.2 External Signals Descriptions

Four to six signals are required for TDM full-duplex operation, depending on the selected operation mode. The function of the I/O pins for the TDM is determined by a number of control bits. The serial transmit data (TD) signal and serial control (RCK and RFS) signals are fully synchronized to the clock if they are programmed as transmit-data signals.

In addition to the functional configuration of the I/O pins, the direction of clocks and frame syncs can be configured. Each can be configured as an input or output. The frame syncs and the clocks can also be shared with other TDM modules on-chip.

Table 25-2. TDM External Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|---------|-----|---|
| TDM_TCK | I/O | TDM transmit clock. This pin can be configured as either an input or an output pin. This clock signal is used by the transmitter and the receiver in shared modes or by the transmitter in independent modes. Note: Although an external serial clock can be independent of and asynchronous to the CSB system clock, the external TDM clock frequency must not exceed 50 MHz. |
| TDM_TFS | I/O | TDM transmit frame sync. This pin is used for frame sync I/O and can be configured as either an input or an output pin. The direction of this signal is determined by the configuration register. The frame sync is used by the transmitter to synchronize the transfer of data in independent mode. It is used by both the receiver and transmitter in shared mode. The frame sync signal can be one bit, or one word in length. The start of the frame sync can occur 0-3 bits before the transfer of data. |
| TDM_RCK | I/O | TDM receive clock. The function of this signal is determined by the programmer selecting either shared or independent mode. In independent mode, this signal is used for the receive clock I/O. In shared mode, this signal can be used as a GPIO. |

Table 25-2. TDM External Signals—Detailed Signal Descriptions (continued)

| Signal | I/O | Description |
|---------|-----|---|
| TDM_RFS | I/O | TDM receive frame sync. The function of this signal is determined by the programmer selecting either shared or independent modes. In the independent mode, such as a single codec with independent transmit and receive, the RFS is the receiver frame sync I/O. In shared mode, this signal can be used as a GPIO. |
| TDM_TD | O | TDM transmit data. The TD signal is used for transmitting data from the serial transmit shift register. TD is an output pin when data is being transmitted from the transmit shift register. The TD pin can be tri-stated after transmitting the last data bit when another data word does not follow immediately. |
| TDM_RD | I | TDM receive data. The RD signal receives serial data, then transfers the data to the TDM receive data shift register. |

25.3 TDM Overview

Figure 25-1 illustrates the block diagram of the TDM. It consists of control registers to set-up the port, status registers, separate transmit and receive circuits with FIFO registers, and separate serial clock and frame sync generation for the transmit and receive sections.




Figure 25-1. TDM Block Diagram

25.3.1 TDM Basics

Multiple TDM channels are transferred sequentially in a structure called a frame. The frame start is identified by a frame sync signal that is briefly asserted at the beginning of every frame. The TDM module can receive or transmit up to 128 channels total channels at a granularity of two channels. There is also a single channel mode, which only transmits and receives one channel of data. The number of receive channels is determined by the RNCF field in the TDM receive frame parameters register (TDMRFP). The number of transmit channels is determined by the TNCF field in the TDM transmit frame parameters register (TDMTFP).

The size of all the channels is unified and it can be 8 or 16 bits. The receive channel size is determined by the RCS field in the TDMRFP; the transmit channel size is determined by the TCS field in the TDMTFP.

When the TDM connects to a T1 framer, the RT1 field in the TDM receive frame parameters register (TDMRFP) and the TT1 field in the TDM transmit frame parameters register (TDMTFP) should be set. Also the number of channels in the RCS and TCS fields should be set to 24. The T1 frame contains 193 bits (24 channels of 8 bits each) where the first bit of the frame is a frame alignment bit that is not used by the TDM. At the T1 received frame, the frame alignment bit is removed and does not transfer to the data registers. At the transmit T1 frame, the bit is not driven out.

Figure 25-2 shows an example of a TDM interface. The receive frame contains two 8-bit channels. The transmit frame contains four 8-bit channels. Figure 25-3 shows an example of T1 frame. Note the first bit of the frame is not used by the receive and transmit TDM.

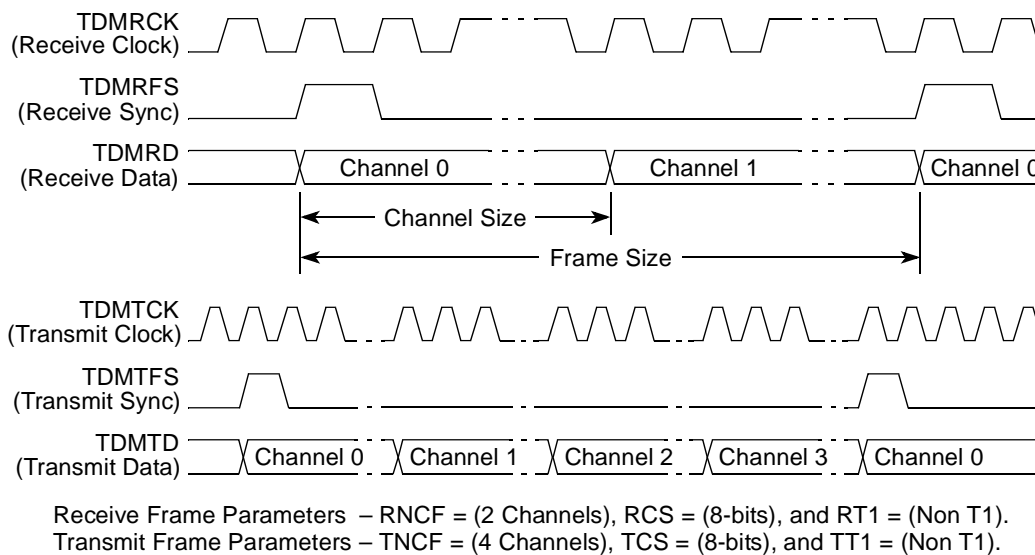
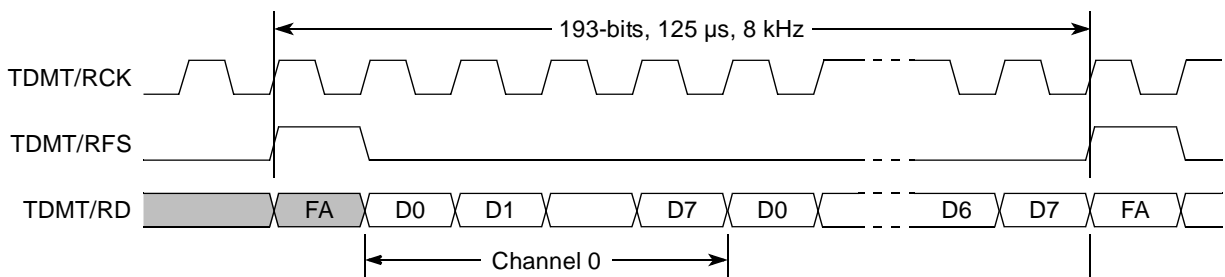


Figure 25-2. TDM Frames



Receive Frame Parameters – RNC[7:0] = (24 Channels), RCS = (8-bits), and RT1 = (T1 Mode).
 Transmit Frame Parameters – TNC[7:0] = (24 Channels), TCS = (8-bits), and TT1 = (T1 Mode).

Figure 25-3. T1 Frame

25.3.2 Common Signals for the TDM Modules

The sync and clock signals can be shared among the TDM modules or separate for each TDM module. The TDM modules share sync and clock signals; the common signals connect to the following signal lines:

- The transmit sync/frame sync connects to TFS (receive and transmit share the same sync signal for all TDMs).
- The transmit clock/frame clock connects to TCK (receive and transmit share the same clock signal for all TDMs).

The configuration registers should be identical for the TDM modules that share signals. [Figure 25-4](#) illustrates common receive sync, receive clock, transmit sync, and transmit clock signals.

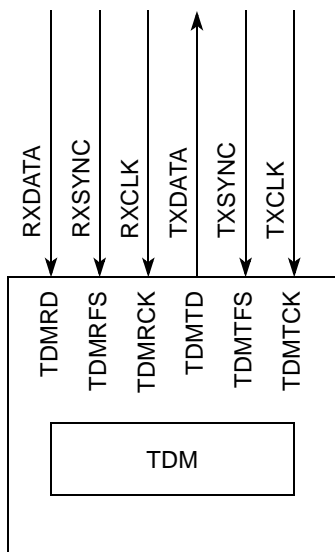


Figure 25-4. TDM Modules Shared Mode

25.4 TDM Programming Model

The handshake between the TDM module and the core occurs via a set of registers, interrupts, and DMA requests. The TDM registers are mapped into the SB and AHB address space. [Table 25-3](#) shows the location of the SB, AHB, and DMAC memory maps in this chapter.

Table 25-3. TDM Module Memory Map Summary

| TDM Block | Section/Page |
|--|-------------------------------|
| TDM configuration registers (SB interface) | 25.4.2/25-7 |
| TDM data registers (AHB interface) | 25.4.4/25-28 |
| TDM clock control registers | 25.9.3/25-52 |
| TDM DMA controller (DMAC) registers | 25.10.2/25-55 |

25.4.1 SB Interface

This section details the registers that are located in the SB memory map. Refer to [Section 25.4.3, “AHB Interface,”](#) for details on the registers that are located in the AHB memory map. The SB address of a TDM register is composed of two concatenated parts:

- SB base address:
 - TDM module: TDM_BASE, refer to the memory map in [Section 25.4.2, “SB Memory Map.”](#)
 - 8-bit register offset. The TDM module registers are divided into configuration registers, control registers, status, and data registers. The registers detailed in this section are as follows:
- Configuration registers to set the operation modes and provide indications for all channels. They are set before the TDM is enabled and should not be changed while the TDM is active.
 - Control registers to set the channel specific parameters individually for each channel and to set the threshold pointers. These registers can also be changed during operation. The channel enable and mask registers are sampled at the beginning of the frame.
 - Status registers, which are read-only registers and can be accessed any time.

25.4.2 SB Memory Map

The TDM registers are listed in [Table 25-4](#).

Table 25-4. TDM SB Memory-Mapped Registers

| Offset | Name | Access | Reset | Section/Page |
|--|--|--------|-------------|----------------------------------|
| TDM Configuration—Block Base Address 0x1_6000 | | | | |
| TDM Configuration | | | | |
| 0x000 | TDM general interface register (TDMGIR) | R/W | 0x0000_0000 | 25.4.2.1.1/25-8 |
| 0x004 | TDM receive interface register (TDMRIR) | R/W | 0x0000_0000 | 25.4.2.1.2/25-9 |
| 0x008 | TDM transmit interface register (TDMTIR) | R/W | 0x0001_0000 | 25.4.2.1.3/25-12 |

Table 25-4. TDM SB Memory-Mapped Registers (continued)

| Offset | Name | Access | Reset | Section/Page |
|--|--|--------|-------------|----------------------------------|
| 0x00C | TDM receive frame parameters (TDMRFP) | R/W | 0x0000_0000 | 25.4.2.1.4/25-14 |
| 0x010 | TDM transmit frame parameters (TDMTFP) | R/W | 0x0000_0000 | 25.4.2.1.5/25-16 |
| TDM Channel, Interrupt, and Rx/Tx Enables | | | | |
| 0x020–0x02C | TDM receive channel enable registers (TDMRCEN0–3) | R/W | 0x0000_0000 | 25.4.2.2.1/25-17 |
| 0x040–0x04C | TDM transmit channel enable registers (TDMTCEN0–3) | R/W | 0x0000_0000 | 25.4.2.2.2/25-17 |
| 0x060–0x06C | TDM transmit channel mask registers (TDMTCMA0–3) | R/W | 0x0000_0000 | 25.4.2.2.3/25-18 |
| 0x080 | TDM receive control register (TDMRCR) | R/W | 0x0000_0000 | 25.4.2.2.4/25-19 |
| 0x084 | TDM transmit control register (TDMTCR) | R/W | 0x0000_0000 | 25.4.2.2.5/25-19 |
| 0x088 | TDM receive interrupt enable register (TDMRIER) | R/W | 0x0000_0000 | 25.4.2.2.6/25-20 |
| 0x08C | TDM transmit interrupt enable register (TDMTIER) | R/W | 0x0000_0000 | 25.4.2.2.7/25-21 |
| TDM Status | | | | |
| 0x0A0 | TDM receive event register (TDMRER) | R/W | 0x0000_0000 | 25.4.2.3.1/25-22 |
| 0x0A4 | TDM transmit event register (TDMTER) | R/W | 0x0000_0000 | 25.4.2.3.2/25-24 |
| 0x0A8 | TDM receive status register (TDMRSR) | RO | 0x0000_0000 | 25.4.2.3.3/25-26 |
| 0x0AC | TDM transmit status register (TDMTSR) | RO | 0x0000_0000 | 25.4.2.3.4/25-26 |

25.4.2.1 Configuration Registers

This section describes the TDM module registers.

25.4.2.1.1 TDM General Interface Register (TDMGIR)

The TDM general interface register, shown in [Figure 25-5](#), defines the TDM interface operation mode.




Figure 25-5. TDM General Interface Register

Table 25-7 defines the bit fields of TDMRIR.

Table 25-7. TDMRIR Bit Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–13 | — | Reserved. Write to zero for future compatibility. |
| 14–15 | RFWM | Receive FIFO watermark. Determines when the receive full FIFO event will occur. 00 The FIFO will be full with 1 or more elements. 01 The FIFO will be full with 2 or more elements. 10 The FIFO will be full with 3 or more elements. 11 The FIFO will be full with 4 or more elements. |
| 16 | RFEN | Receive FIFO enable. Determines whether the receive FIFO is used. 0 Rx FIFO is not used. 1 Rx FIFO is used. |
| 17 | RWEN | Receive wide FIFO enable. Determines whether the receive FIFO is used as 1 sample (8/16 bit data) or as 64 bits packed with 8 8-bit data or four 16 bit data per entry in the FIFO. Only set when RFEN is set. 0 Rx FIFO is used as 16 bit or 8 bit wide. 1 Rx FIFO is used as 64 bit wide, packed with 8-8bit data or 4-16 bit data. |
| 18 | RSO | Receive sync output. Determines whether the receive sync is driven out by the TDM receiver or is input to the TDM receiver. For details, see Section 25.7.1.1, “Sync Out Configuration.” 0 Receive sync is input. 1 Receive sync is output. |
| 19 | — | Reserved. Write to zero for future compatibility. |
| 20 | RSL | Receive sync out length. Indicates whether the TDMRFS is asserted for one cycle of TDMRCK or is asserted for the duration of the first channel in the frame. Note: must be set to 0 when there is only 1 receive channel in a frame. For details, see Section 25.7.1.1, “Sync Out Configuration.” 0 The receive sync_out width is one bit. 1 The receive sync_out width is equal to the channel width. |
| 21 | — | Reserved. Write to zero for future compatibility. |
| 22 | RCOE | Receive clock output enable. Determines whether the Receive Clock out signal TDMRCK is driven out from the appropriate timer. For details, see Section 25.5.1, “TDM Clock and Frame Sync Generation.” 0 The receive clock is an input. 1 The receive clock is an output. |
| 23–24 | — | Reserved. Write to zero for future compatibility. |
| 25 | RDMA | Receive DMA enable when set, the TDM will request a DMA module transfer of the received data when: the Rx data ready (TDMRER[RDR]) bit is set (Rx FIFO disabled) or the receive FIFO full (TDMRER[RFF]) bit is set (Rx FIFO enabled). 0 The TDM will not request DMA service for the receiver. 1 The TDM will request DMA service for the receiver. |
| 26–27 | RFSD | Receive frame sync delay with the RDE and the RFSE bits, determines the number of clocks between the receive sync signal and the first data bit of the receive frame. For examples, see Section 25.7.1.2, “Sync In Configuration.” Refer to Table 25-8 . |
| 28 | RSA | Receive sync active. Determines the polarity of the receive sync signal. For details, see Section 25.7.1.2, “Sync In Configuration.” 0 The receive sync is active on logic 1. 1 The receive sync is active on logic 0. |

Table 25-7. TDMRIR Bit Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 29 | RDE | Receive data edge. Determines whether the receive data is driven out on the rising or falling edge of the receive clock. For details, see Section 25.7.1.2, “Sync In Configuration.” 0 The receive data is driven out on the rising edge of the receive clock. 1 The receive data is driven out on the falling edge of the receive clock. |
| 30 | RFSE | Receive frame sync edge. Determines whether the receive frame sync signal is driven with the rising or falling edge of the receive clock. For details, see Section 25.7.1.2, “Sync In Configuration.” 0 The receive frame sync signal is driven on the rising edge of the receive clock. 1 The receive frame sync signal is driven on the falling edge of the receive clock. |
| 31 | RFDO | Receive reversed data order. For examples, see Section 25.7.1.4, “Reverse Data Order.” 0 The first bit of a received channel is stored as the least significant bit in the on-chip memory. 1 The first bit of a received channel is stored as the most significant bit in the on-chip memory. |

Table 25-8. Received Data Delay for Receive Frame Sync

| Frame Sync Delay | Frame Sync Edge | Data Edge | Receive Clocks ¹ |
|------------------|-----------------|-----------|-----------------------------|
| 00 | 0 | 0 | 0 |
| 00 | 0 | 1 | 0.5 |
| 00 | 1 | 0 | 0.5 |
| 00 | 1 | 1 | 0 |
| 01 | 0 | 0 | 1 |
| 01 | 0 | 1 | 1.5 |
| 01 | 1 | 0 | 1.5 |
| 01 | 1 | 1 | 1 |
| 10 | 0 | 0 | 2 |
| 10 | 0 | 1 | 2.5 |
| 10 | 1 | 0 | 2.5 |
| 10 | 1 | 1 | 2 |
| 11 | 0 | 0 | 3 |
| 11 | 0 | 1 | 3.5 |
| 11 | 1 | 0 | 3.5 |
| 11 | 1 | 1 | 3 |

¹ Receive clocks is the number of receive clocks between the first edge of the receive frame sync and the drive of the first data bit of the received frame.

Table 25-9. TDMTIR Bit Descriptions (continued)

| Bits | Name | Description |
|-------|------|---|
| 21 | — | Reserved. Write to zero for future compatibility. |
| 22 | TCOE | Transmit clock output enable. Determines whether the transmit clock out signal TDMTCK is driven out from the appropriate timer. For details, see Section 25.5.1, “TDM Clock and Frame Sync Generation.” 0 The transmit clock is an input. 1 The transmit clock is an output. |
| 23–24 | — | Reserved. Write to zero for future compatibility. |
| 25 | TDMA | Transmit DMA Enable. When set, the TDM will request a DMA module transfer of the transmit data when: the Tx data register empty (TDMTER[TDR]) bit is set (Tx FIFO disabled) or the transmit FIFO empty (TDMTER[TFE]) bit is set (Tx FIFO enabled). 0 The TDM will not request DMA service for the transmitter. 1 The TDM will request DMA service for the transmitter. |
| 26–27 | TFSD | Transmit frame sync delay. With the TDE and the TFSE bits, determines the number of clocks between the transmit sync signal and the first data bit of the transmit frame. For examples, see Section 25.7.1.2, “Sync In Configuration.” Refer to Table 25-10 . |
| 28 | TSA | Transmit sync active. Determines the polarity of the transmit sync signal. For details, see Section 25.7.1.2, “Sync In Configuration.” 0 The transmit sync is active on logic 1. 1 The transmit sync is active on logic 0. |
| 29 | TDE | Transmit data edge. Determines whether the transmit data is driven out on the rising or falling edge of the transmit clock. For details, see Section 25.7.1.2, “Sync In Configuration.” 0 The transmit data is driven out on the rising edge of the transmit clock. 1 The transmit data is driven out on the falling edge of the transmit clock. |
| 30 | TFSE | Transmit frame sync edge. Determines whether the transmit frame sync signal is driven with the rising or falling edge of the transmit clock. For details, see Section 25.7.1.2, “Sync In Configuration.” 0 The transmit frame sync signal is driven on the rising edge of the transmit clock. 1 The transmit frame sync signal is driven on the falling edge of the transmit clock. |
| 31 | TRDO | Transmit reversed data order. For examples, see Section 25.7.1.4, “Reverse Data Order.” 0 The least significant bit of the memory is sent out as the first transmit data bit. 1 The most significant bit of the memory is sent out as the first transmit data bit. |

Table 25-10. Transmit Data Delay for Transmit Frame Sync

| Frame Sync Delay | Frame Sync Edge | Data Edge | Transmit Clocks ¹ |
|------------------|-----------------|-----------|------------------------------|
| 00 | 0 | 0 | 0 |
| 00 | 0 | 1 | 0.5 |
| 00 | 1 | 0 | 0.5 |
| 00 | 1 | 1 | 0 |
| 01 | 0 | 0 | 1 |
| 01 | 0 | 1 | 1.5 |
| 01 | 1 | 0 | 1.5 |
| 01 | 1 | 1 | 1 |
| 10 | 0 | 0 | 2 |

Table 25-10. Transmit Data Delay for Transmit Frame Sync (continued)

| Frame Sync Delay | Frame Sync Edge | Data Edge | Transmit Clocks ¹ |
|------------------|-----------------|-----------|------------------------------|
| 10 | 0 | 1 | 2.5 |
| 10 | 1 | 0 | 2.5 |
| 10 | 1 | 1 | 2 |
| 11 | 0 | 0 | 3 |
| 11 | 0 | 1 | 3.5 |
| 11 | 1 | 0 | 3.5 |
| 11 | 1 | 1 | 3 |

¹ Transmit clocks is the number of transmit clocks between the first edge of the transmit frame sync and the first data bit of the frame that is driven out.

25.4.2.1.4 TDM Receive Frame Parameters (TDMRFP)

The TDM receive frame parameters register, shown in [Figure 25-8](#), defines the TDM receive frame parameters.




Figure 25-8. TDM Transmit Interface Register

Table 25-11 defines the bit fields of TDMRFP.

Table 25-11. TDMTIR Bit Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–7 | — | Reserved. Write to zero for future compatibility. |
| 8–15 | RNCF | <p>Receive number of channels in a TDM frame. Specifies the total number of channels that are received in the TDM modules. One TDM frame can contain 1–128 channels at a granularity of two.</p> <p>Note: If RNCF field is clear then the minimum number of channels is one.</p> <p>8x00 1 received channels 8X01 2 received channels 8x02 Reserved 8x03 4 received channels 8x04 Reserved ... 8x7D 126 received channels 8x7F 128 received channels 8x80-8xFF Reserved</p> <p>Note: The even values are reserved, except x00.</p> |
| 16–25 | — | Reserved. Write to zero for future compatibility. |
| 26–27 | RCS | <p>Receive channel size. Determines the receiver channel size for all channels in the frame, including decoding μ-Law and A-Law.</p> <p>00 The receive channel size is 8 bits. 01 The receive channel size is 8 bits using μ-Law decoding. 10 The receive channel size is 8 bits using A-Law decoding. 11 The receive channel size is 16 bits.</p> |
| 28–29 | — | Reserved. Write to zero for future compatibility. |
| 30 | RT1 | <p>Receive T1 frame. Determines whether the receive frame is T1 frame or non T1. Note: In T1 mode the channel size must be 8 (RCS = 0) and the number of channels must be 24 (RNCF = 0x018). For details, see Section 25.3, “TDM Overview”.</p> <p>0 The receive frame is a non T1 frame. 1 The receive frame is a T1 frame.</p> |
| 31 | — | Reserved. Write to zero for future compatibility. |

25.4.2.1.5 TDM Transmit Frame Parameters (TDMTFP)

The TDM transmit frame parameters register, shown in [Figure 25-9](#), defines the TDM transmit frame parameters.




Figure 25-9. TDM Transmit Frame Parameters

[Table 25-12](#) defines the bit fields of TDMTFP.

Table 25-12. TDMTFP Bit Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–7 | — | Reserved. Write to zero for future compatibility. |
| 8–15 | TNCF | Transmit number of channels in a TDM frame. Specifies the total number of channels that are transmitted in the TDM module. One TDM frame can contain 1–128 channels at a granularity of two. Note: If TNCF field is clear then the minimum number of channels is one. 8x00 1 transmit channels 8X01 2 transmit channels 8x02 Reserved 8x03 4 transmit channels 8x04 Reserved ... 8x7D 126 transmit channels 8x7F 128 transmit channels 8x80–8xFF Reserved Note: The even values are reserved, except x00. |
| 16–25 | — | Reserved. Write to zero for future compatibility. |
| 26–27 | TCS | Transmit channel size. Determines the transmitter channel size for all channels in the Frame. 00 The transmit channel size is 8 bits. 01 The transmit channel size is 8 bits using μ -Law decoding. 10 The transmit channel size is 8 bits using A-Law decoding. 11 The transmit channel size is 16 bits. |
| 28–29 | — | Reserved. Write to zero for future compatibility. |
| 30 | TT1 | Transmit T1 frame. Determines whether the transmit frame is T1 frame or non T1. Note: : In T1 mode the channel size must be 8 (TCS = 0) and the number of channels must be 24 (TNCF = 0x018). For details, see Section 25.3, “TDM Overview.” 0 The transmit frame is a non T1 frame. 1 The transmit frame is a T1 frame. |
| 31 | — | Reserved. Write to zero for future compatibility. |

25.4.2.2 Control Registers

This section discusses the following registers:

- Section 25.4.2.2.1, “TDM Receive Channel Enable n (TDMRCEN n)”
- Section 25.4.2.2.2, “TDM Transmit Channel Enable n (TDMTCEN n)”
- Section 25.4.2.2.3, “TDM Transmit Channel Mask n (TDMTCMA n)”
- Section 25.4.2.2.4, “TDM Receive Control Register (TDMRCR)”
- Section 25.4.2.2.5, “TDM Transmit Control Register (TDMTCR)”
- Section 25.4.2.2.6, “TDM Receive Interrupt Enable Register (TDMRIER)”
- Section 25.4.2.2.7, “TDM Transmit Interrupt Enable Register (TDMTIER)”

25.4.2.2.1 TDM Receive Channel Enable n (TDMRCEN n)

Figure 25-10 shows the TDM receive channel enable registers, 0–3.




Figure 25-10. TDM Receive Channel Enable n

Table 25-13 defines the bit fields of TDMRCEN n .

Table 25-13. TDMRCEN n Bit Descriptions

| Bits | Name | Description |
|------|----------|---|
| 0–31 | RCEN n | Receive channel active enable group n [0–31]. Set when the receive channel m is active. Each bit corresponds to each channel, as follows: <ul style="list-style-type: none"> • $n = 0$, bit 0 is the active bit for channel 0 and bit 1 is the active bit for channel 1. • $n = 1$, bit 0 is the active bit for channel 32 and bit 1 is the active bit for channel 33. • $n = 2$, bit 0 is the active bit for channel 64 and bit 1 is the active bit for channel 65. • $n = 3$, bit 0 is the active bit for channel 96 and bit 1 is the active bit for channel 97. 0 The channel m is not active. 1 The channel m is active. |

25.4.2.2.2 TDM Transmit Channel Enable n (TDMTCEN n)

To drive and account for time slots, there is a double layer mask for transmission of data. Two set of registers are defined for each channel—TDMTCEN n and TDMTCMA n (Section 25.4.2.2.3, “TDM Transmit Channel Mask n (TDMTCMA n)”).

If all bits of the TDMTCMA n register are cleared, the TDM transmitter will continue to operate based solely on the TDMTCEN n registers. The TDMTCMA n registers are used to output data onto specific

channels. Therefore, $TDMTCMA_n$ register bits should only be set on the channels that have their corresponding $TDMTCEN_n$ bit already cleared; otherwise the channel is disabled.

Once the $TDMTCEN_n$ register is cleared, the $TDMTCMA_n$ register bit is set to 1 in the desired channel bit location for data to be discarded by the TDM. This allows for memory structures to keep channels open that are physically closed on the TDM, which are transferred to the TDM via a DMA, where the TDM discards the data.

If the $TDMTCMA_n$ bit is set on a channel that is enabled, then the data for that channel is discarded by the TDM and the TDM will instead transmit all 1's. Disabling a channel in this manner causes the time slot for that channel to be ignored by the TDM. This means no data is transferred to the transmit shift register.

Figure 25-11 shows the TDM transmit channel enable registers, 0–3.




Figure 25-11. TDM Transmit Channel Enable n

Table 25-14 defines the bit fields of $TDMTCEN_n$.

Table 25-14. $TDMTCEN_n$ Bit Descriptions

| Bits | Name | Description |
|------|----------|---|
| 0–31 | $TCEN_n$ | Transmit channel active enable group n [0–31]. Set when the transmit channel m is active. Each bit corresponds to each channel, as follows: <ul style="list-style-type: none"> • $n = 0$, bit 0 is the active bit for channel 0 and bit 1 is the active bit for channel 1. • $n = 1$, bit 0 is the active bit for channel 32 and bit 1 is the active bit for channel 33. • $n = 2$, bit 0 is the active bit for channel 64 and bit 1 is the active bit for channel 65. • $n = 3$, bit 0 is the active bit for channel 96 and bit 1 is the active bit for channel 97. 0 The channel m is not active. 1 The channel m is active. |

25.4.2.2.3 TDM Transmit Channel Mask n ($TDMTCMA_n$)

The transmit channel mask registers 0–3 are shown in Figure 25-12. See Section 25.4.2.2.2, “TDM Transmit Channel Enable n ($TDMTCEN_n$)” for a description on how these registers are used to output data onto specific channels.



Figure 25-12. TDM Transmit Channel Mask n

Table 25-15 defines the bit fields of TDMTTCMA n .

Table 25-15. TDMTTCMA n Bit Descriptions

| Bits | Name | Description |
|------|----------|---|
| 0–31 | TCMA n | Transmit channel mask group 0 [0–31]. Set when the transmit channel m data is to be ignored when received in the transmit data register (TDREG). Each bit corresponds to each channel, as follows: <ul style="list-style-type: none"> • $n = 0$, bit 0 is the active bit for channel 0 and bit 1 is the active bit for channel 1. • $n = 1$, bit 0 is the active bit for channel 32 and bit 1 is the active bit for channel 33. • $n = 2$, bit 0 is the active bit for channel 64 and bit 1 is the active bit for channel 65. • $n = 3$, bit 0 is the active bit for channel 96 and bit 1 is the active bit for channel 97. 0 The channel m is transmitted according to the corresponding TDMTTCEN0–3 register. 1 The channel m data is discarded. |

25.4.2.2.4 TDM Receive Control Register (TDMRCR)

The TDM receive control register, shown in Figure 25-13, controls the activation/deactivation of the TDM receiver. The activation of the receiver is valid only when the RENS bit is clear.




Figure 25-13. TDM Receive Control Register

Table 25-16 defines the bit fields of TDMRCR.

Table 25-16. TDMRCR Bit Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–30 | — | Reserved. Write to zero for future compatibility. |
| 31 | REN | Receive enable. Determines whether the receive TDM is enabled or disabled. Note: Setting this bit is the last step in initializing the receiver. 0 Receiver is disabled 1 Receiver is enabled |

25.4.2.2.5 TDM Transmit Control Register (TDMTCR)

The TDM transmit control register, shown in Figure 25-14, controls the activation/deactivation of the TDM transmitter. The activation of the transmitter is valid only when the TENS bit is clear.



Figure 25-14. TDM Transmit Control Register

Table 25-17 defines the bit fields of TDMTCR.

Table 25-17. TDMTCR Bit Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–30 | — | Reserved. Write to zero for future compatibility. |
| 31 | TEN | Transmit enable. Determines whether the transmit TDM is enabled or disabled. Note: Setting this bit is the last step in initializing the transmitter. 0 Transmitter is disabled. 1 Transmitter is enabled. |

25.4.2.2.6 TDM Receive Interrupt Enable Register (TDMRIER)

The TDM receive interrupt enable register, shown in Figure 25-15, has the same bit format as the TDMRER registers. If an RIER bit is clear, the corresponding event in the TDMRER registers is masked.




Figure 25-15. TDM Receive Interrupt Enable Register

Table 25-18 defines the bit fields of TDMRIER.

Table 25-18. TDMRIER Bit Descriptions

| Bits | Name | Description |
|------|-------|---|
| 0–22 | — | Reserved. Write to zero for future compatibility. |
| 23 | RCEUE | Receive channel enable update enable. Enable assertion of an interrupt when the receive channel enable update (TDMRER[RCEU]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 24 | RLCEE | Receive last channel event enable. Enable assertion of an interrupt when the receive last channel (TDMRER[RLC]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 25 | RFSEE | Receive frame sync event enable. Enable assertion of an interrupt when the receive frame sync (TDMRER[RFS]) bit is set (see page 19-57). 0 Event is masked. 1 Event is enabled. |

Table 25-18. TDMRIER Bit Descriptions (continued)


| Bits | Name | Description |
|-------|-------|---|
| 26 | RFFEE | Receive FIFO full event enable. Enable assertion of an interrupt when the receive FIFO Full (TDMRER[RFF]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 27 | RDREE | Receive data ready event enable. Enable assertion of an interrupt when the receive data ready (TDMRER[RDR]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 28 | RSEEE | Receive sync error event enable. Enable assertion of the receive error interrupt when the receive sync error (TDMRER[RSE]) bit is set. 0 Receive sync error is masked. 1 Receive sync error is enabled. |
| 29 | ROEE | Receive overrun event enable. Enable assertion of an interrupt when the receive overrun event (TDMRER[ROE]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 30–31 | — | Reserved. Write to zero for future compatibility. |

25.4.2.2.7 TDM Transmit Interrupt Enable Register (TDMTIER)

The TDM transmit interrupt enable register, shown in [Figure 25-16](#), has the same bit format as the vTER registers. If a TDMTIER bit is clear, the corresponding event in the TDMTER is masked.

Offset 0x08C

Access: Read/Write


Figure 25-16. TDM Transmit Interrupt Enable Register

[Table 25-19](#) defines the bit fields of TDMTIER.

Table 25-19. TDMTIER Bit Descriptions

| Bits | Name | Description |
|------|-------|---|
| 0–22 | — | Reserved. Write to zero for future compatibility. |
| 23 | TCEUE | Transmit channel enable update enable. Enable assertion of an interrupt when the transmit channel enable update (TDMTER[TCEU]) bit is set. 0 Event is masked. 1 Event is enabled. |

Table 25-19. TDMTIER Bit Descriptions (continued)

| Bits | Name | Description |
|-------|-------|--|
| 24 | TLCEE | Transmit last channel event enable. Enable assertion of an interrupt when the transmit last channel (TDMTER[TLC]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 25 | TFSEE | Transmit frame sync event enable. Enable assertion of an interrupt when the transmit frame sync (TDMTER[TFS]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 26 | TFEEE | Transmit FIFO empty event enable. Enable assertion of an interrupt when the transmit FIFO empty (TDMTER[TFE]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 27 | TDREE | Transmit data register empty event enable. Enable assertion of an interrupt when the transmit data ready (TDMTER[TDR]) bit is set. 0 Event is masked. 1 Event is enabled. |
| 28 | TSEEE | Transmit sync error event enable. Enable assertion of the transmit error interrupt when the transmit sync error (TDMTER[TSE]) bit is set. 0 Transmit sync error interrupt is masked. 1 Transmit sync error interrupt is enabled. |
| 29 | TUEE | Transmitter underrun error enable. Enable assertion of an interrupt when the transmitter underrun error (TDMTER[TUE]) bit is set. 0 Underrun error is masked. 1 Underrun error is enabled. |
| 30–31 | — | Reserved. Write to zero for future compatibility. |

25.4.2.3 Status Registers

This section discusses the following registers:

- [Section 25.4.2.3.1, “TDM Receive Event Register \(TDMRER\)”](#)
- [Section 25.4.2.3.2, “TDM Transmit Event Register \(TDMTER\)”](#)
- [Section 25.4.2.3.3, “TDM Receive Status Register \(TDMRSR\)”](#)
- [Section 25.4.2.3.4, “TDM Transmit Status Register \(TDMTSR\)”](#)

25.4.2.3.1 TDM Receive Event Register (TDMRER)

The TDM receive event register, shown in [Figure 25-17](#), contains the status of the receive data buffers and general receive events. The register can be read at any time. Depending on the event, some bits are cleared by writing a 1 to the register and some bits are cleared by reading the RDREG register. Bits are cleared as defined for each bit in [Table 25-20](#). If they are cleared by writing a 1, then writing a 0 has no effect. Only the receive sync error and the receive overrun error cause a receive error interrupt. All others cause a normal receive interrupt. Interrupts occur only when the receive sync synchronization status is in the “SYNC” state (refer to TDMRSR[RSSS]) except as defined below.




Figure 25-17. TDM Receive Event Register

Table 25-20 defines the bit fields of TDMTIER.

Table 25-20. TDMRER Bit Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–22 | — | Reserved. Write to zero for future compatibility. |
| 23 | RCEU | Receive channel enable update. The RCEU bit is set when the RXCHEN0–3 register is updated for the current frame, even if it is the same value. Bit is cleared by writing a 1 to this bit location in the register. Valid during all states except “HUNT” (refer to TDMRSR[RSSS]). 0 The RXCHEN0–3 was not updated for the current frame. 1 The RXCHEN0–3 was updated for the current frame. |
| 24 | RLC | Receive last channel. The RLC bit is set when the start of the last channel of the frame is being received, regardless of whether the last channel is enabled. It can happen between the first and fourth bit of the last received channel. The next word in the RDREG will be the last channel, if the channel is enabled. Bit is cleared by writing a 1 to this bit location in the register. Valid while the receiver is enabled. 0 The last channel is not being received. 1 The last channel is being received. |
| 25 | RFS | Receive frame sync. The RFS bit is set when the Rx frame sync is received. Bit is cleared by writing a 1 to this bit location in the register. Valid during all states except “HUNT” (refer to TDMRSR[RSSS]). 0 The first channel is not being received. 1 The first channel is being received. |
| 26 | RFF | Receive FIFO full. When the receiver is programmed to use the Rx FIFO, this bit will be set when the Rx FIFO has reached the Rx FIFO watermark. The RFF bit is cleared by reading the RDREG. 0 The Rx FIFO has not reached the Rx FIFO watermark. 1 The Rx FIFO has reached the Rx FIFO watermark. |
| 27 | RDR | Receive data ready. This flag bit is set when receive data register (RDREG) or receive FIFO is loaded with a new value. RDR is cleared by reading the RDREG register. If RXFIFO is enabled, RDR is cleared when receive FIFO is empty. 0 No data loaded into the RDREG or Rx FIFO. 1 New data loaded into the RDREG or Rx FIFO. |

Table 25-20. TDMRER Bit Descriptions (continued)

| Bits | Name | Description |
|-------|------|--|
| 28 | RSE | Receive sync error. Indicates whether a sync error has occurred. RSE is set when the receive frame synchronization is lost (the synchronization state change from SYNC to HUNT state) because a frame sync arrived early or it not was not received at the expected position. During operation, this bit indicates glitches on the receive pins of the TDM module. For details, see Section 25.7.1.3, “Serial Interface Synchronization.” Bit is cleared by writing a 1 to this bit location in the register. 0 Normal operation. No receive error has occurred. 1 Receive sync error has occurred. |
| 29 | ROE | Receive overrun error. Indicates whether an overrun event has occurred in the TDM. It indicates that the Rx shift register loaded the current data by overwriting the oldest data in the RDREG or the Rx FIFO. If running in wide mode, the entire line is thrown out when the current channel is loaded. This may cause a decrease in the Rx FIFO count until the current line is filled. This error should not occur during normal operation. Bit is cleared by reading from the RDREG and then by writing a 1 to this bit location in the register. 0 No overrun event has occurred in the TDM local memory. 1 An overrun event has occurred in the TDM local memory. |
| 30–31 | — | Reserved. Write to zero for future compatibility. |

25.4.2.3.2 TDM Transmit Event Register (TDMTER)

The TDM transmit event register, shown in [Figure 25-18](#), contains the status of the transmit data buffers and general transmit events. The register can be read at any time. Depending on the event, some bits are cleared by writing a 1 to the register and some bits are cleared by writing to the TDREG register. Bits are cleared as defined for each bit in [Figure 25-18](#). If they are cleared by writing a 1, then writing a 0 has no effect. Only the transmit sync error and the transmit underrun error cause a transmit error interrupt. All others cause a normal transmit interrupt. Interrupts occur only when the transmit sync synchronization status is in the “SYNC” state (refer to TDMTSR[TSSS]) except as defined below.




Figure 25-18. TDM Transmit Event Register

Table 25-21 defines the bit fields of TDMTER.

Table 25-21. TDMTER Bit Descriptions

| Bits | Name | Description |
|-------|------|--|
| 0–22 | — | Reserved. Write to zero for future compatibility. |
| 23 | TCEU | Transmit channel enable update. The TCEU bit is set when the TCEN0–3 or TCMA0–3 registers are updated for the current frame, even if it is the same value. Bit is cleared by writing a 1 to this bit location in the register. Valid during all states except “HUNT” (refer to TDMTSR[TSSS]). 0 The TXCHEN0–3 was not updated for the current frame. 1 The TXCHEN0–3 was updated for the current frame. |
| 24 | TLC | Transmit last channel. The TLC bit is set when the start of the last channel of the frame is being transmitted, regardless of whether the last channel is enabled. It can happen between the first and fourth bit of the last transmitted channel. The next word in the TDREG will be the first channel, if that channel is enabled. Bit is cleared by writing a 1 to this bit location in the register. Valid during all states except “HUNT” (refer to TDMTSR[TSSS]). 0 The last channel is not being received. 1 The last channel is being received. |
| 25 | TFS | Transmit frame sync. The TFS bit is set when the transmit frame sync is received. Bit is cleared by writing a 1 to this bit location in the register. Valid while the transmitter is enabled. 0 The first channel is not being received. 1 The first channel is being received. |
| 26 | TFE | Transmit FIFO empty. When the transmitter is programmed to use the Tx FIFO, then this bit will be set when the Tx FIFO falls below the Tx FIFO watermark. The TFF bit is cleared by writing data to the TDREG or during TDM reset. 0 The Tx FIFO has not reached the Tx FIFO watermark. 1 The Tx FIFO has reached the Tx FIFO watermark. |
| 27 | TDR | Transmit data register empty. This flag bit is set when transmit data register (TDREG) or transmit FIFO is empty and has no values to load into the Tx shift register. TDR is cleared on a write to the TDREG. 0 Data is waiting to be transmitted in the TDREG or Tx FIFO. 1 No data is waiting to be transmitted. |
| 28 | TSE | Transmit sync error. Indicates whether a sync error has occurred. TSE is set when the transmit frame synchronization is lost (the synchronization state change from SYNC to HUNT state) because a transmit frame sync arrived early or not at the expected position. During operation, this bit indicates glitches on the transmit pins of the TDM module. For details, see Section 25.7.1.3, “Serial Interface Synchronization.” Bit is cleared by writing a 1 to this bit location in the register. 0 Normal operation. No receive error has occurred. 1 A transmit sync error has occurred. |
| 29 | TUE | Transmitter underrun error. Indicates whether an underrun event has occurred in the TDM. This error should not occur during normal operation. It indicates that the TDM has not received enough data to transmit. Bit is cleared by first writing data to the TDREG, then by writing a 1 to this bit location in the register. 0 No underrun error has occurred. 1 An underrun error has occurred. |
| 30–31 | — | Reserved. Write to zero for future compatibility. |

25.4.2.3.3 TDM Receive Status Register (TDMRSR)

The TDM receive status register, shown in [Figure 25-19](#), contains the receiver status. It indicates whether the receiver is synchronized on the receive sync, the receiver is enabled or disabled, and the Rx FIFO status.




Figure 25-19. TDM Receive Status Register

[Table 25-22](#) defines the bit fields of TDMRSR.

Table 25-22. TDMRSR Bit Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 0–25 | — | Reserved. Write to zero for future compatibility. |
| 26–28 | RFCNT | Receive FIFO counter. Number of data words in the Rx FIFO. 000 0 words in the Rx FIFO 001 1 words in the Rx FIFO 010 2 words in the Rx FIFO 011 3 words in the Rx FIFO 100 4 words in the Rx FIFO |
| 29–30 | RSSS | Receive sync synchronization status. Indicates the status of the receive sync synchronization. When the synchronization state is SYNC, the serial part synchronized on the received sync and the received data transfer to the buffer in main memory for processing. For details, see Section 25.7.1.3, “Serial Interface Synchronization.” 00 HUNT state 01 WAIT state 11 PRESYNC state 10 SYNC state |
| 31 | RENS | Receive enable status. Indicates whether all the receiver parts are enabled/disabled. The propagation of the enable/disable may be delayed because of the different clocks domains. 0 The receiver machine is disabled 1 The receiver machine is enabled |

25.4.2.3.4 TDM Transmit Status Register (TDMTSR)

The TDM transmit status register, shown in [Figure 25-20](#), contains the status of the transmitter. It indicates whether the transmitter is synchronized on the transmit sync, whether it is enabled or disabled, and the Tx FIFO status.



Figure 25-20. TDM Transmit Status Register

Table 25-23 defines the bit fields of TDMTSR.

Table 25-23. TDMTSR Bit Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0–25 | — | Reserved. Write to zero for future compatibility. |
| 26–28 | TFCNT | Transmit FIFO counter. Number of data words in the Tx FIFO. 000 0 words in the Tx FIFO 001 1 words in the Tx FIFO 010 2 words in the Tx FIFO 011 3 words in the Tx FIFO 100 4 words in the Tx FIFO |
| 29–30 | TSSS | Transmit sync synchronization status. Indicates the transmit sync synchronization status. When the synchronization state is SYNC, the serial part is synchronized on the transmit sync and new transmit data is driven out. For details, see Section 25.7.1.3, “Serial Interface Synchronization.” 00 HUNT state 01 WAIT state 11 PRESYNC state 10 SYNC state |
| 31 | TENS | Transmit enable status. Indicates whether all the transmitter parts are enabled/disabled. The propagation of the enable/disable may be delayed because of the different clocks domains. 0 The transmit machine is disabled. 1 The transmit machine is enabled. |

25.4.3 AHB Interface

The handshake between the TDM module and the core occurs via a set of registers, interrupts, and DMA requests. The TDM registers are mapped into the SB and AHB address space. This section details the registers that are located in the AHB memory map. Refer to [Section 25.4, “TDM Programming Model,”](#) for details on the registers that are located in the SB memory map. The AHB address of a TDM register is composed of two parts:

- AHB base address:
— TDM module: TDMAHB_BASE

- 8-bit register offset. The TDM module registers are divided into configuration registers, control registers, status, and data registers. The registers detailed in this section are as follows:
 - Data registers, which are the registers for the transmit and receive data and can be accessed any time.

25.4.4 AHB Memory Map

This section describes the TDM module registers in the AHB memory space, which are listed in [Table 25-24](#).

Table 25-24. TDM AHB Memory-Mapped Registers

| Address Offset | Name | Access | Reset | Section/Page |
|---|--|--------|-------------|----------------------------------|
| TDM Data—Block Base Address 0x1_6100 | | | | |
| 0x000 | TDM receive data registers (TDMRDREG) | RO | 0x0000_0000 | 25.4.4.1.1/25-28 |
| 0x008 | TDM transmit data registers (TDMTDREG) | WO | 0x0000_0000 | 25.4.4.1.2/25-29 |

This section describes the TDM module registers, which are listed below.

25.4.4.1 Data Registers

This section discusses the TDM receive data register and the TDM transmit data register.

25.4.4.1.1 TDM Receive Data Register (TDMRDREG)

The TDM receive data register, shown in [Figure 25-21](#), is the read-only 64-bit register which contains the data received by the receiver. The data in the register is based on the channel size and any encoding features. If the RXFIFO wide mode is not used and it is 16-bit data or 8-bit data decoded to 16 bits, then there is one word in the RDREG register, which must be read as 16 bits. If the RXFIFO wide mode is not used and it is 8-bit data, there is 1 byte in the RDREG, which must be read as 8 bits. If the RXFIFO wide mode is used, there is either eight 8-bit data samples or four 16-bit samples in both the RDREG, which must be read simultaneously in one 64 bit read. The data is loaded from the receive FIFO or receive shift register when any read occurs and the data is available. See [Section 25.7.5, “FIFO Configuration,”](#) for more information on the Receive FIFO wide mode.




Figure 25-21. TDM Receive Data Register

Table 25-25 defines the bit fields of TDMRDREG.

Table 25-25. TDMRDREG Bit Descriptions

| Bits | Name | Description |
|------|-------|--|
| 0–63 | RDREG | Receive channel data register 0[0–63]. The data received by TDM. This is loaded with top of stack FIFO data. |

25.4.4.1.2 TDM Transmit Data Register (TDMTDREG)

The TDM transmit data register, shown in Figure 25-22, is the write-only 64-bit register which contains the data to be transmitted for the transmitter. If the TXFIFO wide mode is not used and it is 16-bit data or 16-bit data to be encoded to eight bits, then one word must be written to the TDREG register. If the TXFIFO wide mode is not used and it is 8-bit data, then 1 byte must be written to the TDREG. If the TXFIFO wide mode is used then there is either eight 8-bit data samples or four 16-bit samples to be loaded in the TDREG, which must be written simultaneously in one 64-bit write. The data is loaded from the transmit FIFO or transmit shift register as needed after any write. See Section 25.7.5, “FIFO Configuration,” for more information on the transmit FIFO wide mode.




Figure 25-22. TDM Transmit Data Register

Table 25-26 defines the bit fields of TDMTDREG.

Table 25-26. TDMTDREG Bit Descriptions

| Bits | Name | Description |
|------|-------|--|
| 0–63 | TDREG | Transmit channel data register 0[0–63]. The data to be transmitted out of the TDM. This data goes into the FIFO or shift register. |

25.5 Clocks and Reset

The following is for conceptual purposes only and may not be the actual implementation. The TDM uses the following three clocks, illustrated in Figure 25-23 and Figure 25-24:

- Bit clock—Used to serially clock the data bits in and out of the TDM port
- Word clock—Used to count the number of data bits per word (8 or 16 bits)
- Frame clock—Used to count the number of words in a frame




Figure 25-23. TDM Clocking (8-Bit Words, 4 Time Slots/Frame)

The bit clock is used to serially clock the data. The clock is visible on the TCK, for independent transmit or shared modes and RCK, for independent receive clock operation pins. The word clock is an internal clock used to determine when transmission of an 8- or 16-bit word has completed. The word clock in turn then clocks the frame clock, which marks the beginning of each frame. The frame clock can be viewed on the TFS and RFS pins. The bit clock can be received from an TDM clock pin or can be generated from the peripheral clock passed through a divider, illustrated in [Figure 25-24](#).




Figure 25-24. TDM Clock Generation

[Table 25-27](#) shows the clock summary.

Table 25-27. Clock Summary

| Clock | Source | Characteristics |
|-------|-------------------|---|
| TCK | Internal/External | Transmit data is changed on the edge of this clock. The control register can invert the clock if required. |
| RCK | Internal/External | Receive data is captured on the edge of this clock. The control register can invert the clock if required. |
| TFS | Internal/External | Transmit frames begin with this signal. See the control register for timing options. The control register can invert this signal if required. |
| RFS | Internal/External | Receive frames begin with this signal. See the definition of the control register for timing options. |

25.5.1 TDM Clock and Frame Sync Generation

Data clock and frame sync signals can be generated internally for the TDM from an on-chip clock source or can be obtained from external from a TDM pin. If internally generated by the timer, the TDM derives bit clock and frame sync signals from this internal timer. The sources change according to the Receive and Transmit sharing as well as common TDM definitions. A programmable frame rate divider and a word length divider are used for frame rate sync signal generation.

Figure 25-25 shows the clock generator block for the transmit section. The receive section contains an equivalent circuit for its frame sync generator.




Figure 25-25. TDM Transmit Clock Generator Block Diagram

Figure 25-26 demonstrates the frame sync generator block for the transmit section. When internally generated, both receive and transmit frame sync are generated from the word clock and are defined by the TDM transmit interface register (TDMTIR) and the TDM transmit frame parameters register (TDMTFP). The receive section contains an equivalent circuit for its frame sync generator.




Figure 25-26. TDM Transmit Frame Sync Generator Block Diagram

25.5.2 Reset

The TDM is reset by a system reset.

When the TDM resets all control registers are reset and the TDM is disabled.

25.6 TDM Configurations

Figure 25-27 and Figure 25-28 illustrate the main TDM configurations. These pins support all transmit and receive functions as shown. Section 25.3, “TDM Overview,” describes the clock, frame sync, and data

timing relationships in each of the modes available. Some modes do not require the use of all six pins. In this case, these pins can be used as GPIO pins, if desired.

NOTE

The GPIO interface is a separate module which alternatively controls the function and state of the I/O pins. See the GPIO module definition for alternate functions of the I/O pins defined here.

25.6.1 Typical Configurations

The TDM connects in various configurations. [Figure 25-27](#) shows two devices that connect point-to-point. Data transmits from the device on the left to the device on the right or vice versa.




Figure 25-27. TDM Point-to-Point Configuration

[Figure 25-28](#) depicts a TDM point to multipoint configuration. Multiple devices connect on the same TDM bus, which connects to the network through a framer.

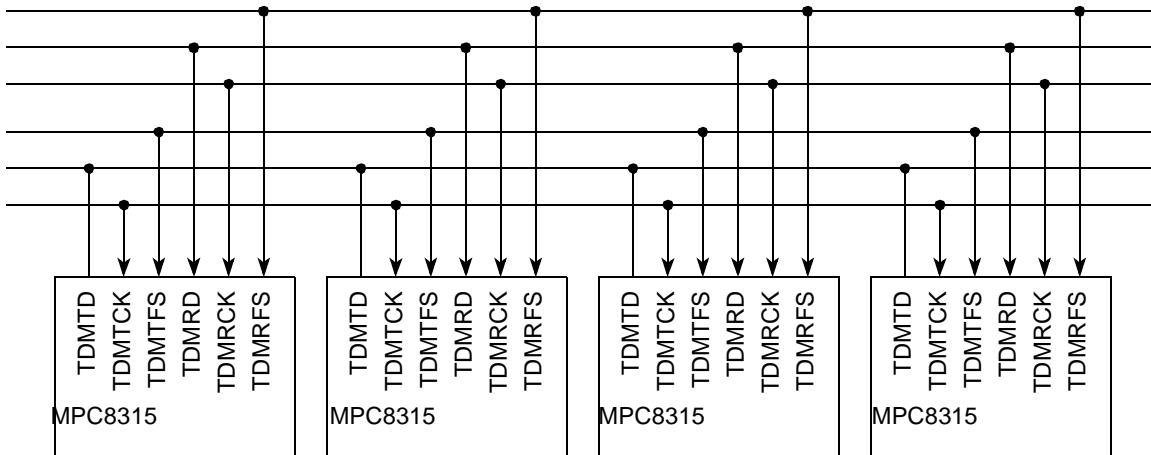


Figure 25-28. TDM Network Configuration

[Figure 25-29](#) depicts an application in which all the TDM modules share the sync and the clock. Therefore, each TDM module supports one active links. In this example, six receive links and six transmit links connect to two devices.

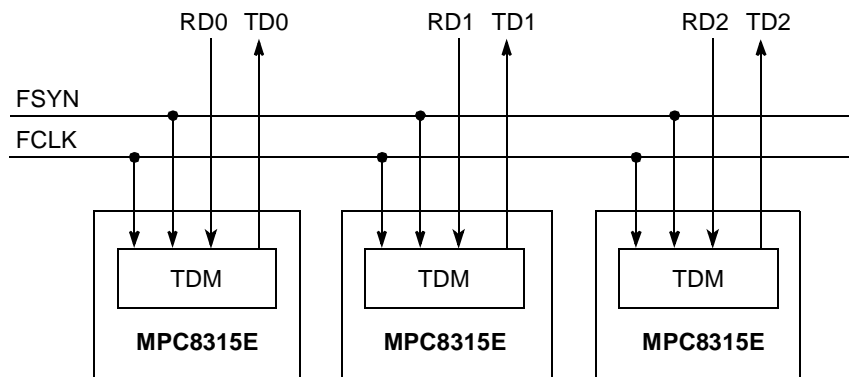


Figure 25-29. TDM Shared Network Configuration

25.7 TDM Detailed Operation

This section discusses the following:

- [Section 25.7.1, “Serial Interface”](#)
- [Section 25.7.2, “Receiver and Transmitter Independent or Shared Operation”](#)
- [Section 25.7.3, “TDM Multichannel \(Network\) Mode”](#)
- [Section 25.7.4, “Data Structures”](#)
- [Section 25.7.5, “FIFO Configuration”](#)
- [Section 25.7.6, “DMA Configuration”](#)

25.7.1 Serial Interface

This section covers issues related to the serial interface, such as how to configure the frame sync and how to control the data order of the bits in the channel word.

25.7.1.1 Sync Out Configuration

TFS, RFS are programmed as either an input or output by writing 1 to the TSO bit in the transmit interface register (TDMTIR) or writing 1 to the RSO bit in the receive interface register (TDMRIR). When the TSO, RSO bit value is equal to 1, the TFS and RFS is internally generated. Additionally, the edge the TFS and RFS are driven out is controlled by writing to the TSA bit in the transmit interface register (TDMTIR) or writing to the RSA bit in the receive interface register (TDMRIR). When these bits are set to 0, the corresponding TFS or RFS are driven out on the rising edge. When these bits are set to 1, the corresponding TFS or RFS are driven out on the falling edge.

[Figure 25-30](#) shows the sync length selection.




Figure 25-30. Sync Length Selection

25.7.1.2 Sync In Configuration

TFS, RFS are inputs that identify the beginning of the frame. [Figure 25-31](#) and [Figure 25-32](#) illustrate the relationship between the data, the sync, and the clock for various configurations. The receive data and frame sync are sampled with the rising or falling edge of the receive clock. The transmit frame sync, TFS, is sampled with the rising or falling edge of the transmit clock. The transmit data drives out at the rising or falling edge of the transmit clock.

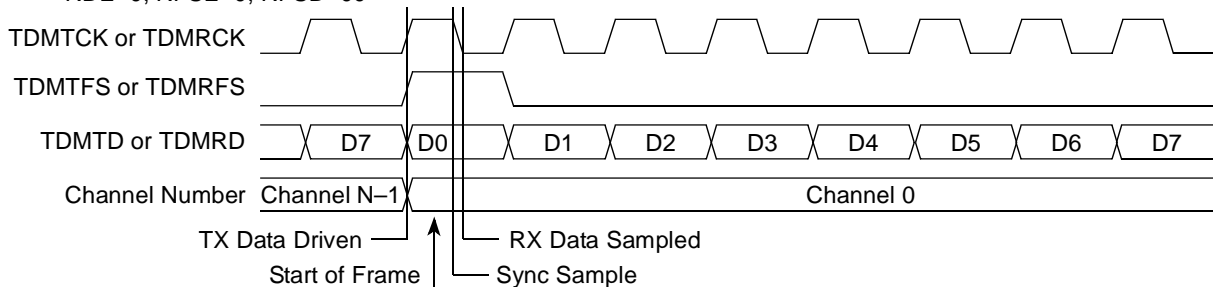
To determine the (R/T)DE, (R/T)FSE, and the (R/T)FSD bits:

1. Determine the edge that the frame sync is driven on (or changes), whether or not it is internally generated, to set the appropriate (R/T)FSE accordingly.
2. Determine the edge that the data is driven on (or changes), whether or not it is internally generated, to set the appropriate (R/T)DE accordingly.
3. If the (R/T)DE = (R/T)FSE, then set the (R/T)FSD to the number of clock cycles that the first channel's first bit of data arrives after the first part of the frame sync edge. Otherwise, set the (R/T)FSD to the number of clock cycles that the first channel's first bit of data arrives after the first clock edge that the frame sync is active and stable.


The TDM frame syncs when configured as inputs only require that the frame sync is inactive 2 cycles before being active for a minimum of 1 cycle. Thus the TDM can support frame syncs of any length such that these requirements are met.

Data and Sync Change on the Rising Edge (No Sync Delay)


TDE=0, TFSE=0, TFSD=00
RDE=0, RFSE=0, RFSD=00


Data and Sync Change on the Falling Edge (No Sync Delay)

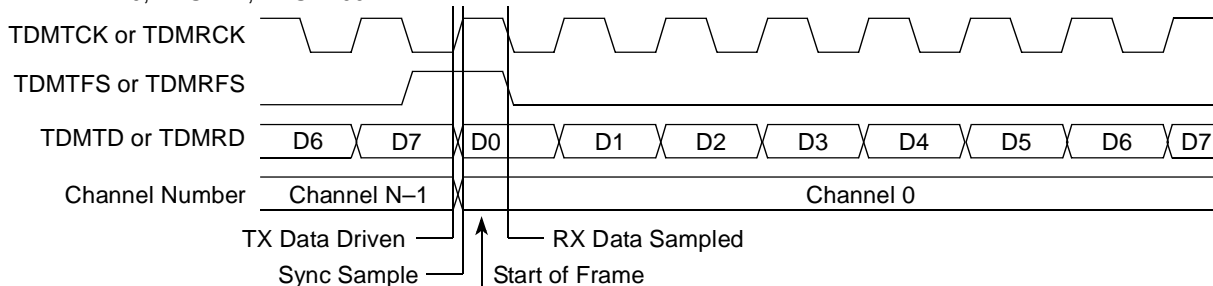
TDE=1, TFSE=1, TFSD=00
RDE=1, RFSE=1, RFSD=00


Data Changes on the Falling Edge, Sync on the Rising Edge (No Sync Delay)

TDE=1, TFSE=0, TFSD=00
RDE=1, RFSE=0, RFSD=00

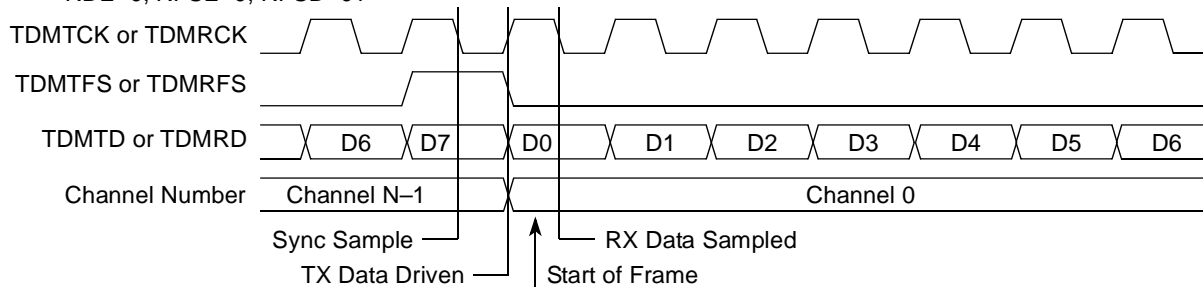

Data Changes on the Rising Edge, Sync on the Falling Edge (No Sync Delay)

TDE=0, TFSE=1, TFSD=00
RDE=0, RFSE=1, RFSD=00


Figure 25-31. Frame Sync Configurations without Sync Delays

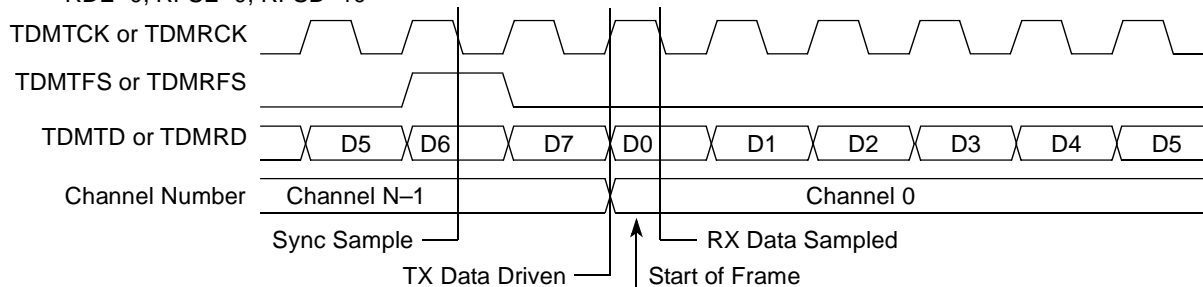
Data and Sync Change on the Rising Edge (One Bit Delay)

TDE=0, TFSE=0, TFSD=01
RDE=0, RFSE=0, RFSD=01



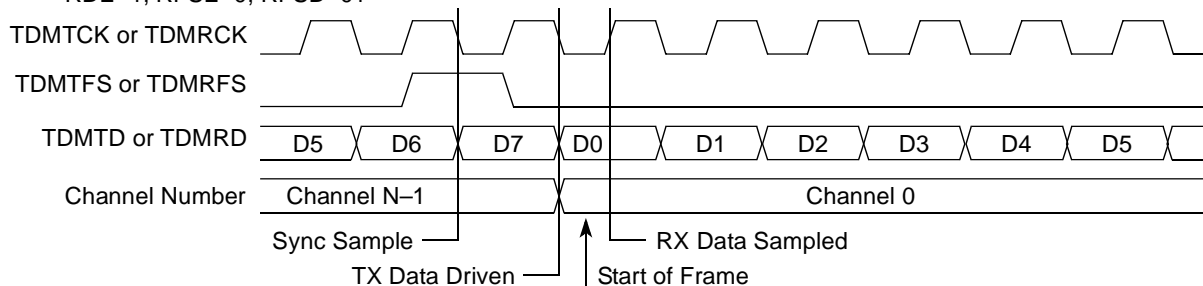
Data and Sync Change on the Rising Edge (Two Bit Delay)

TDE=0, TFSE=0, TFSD=10
RDE=0, RFSE=0, RFSD=10



Data Changes on the Falling Edge, Sync on the Rising Edge (One Bit Delay)

TDE=1, TFSE=0, TFSD=01
RDE=1, RFSE=0, RFSD=01



Data Changes on the Falling Edge, Sync on the Rising Edge (One Bit Delay)

TDE=1, TFSE=0, TFSD=01
RDE=1, RFSE=0, RFSD=01

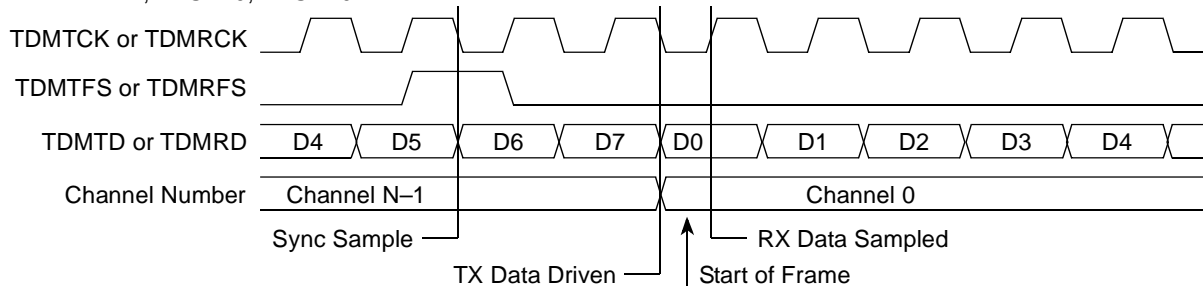


Figure 25-32. Frame Sync Configurations with Sync Delays

No Sync Delay (the Data and the Sync Sample with the Same Edge)

RDE=0, RFSE=0, RFSD=00

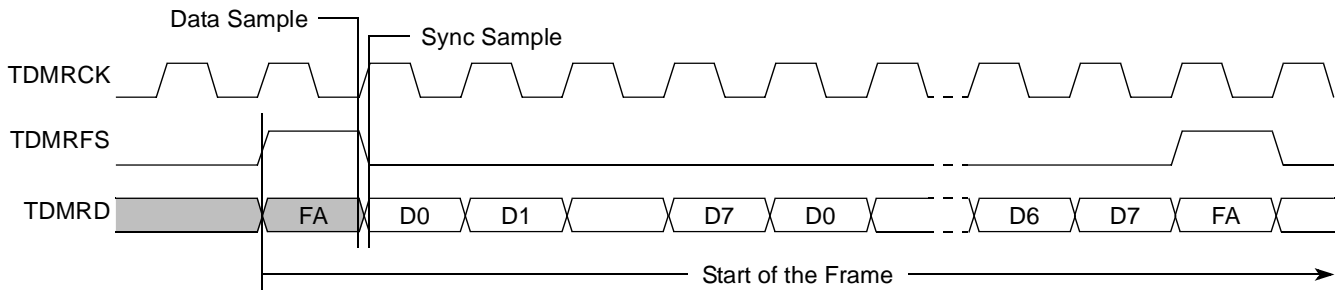



Figure 25-33. Frame Sync Configuration At T1 mode

Figure 25-34 illustrates how the polarity of the receive sync and the transmit sync signals is controlled by the TDMRIR[RSA] and the TDMTIR[TSA] bits.

Sync Level Active High (RSA/TSA=0)



Sync Level Active Low (RSA/TSA=1)




Figure 25-34. Frame Sync Polarity

25.7.1.3 Serial Interface Synchronization

The TDM module enables communication among many devices over a single bus. The receive and transmit of each TDM frame is identified by a frame sync signal that is asserted at the beginning of every frame. The frame sync synchronization is necessary when more than one device drives the bus.

Figure 25-35 shows the state diagram of the frame sync synchronization.

The details of the state diagram are as follows:

- HUNT (0b00). A sync event is constantly sought. As soon the sync event is detected, the state machine changes to a WAIT state. During the Hunt state, data is neither received nor transmitted.
- WAIT (0b01). At least one sync has been detected. The next sync event is accepted after one TDM frame. If the sync appears in the correct position, the state changes to the PRESYNC state (0b11). If the sync does not appear, the state returns to the hunt state. During the WAIT state, data is neither received nor transmitted.
- PRESYNC (0b11). Two sync events have been detected and the distance between the syncs is one TDM frame. If the sync event is recognized early, the state returns to the WAIT state. Otherwise,

the machine transfers to the SYNC state at the last bit of the TDM frame. During PRESYNC state, data is neither received nor transmitted.

- SYNC (0b10). At least one sync event has appeared exactly where it was expected. This state is maintained as long as the sync event continues to appear where expected. If a sync is missed or a sync event is recognized early, the state changes to the HUNT state (0b00). During the SYNC state, data is both received and transferred.




Figure 25-35. Frame Sync Synchronization State Diagram

The TDM receiver synchronizes on the receive frame sync (RFS). The state of the receive frame sync synchronization is indicated by the TDMRSR[RSSS] field. During the HUNT, WAIT, and PRESYNC states, the received data is not transferred to the Receive Data Register. When the receive sync synchronization is lost, the state transfers from SYNC to HUNT (the TDMRER[RSE] bit is asserted). If the TDMREIR[RSEEE] bit is also set, a receive error interrupt is generated.

The transmit frame sync synchronization state is indicated by the TDMTSR[TSSS] field. During the HUNT, WAIT, and PRESYNC states, new data is not driven out. If the Transmit Always Out (TDMTIR[TAO]) field is set, then all 1's are driven out until the frame sync synchronization state returns to SYNC state. If the TDMTIR[TAO] bit is clear, data is not driven out and TD is three-stated. When the transmit sync synchronization is lost, the TDMTER[TSE] bit is asserted. If the TDMTIER[TSEEE] bit is also set, a transmit error interrupt is generated.

The frame sync synchronization state can identify different problems. In the initial design stages, the frame sync summarization state indicates whether the TDM programming matches the actual TDM stream. During operation, the synchronization state and the error interrupts may indicate errors in the TDM module signal processing.

25.7.1.4 Reverse Data Order

Figure 25-36 illustrates how the bit order of the stored data relates to the bit order of the receive or the transmit data. The TDMRIR[RRDO] bit defines how the receive channel data is stored in memory. If

TDMRIR[RRDO] is clear, the first bit of the received channel data is stored as the least significant bit. The TDMTIR[TRDO] bit selects the transmit data bits order. If TDMTIR[TRDO] is clear, the least significant bit of the memory is transmitted as the first transmit data.

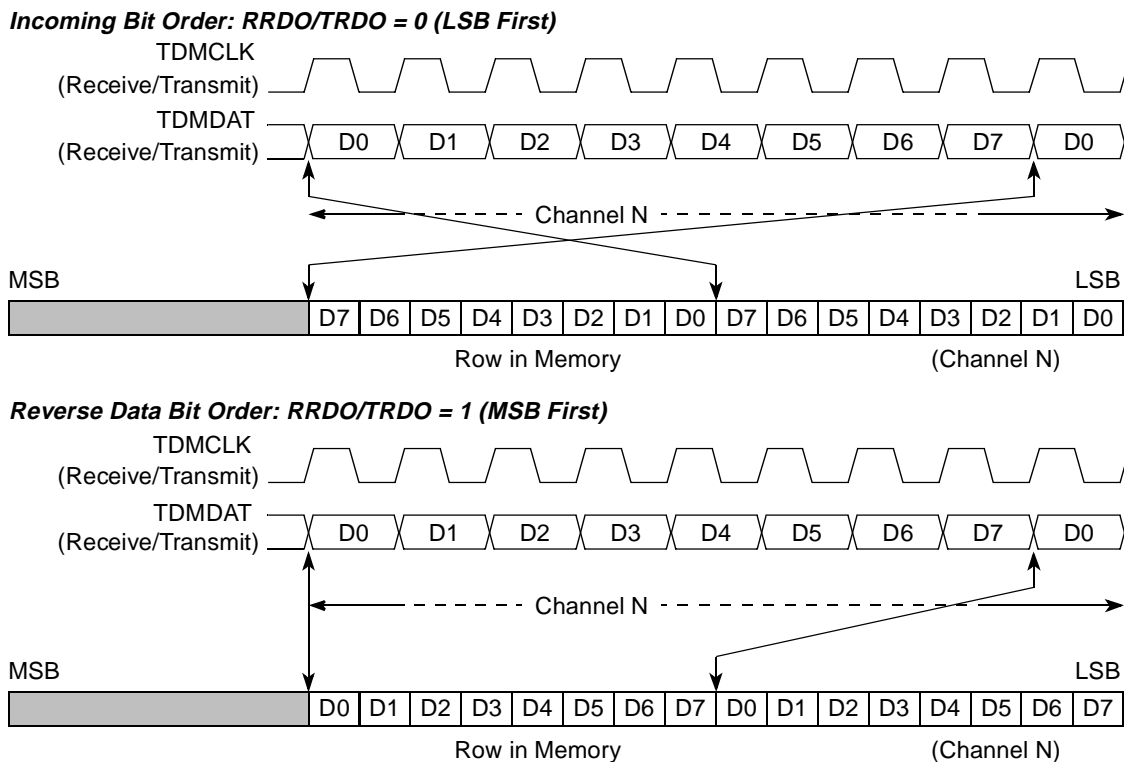
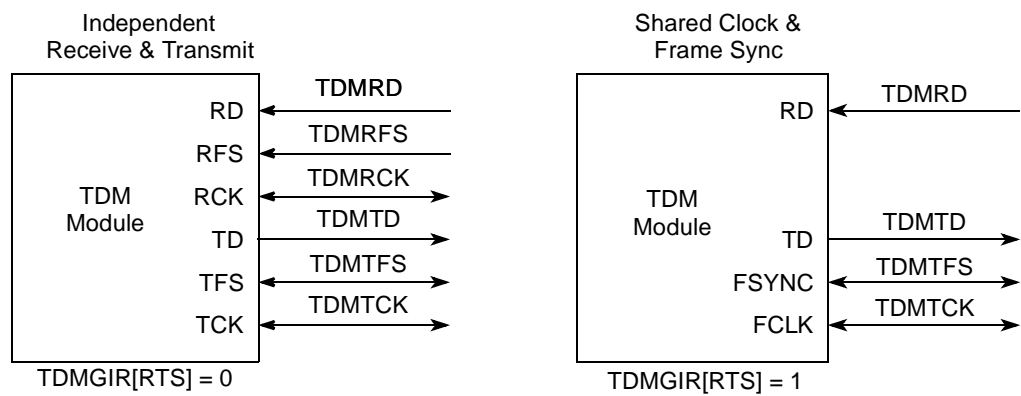


Figure 25-36. Reserve Bit Order

25.7.2 Receiver and Transmitter Independent or Shared Operation

The TDM operates with the transmit and receive operations running either independently or shared, as illustrated in Figure 25-37. When the TDMGIR[RTS] bit is cleared, the receive and the transmit are independent as illustrated on the left side of Figure 25-37. When TDMGIR[RTS] is set, the transmit and the receive share the frame sync (FSYN) and the frame clock (FCLK) signals with the transmitter being

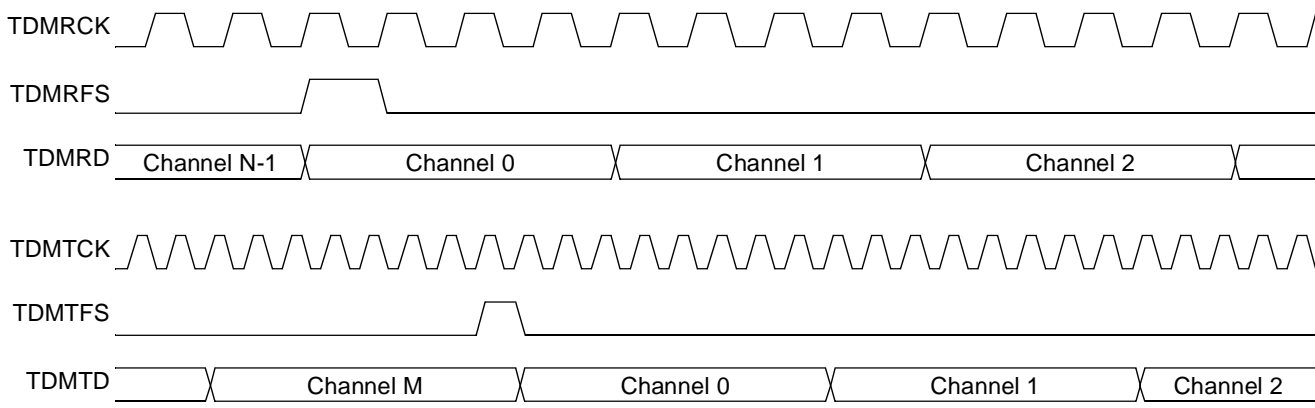
the input or the output, as illustrated on the left side of [Figure 25-37](#). In this mode, there is one input receive data link and one output transmit data link.



FSYNC (Frame Sync) specifies that the receiver and transmitter share the same sync.
 FCLK (Frame Clock) specifies that the receiver and transmitter share the same clock.

Figure 25-37. TDM Module Modes

[Figure 25-38](#) describes the TDM interface when the receive and transmit are totally independent ($TDMGIR[RTS] = 0$). The TDMRCK is not synchronized to the TDMTCK. They differ according to the sync location relative to the beginning of the frame and the number of bits in a word.



n – Defines the TDM number.
 N – The number of channels in the receive TDM frame.
 M – The number of channels in the transmit TDM frame.

Figure 25-38. Receive and Transmit Totally Independent

25.7.3 TDM Multichannel (Network) Mode

Together, [Figure 25-39](#) and [Figure 25-40](#) illustrate sample timing of TDM multi-channel (network) mode transfers. The numbered circles and arrows in the figures identify discussion notes contained in [Table 25-28](#) and [Table 25-29](#).

25.7.3.1 Operation Using Tx Channel Mask Register

If all bits of the TCMA registers are cleared, the TDM transmitter will continue to operate based solely on the TCEN registers. The TCMA registers are used to throw out data on specific channels. Therefore TCMA should only be set on channels that have their corresponding TCEN bits already cleared, or the channel is disabled. Once the TCEN is cleared, the TCMA is set to 1 in the desired channel bit location for data to be discarded by the TDM. This allows for memory structures to keep channels open that are physically closed on the TDM, which are transferred to the TDM via a DMA, where the TDM discards the data. If the TCMA is set on a channel that is enabled, then the data for that channel is discarded by the TDM and the TDM will instead transmit all 1s.

Disabling a channel in this manner causes the time slot for that channel to be ignored by the TDM. This means no data is transferred to the transmit shift register. The transmitter timing, using TCM registers for an 8-bit word, continuous clock with disabled FIFO six words per frame sync in the TDM network mode is illustrated in [Figure 25-39](#). Explanatory notes for the transmit portion of the figure are provided in [Table 25-28](#).

NOTE

In this example there are only three transmit interrupts per frame instead of six as in the previous example where the TCM register is not used.




Figure 25-39. TDM Network Mode Transmit Timing with Mask Register

Table 25-28. Notes for Transmit Timing with Mask Register in Figure 25-39

| Note | Source Signal | Destination Signal | Description |
|------|-------------------|--------------------|--|
| 1 | — | — | Example of a 6 time slot frame, transmitting in channels 0, 3, and 5. |
| 2 | TFS | — | Example of a word length frame sync and standard timing. Frame timing begins with the rising edge of TFS. |
| 3 | — | TDR Interrupt | For enabled time slots, this flag is set at the beginning of each word to indicate the TDREG data has been used and another data word should be supplied by the software. If the transmit interrupt is enabled the processor is interrupted to request the data. |
| 4 | Tx Data Register | Tx Shift Register | On each word clock boundary a decision is made concerning what to transmit on the next time slot. If the TCEN register bit is 0 for the next time slot, the TD pin is either tri-stated or drives previous time slot data and the time slot is ignored. When the TCEN bit is 1 for the next time slot, the contents of the TDREG register are transferred to the TXSR register and this data is shifted out. If the TDREG register has not been written in the previous time slot, the previous data is reused. When neither of these registers were written in the previous time slot (where TCEN = 1), the TUE status bit will be set and the hardware will operate as if the TDREG register had been written. The TD pin will be enabled and the contents of the TDREG will be transmitted again. |
| 5 | Tx Shift Register | TD Pin | On active time slots, the TXSR register contents are shifted out on the TD pin, one bit per rising edge of TCK. On inactive time slots, the TD pin is tri-stated so it can be driven by another device. |
| 6 | Tx Data Register | Tx Shift Register | On disabled time slots (TCEN = 0) and the channel mask is enabled (TCMA = 1), The contents of the TDREG register are discarded and the TDR status flag is set as if the data was shifted out. |

25.7.3.2 Operation Using Rx Channel Enable Register

When the RCEN registers are utilized in the design, interrupt overhead can be reduced. If all bits of the RCEN register are set, the TDM receiver will receive all of the channels in the frame. The RCEN register is used to automatically discard data from selected channels. This is accomplished by writing the RCEN with 0 in the selected channel bit location. This means no data is transferred from the receive data shift register on these channels, no status flags change, and no interrupts are generated.

Receiver timing using RCEN registers for an 8-bit word with a continuous clock, with the disabled FIFO, six words per frame sync. In this example there are only three receive interrupts per frame instead of six

as in the previous example where the RCEN register is not used. This process is illustrated in Figure 25-40 and explained in Table 25-29.




Figure 25-40. TDM Network Mode Receive Timing with Enable Register

Table 25-29. Notes for Receive Timing with Mask Register in Figure 25-40

| Note | Source Signal | Destination Signal | Description |
|------|-------------------|--------------------|---|
| 6 | — | — | Example of a six time slot frame, receive data from channels 0, 2, and 5. The receive hardware will only obtain data on the RD pin when the RCMn bit is set to 1. |
| 7 | RFS | — | Example with bit length frame sync and standard timing. Frame timing begins with the rising edge of RFS. |
| 8 | RD | Rx Shift Register | Data on the RD pin is sampled on the falling edge of RCK and shifted into the Rx shift register. |
| 9 | Rx Shift Register | Rx Data Register | At the word clock, the data in the Rx shift register is transferred to the Rx register, for enabled time slots. |
| 10 | RDR Interrupt | — | This flag is set for each word clock, or time slot, where the RCEN bit is set, indicating data is available to be processed. The software must keep track of the time slots as they occur so it knows which data it is processing. When the receive interrupts are enabled an interrupt will be generated when this status flag is set. The software reads the RDREG register to clear the interrupt. |

25.7.4 Data Structures

There are four types of data received and transmitted by the TDM. There is 16-bit data unencoded, 8 bit data unencoded, 8-bit data A-law encoded and 8 bit data μ -law encoded. The A-/ μ -law conversion is

performed according to the ITU-T recommendation G.711. The channel size and encoding is identical for all channels in the frame.

When the RCS field in the TDMRFP register is set to receive an 8-bit A-law encoded channel, the 8 bit PCM compressed sample that is received is converted into a 13-bit PCM linear sample padded with three zeros on the right to create a 16-bit data structure. When the RCS field in the TDMRFP register is set to receive an 8-bit μ -law encoded channel, the 8 bit PCM compressed sample that is received is converted into a 14-bit PCM linear sample padded with two zeros on the right to create a 16-bit data structure.

When the TCS field in the TDMTFP register is set to transmit an 8-bit A-law encoded channel, the 16 bits (13-bit linear sample padded with three zeros on the right) to be transmitted are converted into a 8-bit PCM A-law compressed sample, which is transmitted out of the TDM. When the TCS field in the TDMTFP register is set to transmit an 8 bit μ -law encoded channel, the 16 bits (14-bit linear sample padded with three zeros on the right) to be transmitted are converted into a 8-bit PCM μ -law compressed sample, which is transmitted out of the TDM.

25.7.5 FIFO Configuration

The TDM has a FIFO in both the Tx and Rx paths. After reset the FIFO is bypassed and disabled. To enable the FIFO set the Rx FIFO Enable bit in the TDMRIR register and the Tx FIFO enable bit in the TDMTIR register. The FIFO watermark registers are used to configure when the FIFO will set event bits and/or interrupts. The RFWM and TFWM in the TDMRIR and TDMTIR registers can be set according to how much data is generally held in the FIFO. The FIFO has a depth of four lines. The Rx FIFO is cleared when the TDM Rx is disabled, and the Tx FIFO is cleared when the TDM Tx is disabled.

The width of each line in the FIFO is set to be either 1 sample of 16-/8-bit data or 64 bits packed with eight 8-bit samples or four 16-bit samples, as seen in [Figure 25-41](#). The width of the data must be accessed in one read/write transaction. This is configured via the RWEN and TWEN bits in the TDMRIR and TDMTIR registers. When using the wide FIFO mode, the channels are packed from the lowest address to the highest address in the RDREG. When a new frame comes in the Rx shift register will stop filling the current line in the FIFO and move to the next line in the FIFO. This only applies to the wide FIFO mode. On the transmit side, the wide FIFO mode will read data sequentially from the lowest address to the highest address. If a frame ends in the middle of a FIFO line, the remainder of the line is bogus data. The

transmitter will not transmit this bogus data. The receiver will not put data in the remainder of the line and should be ignored on reads.

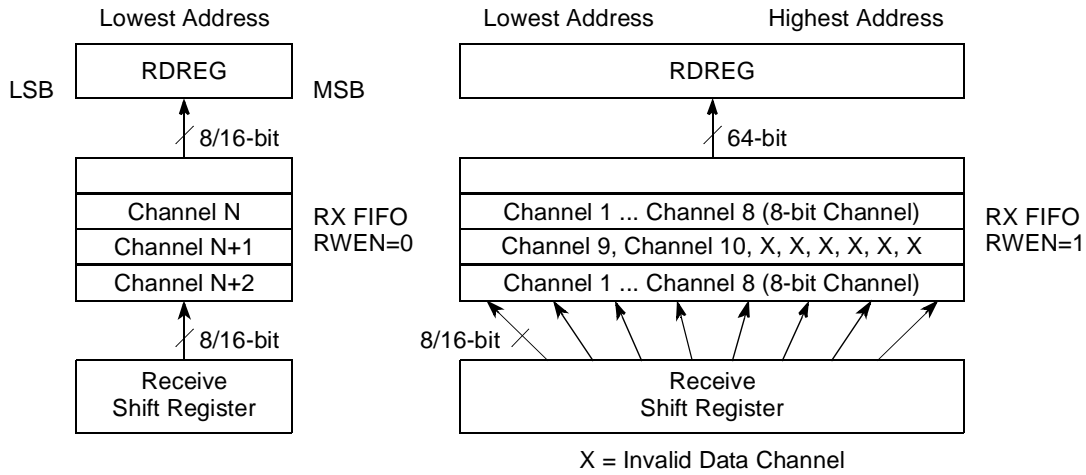


Figure 25-41. TDM Rx FIFO Using RWEN Bit with Ten 8-Bit Channels

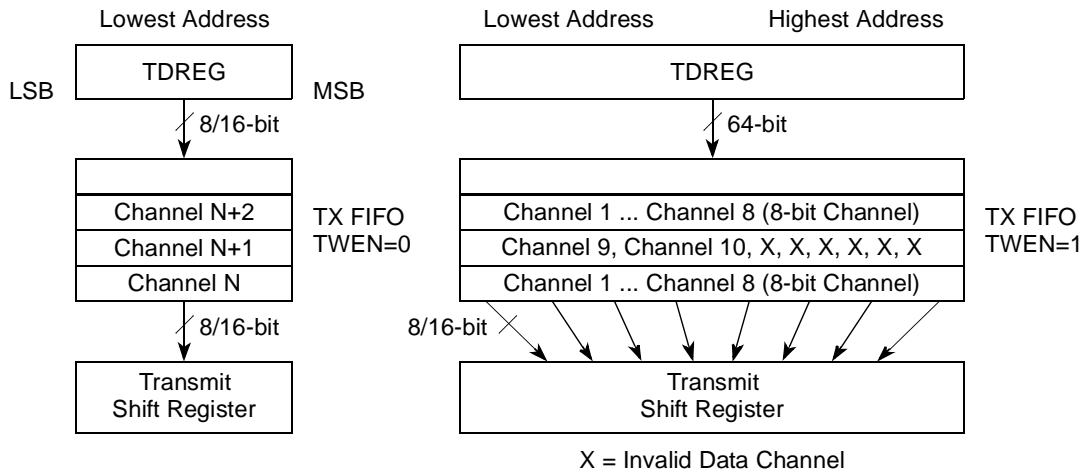


Figure 25-42. TDM Tx FIFO Using TWEN Bit with Ten 8-Bit Channels

25.7.6 DMA Configuration

See [Section 25.10, “DMA Controller \(DMAC\),”](#) for details on the DMA controller.

NOTE

The DMAC channel 0 is used for TDM transmitter and channel 1 is used for TDM receiver. All other references to channel 63–2 in [Section 25.10, “DMA Controller \(DMAC\),”](#) is not available for the MPC8315E.

25.7.6.1 Setting up the TDM for Correct Operation with the DMA

The TDM can be configured to work with the DMA to transfer data in and out of the TDREG and RDREG. The DMA must be configured to transfer the appropriate amount of data for the TDREG or RDREG based

on the channel size, encoding, and whether FIFO wide mode is being used. The DMA interface is enabled by setting the RDMA (see [Section 25.4.2.1.2, “TDM Receive Interface Register \(TDMRIR\)”](#)) or the TDMA (see [Section 25.4.2.1.3, “TDM Transmit Interface Register \(TDMTIR\)”](#)). Once set to be active, [Table 25-30](#) describes the conditions where the DMA would be requested to service the TDM data registers. Once the DMA asserts the DMA acknowledge signal, the requests will de-assert if the event flags clear (no more data to service). See [Chapter 12, “DMA/Messaging Unit,”](#) regarding DMA operation and specific priorities of the channels.

Table 25-30. DMA Request Channels for the TDM

| Source of DMA Request | Event in TDM Causing Request | Description |
|-----------------------|--------------------------------|---|
| TDM—Rx | Rx Data Ready or Rx FIFO Full | If the Rx FIFO is not enabled, the Rx data ready flag will cause the Rx DMA request to go active. If the Rx FIFO is enabled, the Rx FIFO Full flag will cause the Rx DMA request to go active. |
| TDM—Tx | Tx Data Empty or Tx FIFO Empty | If the Tx FIFO is not enabled, the Tx data register empty flag will cause the Tx DMA request to go active. If the Tx FIFO is enabled, the Tx FIFO Empty flag will cause the DMA request to go active. |

25.8 Software Programming

This section discusses the following:

- [Section 25.8.1, “Software Programming Sequence”](#)
- [Section 25.8.2, “Interrupts”](#)

25.8.1 Software Programming Sequence

This section provides useful information for software developers writing code to use the TDM. It contains the following:

- Initialization
 - Shared operation starting on the same frame
 - Non-shared operation
- Dynamic channel configuration while TDM operating—non-shared
- Configuring the TDM for I²S operation
- TDM power down
- Synchronization errors
- Transmit normal and transmit error interrupts

25.8.1.1 Initialization—Shared Operation Starting on the Same Frame

The correct sequence to initialize the TDM for shared operation (that is, TDMGIR[RTS] is set) where both the Tx and Rx machines must start on the same frame is as follows:

1. Bring the chip out of reset.

2. Program the TDM control registers—TDMRIR, TDMRFP, TDMTIR, TDMTFP, and TDMGIR.
3. Program all channel enable registers (TDMRCEN0–3, TDMTCEN0–3) and transmit channel mask registers (TDMTCMA0–3) as disabled.
4. Enable the TDM Rx and Tx by setting the TDMRCCR[REN] and TDMTCCR[TEN] bits.
5. Wait until TDMTSR[TENS] and TDMRSR[RENS] are set.
6. If the TDM is serviced by:
 - a) DMA transfer—immediately use a software started DMA channel to write data to the Tx data register(s) (must be completed before two Tx frame syncs arrive after enabling).
 - b) Interrupts—immediately write data to the Tx data register(s) (must be completed before two Tx frame syncs arrive after enabling).
7. Follow the steps outlined in [Section 25.8.1.3, “Dynamic Channel Configuration while TDM Operating—Shared,”](#) to enable channels.
8. All other writes and reads will be initiated by either an interrupt or a DMA request, depending on the configuration.

This procedure ensures proper operation of the TDM by changing any of the control bits before enabling the TDM. These control bits should not be changed during TDM operation. The only registers that can be changed during operation are the TDM receive channel enable registers (TDMRCEN0–3), TDM transmit channel enable registers (TDMTCEN0–3), and the TDM transmit channel mask registers (TDMTCMA0–3). These registers are sampled at the beginning of the frame to determine how the current frame will be processed.

25.8.1.2 Initialization—Non-Shared Operation

The correct sequence to initialize the TDM is as follows:

1. Bring the chip out of reset.
2. Program the TDM control registers.
3. Enable the TDM Rx and Tx by setting the TDMRCCR[REN] and TDMTCCR[TEN] bits.
4. Wait until TDMTSR[TENS] and TDMRSR[RENS] are set.
5. If the TDM is serviced by:
 - a) DMA transfers—immediately use a software started DMA channel to write data to the Tx data register(s) (must be completed before two Tx frame syncs arrive after enabling).
 - b) Interrupts—immediately write data to the Tx data register(s) (must be completed before two Tx frame syncs arrive after enabling).
6. All other writes and all reads will be initiated by either an interrupt or a DMA request, depending on the configuration.

NOTE

This procedure also applies for shared operation when start the Tx and Rx machines on the same frame is not required.

This procedure ensures proper operation of the TDM by changing any of the control bits before enabling the TDM. These control bits should not be changed during TDM operation. The only registers that can be

changed during operation are the TDM receive channel enable registers (TDMRCEN0–3), TDM transmit channel enable registers (TDMTCEN0–3), and the TDM transmit channel mask registers (TDMTCMA0–3). These registers are sampled at the beginning of the frame to determine how the current frame will be processed.

25.8.1.3 Dynamic Channel Configuration while TDM Operating—Shared

The correct sequence in shared operation (that is, TDMGIR[RTS] is set) to change channels while the TDM is operating is to:

1. Enable the channel enable update interrupts, TDMRIER[RCEUE].
2. Read the values of the TDMRCEN0 into one of the core's data registers.
3. Immediately write this same value back out to the same register—TDMRCEN0. This creates an interrupt for receive on a frame boundary without modifying the values in this register.
4. Wait for the receive channel enable update interrupt, TDMRER[RCEU]. Upon receiving this interrupt, write the new values to the transmit and receive channel enable registers (TDMTCEN0–3 and TDMRCEN0–3) and transmit channel mask registers (TDMTCMA0–3) to turn on/off the appropriate channels. This should be completed within half of a frame after receiving the interrupt.
5. Wait until the next receive channel enable update interrupt occurs. Read all enable and mask registers to verify that the writes in Step 4 were completed in time.
6. Update any data structures for the TDM in memory.
7. Disable the channel enable update interrupt.
8. Process data as normal.

NOTE

The procedure for non-shared dynamic channel configuration must instead be used when there are only one or two physical channels per frame (that is, TDMRFP[RNCF] is \$0 or \$1).

25.8.1.4 Dynamic Channel Configuration while TDM Operating—Non-Shared

The correct sequence to change channels while the TDM is operating is to:

1. Enable the channel enable update interrupts, TDMRIER[RCEUE] and TDMTIER[TCEUE].
2. Read the values of the TDMRCEN0 and TDMTCEN0 into two of the core's data registers.
3. Immediately write these same values back out to the same registers—TDMRCEN0 and TDMTCEN0. This creates interrupts for receive and transmit on the frame boundaries without modifying the values in these registers.
4. Wait for the transmit channel enable update interrupt, TDMTER[TCEU]. Upon receiving this interrupt, write the new values to the transmit channel enable registers (TDMTCEN0–3) and transmit channel mask registers (TDMTCMA0–3) to turn on/off the appropriate channels. This should be completed within half of a frame after receiving the interrupt.

5. Wait for the receive channel enable update interrupt, TDMRER[RCEU]. Upon receiving this interrupt, write the new values to the receive channel enable registers (TDMRCEN0–3) to turn on/off the appropriate channels. This should be completed within half of a frame after receiving the interrupt.

NOTE

Steps 4 and 5 can be done in either order.

6. Wait until the next transmit channel enable update interrupt occurs. Read all enable and mask registers to verify that the writes in Step 4 were completed in time.
7. Wait until the next receive channel enable update interrupt occurs. Read all enable registers to verify that the writes in Step 5 were completed in time.

NOTE

Steps 6 and 7 can be done in either order.

8. Update any data structures for the TDM in memory.
9. Disable the channel enable update interrupts.
10. Process data as normal.

25.8.1.5 Configuring the TDM for I²S Operation

The TDM programming model is flexible enough that it can be configured for I²S operation. This can be done by setting the TDMTIR and TDMRIR registers as follows:

1. Set frame sync out length equal to channel width.
 - Set TDMTIR[TSL] to 1
 - Set TDMRIR[RSL] to 1
2. Frame sync active set as high on right channel and low on left channel.
 - Set to TDMTIR[TSA] to 0, TDMRIR[RSA] to 0
3. Set as data and sync change on the falling edge only (one bit delay)
 - Set TDMTIR[TDE] to 1, TDMTIR[TFSE] to 1, TDMTIR[TFSD] to 01
 - Set TDMRIR[RDE] to 1, TDMRIR[RFSE] to 1, TDMRIR[RFSD] to 01
4. Set as reversed data order, that is, MSB first
 - Set TDMTIR[TRDO] to 1
 - Set TDMRIR[RRDO] to 1

25.8.1.6 TDM Power Down

The correct sequence to shut off the TDM is as follows:

1. Disable all channels by clearing all the Tx and Rx channel enable bits.
2. Read all data out of the receive data registers.
3. Disable the TDM Rx and Tx by clearing the TDMRCR[REN] and TDMTCR[TEN] bits.
4. Clear all TDMRER and TDMTER interrupt/status bits by writing 1's to the register.

5. Verify TDM disabled by reading 0 for both TDMSR[RENS] and TDMTSR[TENS].

25.8.1.7 Synchronization Errors

The correct sequence after a synchronization error is as follows:

1. Read all data out of the receive data registers.
2. If it is desired to clear out the data in the Tx FIFO/ Tx data register, disable the TDM Tx by clearing the TDM Tx enable bit. Then write data to the Tx data register(s).
3. Assuming the TDM will synchronize again, all other writes and all reads will be initiated by either an interrupt or a DMA request, depending on the configuration.

25.8.2 Interrupts

25.8.2.1 Receiver Normal and Receiver Error Interrupts

These interrupts can occur when receive interrupts are enabled via the TDM receive interrupt enable register (TDMRIER). The bits in this register enable interrupts based on that particular event occurring. If more than one interrupt is enabled, then an interrupt occurs for either event. To clear the interrupt a one is written to the corresponding bit of the interrupt in the TDM receive event register (TDMRER) or reading the receive data register (TDMRDREG) depending on the type of interrupt.

25.8.2.2 Transmit Normal and Transmit Error TDMRDREG

These interrupts can occur when receive interrupts are enabled via the TDM transmit interrupt enable register (TDMTIER). The bits in this register enable interrupts based on that particular event occurring. If more than one interrupt is enabled, then an interrupt occurs for either event. To clear the interrupt a one is written to the corresponding bit of the interrupt in the TDM transmit event register (TDMTER) or writing to the transmit data register (TDMTDREG) depending on the type of interrupt.

25.9 TDM Clock Control

The TDM can be configured to use on chip clock for serial transmission/reception of data. This section provides the resultant functionality chosen in the design of the `tdm_clk_gen` module. The purpose for this module is to generate the required Tx/Rx clock frequency from the `ipg_clk`, based on programmed values.

25.9.1 Features

The main features are as follows:

- Two independent counters for receive clock and transmit clock generation
- Supports division by odd and even value
- 50% duty cycle

25.9.2 Theory of Operation

The module creates two counter that counts from 0 to N-1 (division value-1) and always clocks on the rising edge of the ipg_clock. The counters resets to 0 whenever count reaches N-1 or a new value is loaded in div_val_rx/div_val_tx registers. The counter value is used to create the combinatorial signals, en_tff_rx/en_tff_tx. These signal are enable signals for T flip-flop, clk_rx/clk_tx.

Depending on the value N, the following scenarios exist:

- N = 0: This is clock off mode. The timer_rx/timer_tx clock output are held low.
- N = 1: This is bypass mode. The ipg_clk is driven on timer_rx/timer_tx. [Figure 25-43](#) below shows the case where the division value for Rx and Tx serial clock is set to 1.

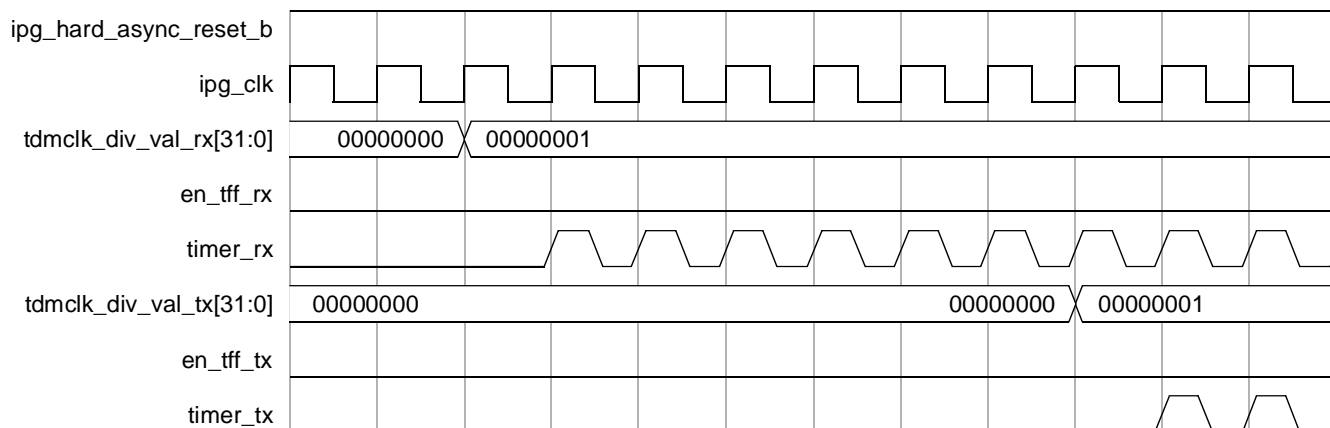


Figure 25-43. Division Value 1

- N = Even: The en_tff_rx/en_tff_tx is driven high (logic 1) when the respective counter value is 0 and N/2. [Figure 25-44](#) below shows the case where division value for Rx and Tx serial clock is set to 6.

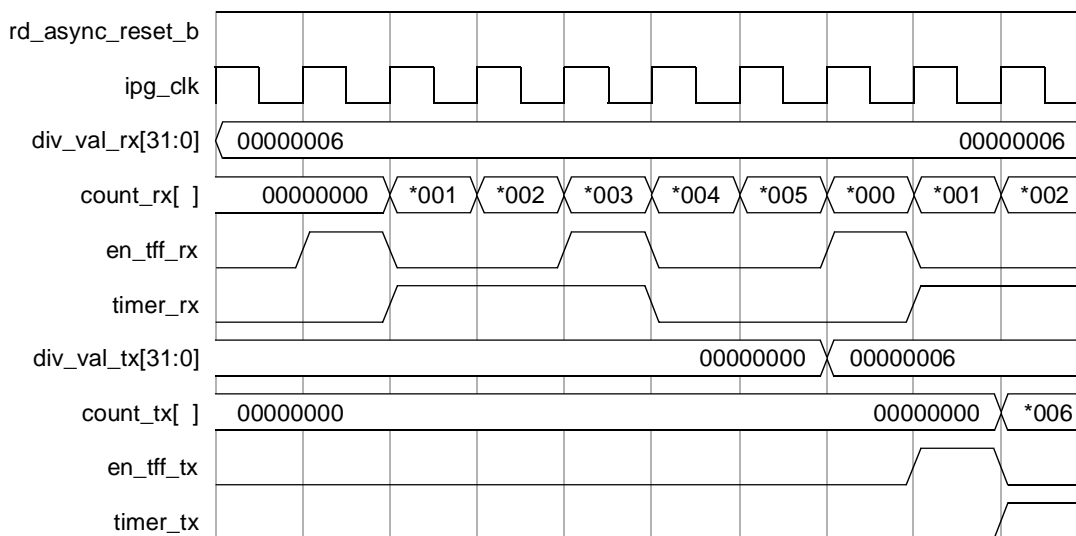


Figure 25-44. Division Value Even

- N = Odd: The en_tff_rx/en_tff_tx is driven high (logic 1) when respective counter value is 0 and N-1/2. Figure 25-45 below shows the case where division value for Rx and Tx serial clock is set to 5.

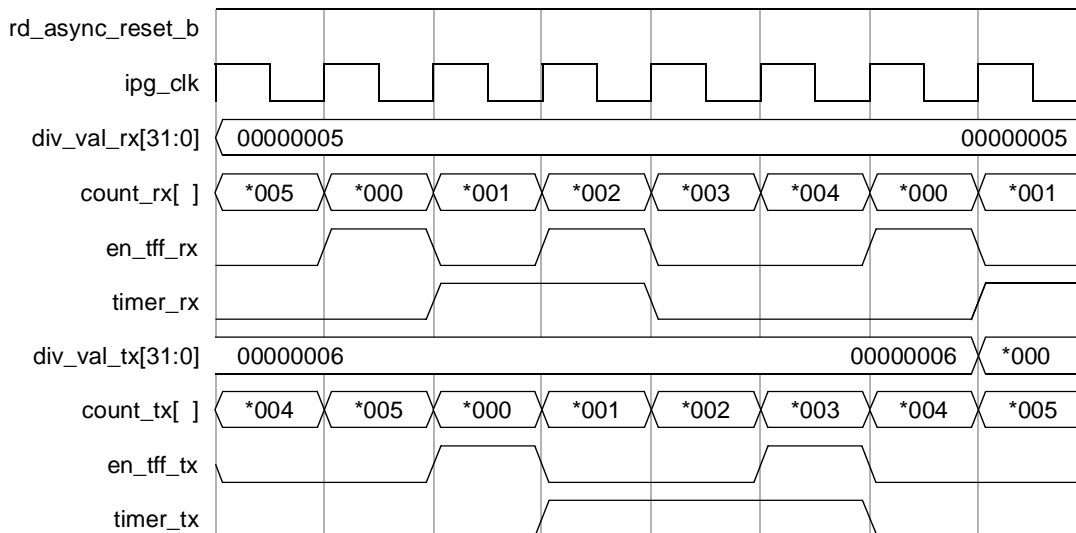


Figure 25-45. Division Value Odd

25.9.3 TDM Clock Memory Map

All the miscellaneous registers are defined in tdmclk_regs module. This section elaborates all the register definition.

Table 25-31. TDM Clock Memory Map

| Offset | Name | Access | Reset | Section/Page |
|--|---|--------|-------------|----------------|
| TDM Clock Control—Block Base Address 0x1_6180 | | | | |
| 0x000 | TDM Division Value Rx Serial Clock Register (TDMCLK_DIV_VAL_RX) | R/W | 0x0000_0000 | 25.9.3.1/25-52 |
| 0x004 | TDM Division Value for Tx Serial Clock Register (TDMCLK_DIV_VAL_TX) | R/W | 0x0000_0000 | 25.9.3.2/25-53 |

25.9.3.1 TDMCLK_DIV_VAL_RX Register

Figure 25-46 shows the bit fields for the TDMCLK_DIV_VAL_RX register.




Figure 25-46. TDMCLK_DIV_VAL_RX Register

Table 25-32 describes the bit fields for the TDMCLK_DIV_VAL_RX register.

Table 25-32. TDMCLK_DIV_VAL_RX Register

| Bits | Name | Description |
|------|--------|---|
| 0–31 | DIV_RX | Division value for Rx serial clock. 0 Rx serial clock off (low) non_0 Rx serial clock is enabled and timer_rx:csb_clk ratio is 1:tdmclk_div_val_rx (csb_clk is at higher frequency)1 |

25.9.3.2 TDMCLK_DIV_VAL_TX Register

Figure 25-47 shows the bit fields for the TDMCLK_DIV_VAL_TX register.




Figure 25-47. TDMCLK_DIV_VAL_TX Register

Table 25-33 describes the bit fields for the TDMCLK_DIV_VAL_TX register.

Table 25-33. TDMCLK_DIV_VAL_TX Register

| Bits | Name | Description |
|------|--------|---|
| 0–31 | DIV_TX | Division value for Tx serial clock. 0 Tx serial clock off (low) non_0 1 Tx serial clock is enabled and timer_tx:csb_clk ratio is 1:tdmclk_div_val_tx (csb_clk is at higher frequency) |

25.10 DMA Controller (DMAC)

The DMA (direct memory access) is a second-generation platform module capable of performing complex data transfers with minimal intervention from a host processor via n programmable channels. Intended for use as part of the Standard Product Platform (SPP), the hardware microarchitecture includes a DMA engine which performs source and destination address calculations, and the actual data movement operations, along with a local memory containing the transfer control descriptors (TCD) for the channels. This SRAM-based implementation is utilized to minimize the overall module size.

Figure 25-48 shows the DMA block diagram.




Figure 25-48. DMA Block Diagram

25.10.1 Overview

The DMA is a highly programmable data transfer engine, which has been optimized to minimize the required intervention from the host processor. It is intended for use in applications where the data size to be transferred is statically known, and is not defined within the data packet itself. The DMA hardware supports:

- Single design with two channels (Tx and Rx)
- Connections to the AMBA-AHB crossbar switch for bus mastering the data movement, IPI-Slave (IPS) slave bus for programming the module
 - Parameterized support for 32- and 64-bit AMBA-AHB datapath widths
- 32-byte transfer control descriptor per channel stored in local memory
- 32 bytes of data registers, used as temporary storage to support burst transfers
- Preliminary size estimation including local SRAM + BIST controller

Throughout this section, *n* is used to reference the channel number. Additionally, data sizes are defined as byte (8-bit), half word (16-bit), word (32-bit) and double word (64-bit).

25.10.1.1 Features

The DMA module supports the following features:

- All data movement via dual-address transfers: read from source, write to destination
 - Programmable source, destination addresses, transfer size, plus support for enhanced addressing modes
- Transfer control descriptor organized to support two-deep, nested transfer operations
 - An inner data transfer loop defined by a minor byte transfer count
 - An outer data transfer loop defined by a major iteration count
- Channel service request via one of three methods:
 - Explicit software initiation
 - Initiation via a channel-to-channel linking mechanism for continuous transfers
 - Independent channel linking at end of minor loop and/or major loop
 - Peripheral-paced hardware requests (one per channel)
 - For all three methods, one service request per execution of the minor loop is required
- Support for fixed-priority and round-robin channel arbitration
- Channel completion reported via optional interrupt requests
 - One interrupt per channel, optionally asserted at completion of major iteration count
 - Error terminations are optionally enabled per channel, and logically summed together to form a small number of error interrupt outputs
- Support for scatter/gather DMA processing
- Support for complex data structures
- Support to cancel transfers via software or hardware

25.10.2 DMAC Memory Map/Register Definition

The DMA programming model is partitioned into two sections, both mapped into the IPI slave space: the first region defines a number of registers providing control functions, while the second region corresponds to the local transfer control descriptor memory. Reading an unimplemented register bit or memory location will return the value of zero. Write the value of zero to unimplemented register bits. Any access to a reserved memory location will result in a bus error. Reserved memory locations are indicated in the memory map. For 16- and 32-channel implementations, reserved memory also includes the high order “H” registers containing channels 63–32 data (that is, DMAERQH, DMAEEIH, DMAINTH, DMAERRH).

Many of the control registers have a bit width that matches the number of channels implemented in the module, that is, 16-, 32-, or 64-bits in size. Registers associated with a 64-channel design are implemented as two 32-bit registers, and include an ‘H’ or ‘L’ suffix, signaling the high and low portions of the control function. The descriptions in this section define the 64-channel implementation. For 16- or 32-channel designs, the unused bits are not implemented: reads return zeros, and writes are ignored.

NOTE

Only two channels (channel 0 and channel 1) are available on the MPC8315E. References to and examples utilizing 64-bits, which refer to 64 channels, are not applicable and should be disregarded.

The DMA module does not include any logic which provides access control. Rather, this function is supported using the standard access control logic provided by the AIPS controller. [Table 25-34](#) is a 32-bit view of the DMA memory map.

Table 25-34. DMAC Register Summary

| Offset | Register | Access | Reset | Section/Page |
|-------------------------------------|---|--------|-------------|----------------------------------|
| Block Base Address: 0x2_C000 | | | | |
| 0x000 | DMACR—DMA Control Register | R/W | 0x0000_E400 | 25.10.2.1/25-57 |
| 0x004 | DMAES—DMA Error Status Register | RO | 0x0000_0000 | 25.10.3/25-59 |
| 0x008 | Reserved | — | — | — |
| 0x00C | DMAERQ—DMA enable request register | R/W | 0x0000_0000 | 25.10.3.1/25-62 |
| 0x010 | Reserved | — | — | — |
| 0x014 | DMAEEI—DMA enable error interrupt register | R/W | 0x0000_0000 | 25.10.3.2/25-63 |
| 0x018 | DMASERQ—DMA set enable request | R/W | 0x00 | 25.10.3.3/25-64 |
| 0x019 | DMACERQ—DMA clear enable request | R/W | 0x00 | 25.10.3.4/25-64 |
| 0x01A | DMASEEI—DMA Set Enable Error Interrupt | R/W | 0x00 | 25.10.3.5/25-65 |
| 0x01B | DMACEEI—DMA Clear Enable Error Interrupt | R/W | 0x00 | 25.10.3.6/25-65 |
| 0x01C | DMACINT—DMA Clear Interrupt Request | R/W | 0x00 | 25.10.3.7/25-66 |
| 0x01D | DMACERR—DMA Clear Error | R/W | 0x00 | 25.10.3.8/25-67 |
| 0x01E | DMASSRT—DMA Set START Bit | R/W | 0x00 | 25.10.3.9/25-67 |
| 0x01F | DMACDNE—DMA Clear DONE Status Bit | R/W | 0x00 | 25.10.3.10/25-68 |
| 0x020 | Reserved | — | — | — |
| 0x024 | DMAINT—DMA interrupt request register | w1c | 0x0000_0000 | 25.10.3.11/25-68 |
| 0x028 | Reserved | — | — | — |
| 0x02C | DMAERR—DMA error register | w1c | 0x0000_0000 | 25.10.3.12/25-69 |
| 0x030 | Reserved | — | — | — |
| 0x034 | DMAHRS—DMA hardware request status register | R/W | 0x000_0000 | 25.10.3.13/25-70 |
| 0x038 | DMAGPOR—DMA general purpose output register | R/W | 0x0000_0000 | 25.10.3.14/25-71 |
| 0x03C–0x0FC | Reserved | — | — | — |

Table 25-34. DMAC Register Summary (continued)

| Offset | Register | Access | Reset | Section/Page |
|-----------------|---|--------|-------|----------------------------------|
| 0x100– 0x13C | DCHPRI _n —DMA Channel <i>n</i> Priority Register | R/W | 0xnn | 25.10.3.15/25-71 |
| 0x140– 0xFFC | Reserved | — | — | — |

25.10.2.1 DMA Control Register (DMACR)

The 32-bit DMACR defines the basic operating configuration of the DMA.

The DMA arbitrates channel service requests in groups of 16 channels. The 64- and 32-channel configurations have 4 groups (3,2,1,0) and 2 groups (1,0), respectively; the 16-channel configuration has only 1 group (0). Group 3 contains channels 63–48, group 2 contains channels 47–32, group 1 contains channels 31–16, and group 0 contains channels 15–0.

Arbitration within a group can be configured to use either a fixed priority or a round robin. In fixed priority arbitration, the highest priority channel requesting service is selected to execute. The priorities are assigned by the channel priority registers (see [Section 25.10.3.15, “DMA Channel *n* Priority \(DCHPRI_n\), *n* = 0, ..., {15,31,63}”](#)). In round robin arbitration mode, the channel priorities are ignored and the channels within each group are cycled through without regard to priority.

The group priorities operate in a similar fashion. In group fixed priority arbitration mode, channel service requests in the highest priority group are executed first where priority level 3 is the highest and priority level 0 is the lowest. The group priorities are assigned in the GRPPri registers. All group priorities must have unique values prior to any channel service requests occur, otherwise a configuration error will be reported. Unused group priority registers, per configuration, are unimplemented in the DMACR. In group round robin mode, the group priorities are ignored and the groups are cycled through without regard to priority.

Minor loop offsets are address offset values added to the final source address (saddr) or destination address (daddr) upon minor loop completion. When minor loop offsets are enabled, the minor loop offset (mloff) is added to the final source address (saddr), or the final destination address (daddr), or both prior to the addresses being written back into the TCD. If the major loop is complete, the minor loop offset is ignored and the major loop address offsets (slast and dlast_sga) are used to compute the next saddr and daddr values.

When minor loop mapping is enabled (DMACR[EMLM] = 1), TCD word2 is redefined. A portion of TCD word2 is used to specify multiple fields: an source enable bit (smloe) to specify the minor loop offset should be applied to the source address (saddr) upon minor loop completion, an destination enable bit (dmloe) to specify the minor loop offset should be applied to the destination address (daddr) upon minor loop completion, and the sign extended minor loop offset value (mloff). The same offset value (mloff) is used for both source and destination minor loop offsets. When either minor loop offset is enabled (smloe set or dmloe set), the nbytes field is reduced to 8 bits. When both minor loop offsets are disabled (smloe cleared and dmloe cleared), the nbytes field becomes a 30-bit vector.

Table 25-35. DMA Control Register (DMACR) Field Descriptions (continued)

| Bits | Name | Description |
|------|------|---|
| 6 | CLM | Continuous link mode. 0 A minor loop channel link made to itself will go through channel arbitration before being activated again. 1 A minor loop channel link made to itself will not go through channel arbitration before being activated again. Upon minor loop completion, the channel will active again if that channel has a minor loop channel link enabled and the link channel is itself. This effectively applies the minor loop offsets and restarts the next minor loop. |
| 5 | HALT | Halt DMA operations. 0 Normal operation. 1 Stall the start of any new channels. Executing channels are allowed to complete. Channel execution will resume when the HALT bit is cleared. |
| 4 | HOE | Halt on error. 0 Normal operation. 1 Any error will cause the HALT bit to be set. Subsequently, all service requests will be ignored until the HALT bit is cleared. |
| 3 | ERGA | Enable round robin group arbitration. 0 Fixed priority arbitration is used for selection among the groups. 1 Round robin arbitration is used for selection among the groups. |
| 2 | ERCA | Enable round robin channel arbitration. 0 Fixed priority arbitration is used for channel selection within each group. 1 Round robin arbitration is used for channel selection within each group. |
| 1 | EDBG | Enable debug. 0 The assertion of the ipg_debug input is ignored. 1 The assertion of the ipg_debug input causes the DMA to stall the start of a new channel. Executing channels are allowed to complete. Channel execution will resume when either the ipg_debug input is negated or the EDBG bit is cleared. |
| 0 | EBW | Enable buffered writes. 0 The bufferable write signal (hprot[2]) is not asserted during AMBA AHB writes. 1 The bufferable write signal (hprot[2]) is asserted on all AMBA AHB writes except for the last write sequence write sequence. |

25.10.3 DMA Error Status (DMAES)

The DMAES register provides information concerning the last recorded channel error. Channel errors can be caused by a configuration error (an illegal setting in the transfer control descriptor or an illegal priority register setting in fixed arbitration mode) or an error termination to a bus master read or write cycle.

A configuration error is caused when the starting source or destination address, source or destination offsets, minor loop byte count and the transfer size represent an inconsistent state. The addresses and offsets must be aligned on 0-modulo-transfer_size boundaries, and the minor loop byte count must be a multiple of the source and destination transfer sizes. All source reads and destination writes must be configured to the natural boundary of the programmed transfer size respectively. In fixed arbitration mode, a configuration error is caused by any two channel priorities being equal within a group, or any group priority levels being equal among the groups. All channel priority levels within a group must be unique and all group priority levels among the groups must be unique when fixed arbitration mode is enabled. If a scatter/gather operation is enabled upon channel completion, a configuration error is reported if the

scatter/gather address (dlast_sga) is not aligned on a 32 byte boundary. If minor loop channel linking is enabled upon channel completion, a configuration error is reported when the link is attempted if the TCD.citer.e_link bit does not equal the TCD.biter.e_link bit. All configuration error conditions except scatter/gather and minor loop link error are reported as the channel is activated and assert an error interrupt request, if enabled. A scatter/gather configuration error is reported when the scatter/gather operation begins at major loop completion when properly enabled. A minor loop channel link configuration error is reported when the link operation is serviced at minor loop completion.

If a system bus read or write is terminated with an error, the data transfer is stopped and the appropriate bus error flag set. In this case, the state of the channel's transfer control descriptor is updated by the dma_engine with the current source address, destination address and current iteration count at the point of the fault. When a system bus error occurs, the channel is terminated after the read or write transaction which is already pipelined after errant access, has completed. If a bus error occurs on the last read prior to beginning the write sequence, the write will execute using the data captured during the bus error. If a bus error occurs on the last write prior to switching to the next read sequence, the read sequence will execute before the channel is terminated due to the destination bus error.

A transfer may be cancelled by software via the DMACR[CX] bit or hardware via the dma_cancel_xfer input signal. When a cancel transfer request is recognized, the dma_engine stops processing the channel. The current read-write sequence is allowed to finish. If the cancel occurs on the last read-write sequence of a major or minor loop, the cancel request is discarded and the channel retires normally.

The error cancel transfer is the same as a cancel transfer except the DMAES register is updated with the cancelled channel number and error cancel bit is set. The TCD of a cancelled channel has the source address and destination address of the last transfer saved in the TCD. It is the responsibility of the user to initialize the TCD again should the channel need to be restarted because the aforementioned fields have been modified by the dma_engine and no longer represent the original parameters. When a transfer is cancelled via the error cancel transfer mechanism (setting the DMACR[ECX] or asserting the dma_err_cancel_xfer input), the channel number is loaded into the ERRCHN field and the ECX and VLD bits are set in the DMAES register. In addition, an error interrupt may be generated if enabled. See [Section 25.10.3.12, “DMA Error Register \(DMAERR\),”](#) for error interrupt details.

The occurrence of any type of error causes the dma_engine to immediately stop, and the appropriate channel bit in the DMA error register to be asserted. At the same time, the details of the error condition are loaded into the DMAES register. The major loop complete indicators, setting the transfer control descriptor done flag and the possible assertion of an interrupt request, are not affected when an error is detected. See [Table 25-36](#) for the DMAES definition.

Table 25-36. DMAES Field Descriptions (continued)

| Bits | Name | Value |
|------|------|--|
| 4 | DOE | Destination offset error. 0 No destination offset configuration error. 1 The last recorded error was a configuration error detected in the TCD.doff field. TCD.doff is inconsistent with TCD.dsize. |
| 3 | NCE | Nbytes/citer configuration error. 0 No nbytes/citer configuration error. 1 The last recorded error was a configuration error detected in the TCD.nbytes or TCD.citer fields. TCD.nbytes is not a multiple of TCD.ssize and TCD.dsize, or TCD.citer is equal to zero, or TCD.citer.e_link is not equal to TCD.biter.e_link. |
| 2 | SGE | Scatter/gather configuration error. 0 No scatter/gather configuration error. 1 The last recorded error was a configuration error detected in the TCD.dlast_sga field. This field is checked at the beginning of a scatter/gather operation after major loop completion if TCD.e_sg is enabled. TCD.dlast_sga is not on a 32 byte boundary. |
| 1 | SBE | Source bus error. 0 No source bus error. 1 The last recorded error was a bus error on a source read. |
| 0 | DBE | Destination bus error. 0 No destination bus error. 1 The last recorded error was a bus error on a destination write. |

25.10.3.1 DMA Enable Request Register (DMAERQ)

DMAERQ provides a bit map for the implemented channels {16,32} to enable the request signal for each channel. The state of any given channel enable is directly affected by writes to this register; it is also affected by writes to DMASERQ and DMACERQ.

DMASERQ and DMACERQ are provided so that the request enable for a single channel can easily be modified without the need to perform a read-modify-write sequence to DMAERQ. Both the DMA request input signal and this enable request flag must be asserted before a channel's hardware service request is accepted. The state of the DMA enable request flag does not affect a channel service request made explicitly through software or a linked channel request.

As a given channel completes the processing of its major iteration count, there is a flag in the transfer control descriptor that may affect the ending state of the DMAERQ bit for that channel. If the TCD.d_req bit is set, then the corresponding DMAERQ bit is cleared, disabling the DMA request; else if the d_req bit is cleared, the state of the DMAERQ bit is unaffected.

Offset 0x00C

Access: Read/Write



Figure 25-51. DMA Enable Request Register (DMAERQ)

Table 25-37 provides the DMAERQ definition.

Table 25-37. DMAERQ Field Descriptions

| Bits | Name | Description |
|------|---------|---|
| 31–0 | ERQ n | Enable DMA request n . 0 The DMA request signal for channel n is disabled. 1 The DMA request signal for channel n is enabled. |

25.10.3.2 DMA Enable Error Interrupt Register (DMAEEI)

DMAEEI provides a bit map for the implemented channels {16,32} to enable the error interrupt signal for each channel. The state of any given channel's error interrupt enable is directly affected by writes to this register; it is also affected by writes to DMASEEI and DMACEEI. DMASEEI and DMACEEI are provided so that the error interrupt enable for a single channel can easily be modified without the need to perform a read-modify-write sequence to DMAEEI.

Both the DMA error indicator and this error interrupt enable flag must be asserted before an error interrupt request for a given channel is asserted.

Offset 0x014

Access: Read/Write



Figure 25-52. DMA Enable Error Interrupt Register (DMAEEI)

Table 25-38 provides the DMAEEI definition.

Table 25-38. DMAEEI Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 31-0 | EEIn | Enable error interrupt <i>n</i> . 0 The error signal for channel <i>n</i> does not generate an error interrupt. 1 The assertion of the error signal for channel <i>n</i> generate an error interrupt request. |

25.10.3.3 DMA Set Enable Request (DMASERQ)

DMASERQ provides a simple memory-mapped mechanism to set a given bit in the DMAERQ register to enable the DMA request for a given channel. The data value on a register write causes the corresponding bit in DMAERQ to be set. A data value of 64 to 127 (regardless of the number of implemented channels) provides a global set function, forcing the entire contents of DMAERQ to be asserted. If bit 7 is set, the command is ignored. This allows multiple-byte registers to be written as a 32-bit word. Reads of this register return all zeros. See Table 25-39 for the DMASERQ definition.




Figure 25-53. DMA Set Enable Request Register

Table 25-39. DMASERQ Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7 | NOP | No operation. 0 Normal operation. 1 No operation, ignore bits 6-0. |
| 6-0 | SERQ | Set enable request. 0-63 Set the corresponding bit in DMAERQ. 64-127 Set all bits in DMAERQ. |

25.10.3.4 DMA Clear Enable Request (DMACERQ)

DMACERQ provides a simple memory-mapped mechanism to clear a given bit in DMAERQ to disable the DMA request for a given channel. The data value on a register write causes the corresponding bit in DMAERQ to be cleared. A data value of 64 to 127 (regardless of the number of implemented channels) provides a global clear function, forcing the entire contents of DMAERQ to be zeroed, disabling all DMA request inputs. If bit 7 is set, the command is ignored. This allows multiple-byte registers to be written as a 32-bit word. Reads of this register return all zeros. See Table 25-40 for the DMACERQ definition.


Figure 25-54. DMA Clear Enable Request Register
Table 25-40. DMACERQ Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7 | NOP | No operation. 0 Normal operation. 1 No operation, ignore bits 6–0 |
| 6–0 | CERQ | Clear enable request. 0–63 Clear corresponding bit in DMAERQ 64–127 Clear all bits in DMAERQ |

25.10.3.5 DMA Set Enable Error Interrupt (DMASEEI)

DMASEEI provides a simple memory-mapped mechanism to set a given bit in the DMAEEI register to enable the error interrupt for a given channel. The data value on a register write causes the corresponding bit in DMAEEI to be set. A data value of 64 to 127 (regardless of the number of implemented channels) provides a global set function, forcing the entire contents of DMAEEI to be asserted. If bit 7 is set, the command is ignored. This allows multiple-byte registers to be written as a 32-bit word. Reads of this register return all zeros. See [Table 25-41](#) for the DMASEEI definition.


Figure 25-55. DMA Set Enable Error Interrupt Register
Table 25-41. DMASEEI Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7 | NOP | No operation. 0 Normal operation. 1 No operation, ignore bits 6–0. |
| 6–0 | SEEI | Set enable error interrupt. 0–63 Set the corresponding bit in DMAEEI. 64–127 Set all bits in DMAEEI. |

25.10.3.6 DMA Clear Enable Error Interrupt (DMACEEI)

The DMACEEI register provides a simple memory-mapped mechanism to clear a given bit in the DMAEEI register to disable the error interrupt for a given channel. The data value on a register write

causes the corresponding bit in DMAEEI to be cleared. A data value of 64 to 127 (regardless of the number of implemented channels) provides a global clear function, forcing the entire contents of DMAEEI to be zeroed, disabling all DMA request inputs. If bit 7 is set, the command is ignored. This allows multiple-byte registers to be written as a 32-bit word. Reads of this register return all zeros. See [Table 25-42](#) for the DMACEEI definition.




Figure 25-56. DMA Clear Enable Error Interrupt Register

Table 25-42. DMACEEI Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 7 | NOP | No operation. 0 Normal operation. 1 No operation, ignore bits 6–0. |
| 6–0 | CEEI | Clear enable error interrupt. 0–63 Clear corresponding bit in DMAEEI. 64–127 Clear all bits in DMAEEI. |

25.10.3.7 DMA Clear Interrupt Request (DMACINT)

DMACINT provides a simple memory-mapped mechanism to clear a given bit in the DMAINT register to disable the interrupt request for a given channel. The given value on a register write causes the corresponding bit in DMAINT to be cleared. A data value of 64 to 127 (regardless of the number of implemented channels) provides a global clear function, forcing the entire contents of DMAINT to be zeroed, disabling all DMA interrupt requests. If bit 7 is set, the command is ignored. This allows multiple-byte registers to be written as a 32-bit word. Reads of this register return all zeros. See [Table 25-43](#) for the DMACINT definition.




Figure 25-57. DMA Clear Interrupt Request Register

Table 25-43. DMACINT Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 7 | NOP | No operation. 0 Normal operation. 1 No operation, ignore bits 6–0. |
| 6–0 | CINT | Clear interrupt request. 0–63 Clear the corresponding bit in DMAINT. 64–127 Clear all bits in DMAINT. |

25.10.3.8 DMA Clear Error (DMACERR)

DMACEER provides a simple memory-mapped mechanism to clear a given bit in the DMAERR register to disable the error condition flag for a given channel. The given value on a register write causes the corresponding bit in DMAERR to be cleared. A data value of 64 to 127 (regardless of the number of implemented channels) provides a global clear function, forcing the entire contents of DMAERR to be zeroed, clearing all channel error indicators. If bit 7 is set, the command is ignored. This allows multiple-byte registers to be written as a 32-bit word. Reads of this register return all zeros. See [Table 25-44](#) for the DMACERR definition.


Figure 25-58. DMA Clear Error Register
Table 25-44. DMACERR Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 7 | NOP | No operation. 0 Normal operation. 1 No operation, ignore bits 6–0. |
| 6–0 | CERR | Clear error indicator. 0–63 Clear corresponding bit in DMAERR. 64–127 Clear all bits in DMAERR. |

25.10.3.9 DMA Set START Bit (DMASSRT)

DMASSRT provides a simple memory-mapped mechanism to set the START bit in the TCD of the given channel. The data value on a register write causes the START bit in the corresponding Transfer Control Descriptor to be set. A data value of 64 to 127 (regardless of the number of implemented channels) provides a global set function, forcing all START bits to be set. If bit 7 is set, the command is ignored. This allows multiple byte registers to be written as a 32-bitword. Reads of this register return all zeros. See [Table 25-45](#) for the TCD START bit definition.




Figure 25-59. DMA Set START Bit Register

Table 25-45. DMASSRT Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 7 | NOP | No operation. 0 Normal operation. 1 No operation, ignore bits 6–0. |
| 6–0 | SSRT | Set START bit (channel service request). 0–63 Set the corresponding channel's TCD.start. 64–127 Set all TCD.start bits. |

25.10.3.10 DMA Clear DONE Status (DMACDNE)

DMACDNE provides a simple memory-mapped mechanism to clear the DONE bit in the TCD of the given channel. The data value on a register write causes the DONE bit in the corresponding transfer control descriptor to be cleared. A data value of 64 to 127 (regardless of the number of implemented channels) provides a global clear function, forcing all DONE bits to be cleared. If bit 7 is set, the command is ignored. This allows multiple byte registers to be written as a 32-bit word. Reads of this register return all zeros. See [Table 25-46](#) for the TCD DONE bit definition.



Figure 25-60. DMA Clear DONE Status Register

Table 25-46. DMACDNE Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 7 | NOP | No operation 0 Normal operation. 1 No operation, ignore bits 6–0. |
| 6–0 | CDNE | Clear DONE status bit. 0–63 Clear the corresponding channel's DONE bit. 64–127Clear all TCD DONE bits. |

25.10.3.11 DMA Interrupt Request Register (DMAINT)

DMAINT provide a bit map for the implemented channels { 16,32}, signaling the presence of an interrupt request for each channel. The dma_engine signals the occurrence of a programmed interrupt upon completion of a data transfer as defined in the transfer control descriptor by setting the appropriate bit in

this register. The outputs of this register are directly routed to the platform's interrupt controller. During the execution of the interrupt service routine associated with any given channel, it is software's responsibility to clear the appropriate bit, negating the interrupt request. Typically, a write to the DMACINT register in the interrupt service routine is used for this purpose.

The state of any given channel's interrupt request is directly affected by writes to this register; it is also affected by writes to the DMACINT register. On writes to DMAINT, a one in any bit position clears the corresponding channel's interrupt request. A zero in any bit position has no effect on the corresponding channel's current interrupt status. The DMAINT register is provided so the interrupt request for a single channel can easily be cleared without the need to perform a read-modify-write sequence to DMAINT.

Offset 0x024 Access: w1c

| | | | | | | | | | | | | | | | | |
|-------|-----------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| | 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 |
| R | INT31 | INT30 | INT29 | INT28 | INT27 | INT26 | INT25 | INT24 | INT23 | INT22 | INT21 | INT20 | INT19 | INT18 | INT17 | INT16 |
| W | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c |
| Reset | All zeros | | | | | | | | | | | | | | | |

| | | | | | | | | | | | | | | | | |
|-------|-----------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| R | INT15 | INT14 | INT13 | INT12 | INT11 | INT10 | INT09 | INT08 | INT07 | INT06 | INT05 | INT04 | INT03 | INT02 | INT01 | INT00 |
| W | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c |
| Reset | All zeros | | | | | | | | | | | | | | | |

Figure 25-61. DMA Interrupt Request Register Low (DMAINT)

See [Table 25-47](#) for the DMAINT definition.

Table 25-47. DMAINT Field Descriptions

| Bits | Name | Description |
|------|---------|---|
| 31–0 | INT n | DMA interrupt request n (write one to clear) 0 The interrupt request for channel n is cleared. 1 The interrupt request for channel n is active. |

25.10.3.12 DMA Error Register (DMAERR)

DMAERRL provides a bit map for the implemented channels {16,32}, signaling the presence of an error for each channel. The dma_engine signals the occurrence of a error condition by setting the appropriate bit in this register. The outputs of this register are enabled by the contents of the DMAEEI register, then logically summed across groups of 16 and 32 channels to form several group error interrupt requests, which are then routed to the platform's interrupt controller. During execution of the interrupt service routine associated with any DMA errors, it is software's responsibility to clear the appropriate bit, negating the error interrupt request; typically, a write to the DMACERR register in the interrupt service routine is used for this purpose. Recall that the normal DMA channel completion indicators, setting the transfer control descriptor done flag and the possible assertion of an interrupt request, are not affected when an error is detected.

The contents of this register can also be polled, and a non-zero value indicates the presence of a channel error regardless of the state of the DMAEEI register. The state of any given channel's error indicators is affected by writes to this register; it is also affected by writes to DMACERR. On writes to DMAERR, a

one in any bit position clears the corresponding channel’s error status; a zero in any bit position has no effect. DMACERR is provided so the error indicator for a single channel can easily be cleared.




Figure 25-62. DMA Error Register (DMAERR)

See [Table 25-48](#) for the DMAERR definition.

Table 25-48. DMAERR Field Descriptions

| Bits | Name | Description |
|------|---------|--|
| 31–0 | INT n | DMA error n (write one to clear) 0 An error in channel n has not occurred. 1 An error in channel n has occurred. |

25.10.3.13 DMA Hardware Request Status Register (DMAHRS)

DMAHRSL provides a bit map for the implemented channels {16,32} to show the current hardware request status for each channel.




Figure 25-63. DMA Hardware Request Status Register (DMAHRS)

See [Table 25-49](#) for the DMAHRS definition.

Table 25-49. DMAHRS Field Descriptions

| Bits | Name | Description |
|------|---------|--|
| 31–0 | INT n | DMA hardware request status n . 0 A hardware service request for channel n is not present. Note: A hardware service request for channel n is present. |

25.10.3.14 DMA General Purpose Output Register (DMAGPOR)

The optional DMAGPOR register provides a general purpose register in the programmer’s model that outputs the register contents. The DMAGPOR performs no functions within the DMA2. This general purpose register is enabled when SPP_DMA2_ENABLE_GPOR is defined. This register may be used by the SoC integrator to define and display configuration information. The contents of this register are exported out of the DMA2.




Figure 25-64. DMA Hardware Request Status Register—High

See [Table 25-56](#) for the DMAGPOR definition.

Table 25-50. DMAGPOR Field Descriptions

| Bits | Name | Description |
|-------|--------------------|--|
| 31–16 | — | Reserved |
| 15–14 | priority_ctrl_high | Priority control high. Indicates the bump up priority value in case of non-sequential transactions. |
| 13–12 | priority_ctrl | Priority control. Normal priority value. |
| 11–7 | — | Reserved |
| 6 | snoop enable | Snoop enable. 0 Transaction is non-cacheable 1 Transaction is cacheable |
| 5–3 | — | Reserved |
| 2 | rd_safe_enable | Read safe enable. 0 Reading extra bytes from the memory may have side effects. 1 Only indicated amount of data beats should be read. |
| 1–0 | — | Reserved |

25.10.3.15 DMA Channel *n* Priority (DCHPRI_{*n*}), *n* = 0,..., {15,31,63}

When the fixed-priority channel arbitration mode is enabled (DMACR[ERCA] = 0), the contents of these registers define the unique priorities associated with each channel within a group. The channel priorities are evaluated by numeric value, that is, 0 is the lowest priority, 1 is the next higher priority, then 2, 3, and so on. Software must program the channel priorities with unique values, otherwise a configuration error will be reported. The range of the priority value is limited to the values of 0 through 15. When read, the GRPPRI bits of the DCHPRI register reflect the current priority level of the group of channels in which

the corresponding channel resides. GRPPRI bits are not affected by writes to the DCHPRI registers. The group priority is assigned in the DMACR. See [Figure 25-49](#) and [Table 25-35](#) for the DMACR definition.

Channel preemption is enabled on a per channel basis by setting the ECP bit in the DCHPRI register. Channel preemption allows the executing channel's data transfers to be temporarily suspended in favor of starting a higher priority channel. Once the preempting channel has completed all of its minor loop data transfers, the preempted channel is restored and resumes execution. After the restored channel completes one read/write sequence, it is again eligible for preemption. If any higher priority channel is requesting service, the restored channel will be suspended and the higher priority channel will be serviced. Nested preemption (attempting to preempt a preempting channel) is not supported. Once a preempting channel begins execution, it cannot be preempted. Preemption is only available when fixed arbitration is selected for both group and channel arbitration modes.

A channel's ability to preempt another channel can be disabled by setting the DPA bit in the DCHPRI register. When a channel's preempt ability is disabled, that channel cannot suspend a lower priority channel's data transfer; regardless of the lower priority channel's ECP setting. This allows for a pool of low priority, large data moving channels to be defined. These low priority channels can be configured to not preempt each other, thus preventing a low priority channel from consuming the preempt slot normally available a true, high priority channel. See [Table 25-51](#) for the DCHPRI definition.




Figure 25-65. DMA Clear DONE Status Register

¹ n = defaults to channel number n after reset

Table 25-51. DCHPRI n Field Descriptions

| Bits | Name | Description |
|------|--------|--|
| 7 | ECP | Enable channel preemption. 0 Channel n cannot be suspended by a higher priority channel's service request. 1 Channel n can be temporarily suspended by the service request of a higher priority channel. |
| 6 | DPA | Disable preempt ability. 0 Channel n can suspend a lower priority channel. 1 Channel n cannot suspend any channel, regardless of channel priority. |
| 5–4 | GRPPRI | Channel n current group priority. Group priority assigned to this channel group when fixed-priority arbitration is enabled. These two bits are read only; writes are ignored. |
| 3–0 | CHPRI | Channel n arbitration priority. Channel priority when fixed-priority arbitration is enabled. |

25.10.3.16 Transfer Control Descriptor (TCD)

Each channel requires a 32-byte transfer control descriptor for defining the desired data movement operation. The channel descriptors are stored in the local memory in sequential order: channel 0, channel

1, ..., channel [n-1]. The definitions of the TCD are presented as eight 32-bit values. [Table 25-52](#) is a 32-bit view of the basic TCD structure.

Table 25-52. TCD 32-Bit Memory Structure

| DMA Offset | Word | TCD Field |
|----------------|------|---|
| 0x1000 + 0x000 | 0 | Source Address (saddr) |
| 0x1000 + 0x004 | 1 | Transfer Attributes (smod, ssize, dmod, dsize) Signed Source Address Offset (soff) |
| 0x1000 + 0x008 | 2 | Signed Minor Loop Offset (smloe, dmloe, mloff) Inner Minor Byte Count (nbytes) |
| 0x1000 + 0x00C | 3 | Last Source Address Adjustment (slast) |
| 0x1000 + 0x010 | 4 | Destination Address (daddr) |
| 0x1000 + 0x014 | 5 | Current Major Iteration Count (citer) Signed Destination Address Offset (doff) |
| 0x1000 + 0x018 | 6 | Last Destination Address Adjustment/Scatter Gather Address (dlast_sga) |
| 0x1000 + 0x01C | 7 | Beginning Major Iteration Count (biter) Channel Control/Status (bwc, major.linkch, done, active, major.e_link, e_sg, d_req, int_half, int_maj, start) |

[Figure 25-66](#) shows the TCD word 0 fields.




Figure 25-66. TCD Word 0 (TCD.saddr) Field

[Table 25-53](#) describes the TCD word 0.

Table 25-53. TCD Word 0 (TCD.saddr) Field Description

| Bits | Name | Description |
|------|-------|---|
| 31-0 | saddr | Source address. Memory address pointing to the source data. |

[Figure 25-67](#) shows the TCD word 1 fields.




Figure 25-67. TCD Word 1 (TCDn.{soff, smod, ssize, dmod, dsize}) Fields

Table 25-54 describes the TCD word 1 fields.

Table 25-54. TCD Word 1 (TCD.{smod, ssize, dmod, dsize, soff}) Field Descriptions

| Bits | Name | Description |
|-------|-------|--|
| 31–27 | smod | Source address modulo. 0 Source address modulo feature is disabled. Other The value defines a specific address bit which is selected to be either the value after saddr + soff calculation is performed or the original register value. This feature provides the ability to easily implement a circular data queue. For data queues requiring power-of-2 'size' bytes, the queue should be based at a 0-modulo-size address and the smod field set to the appropriate value to freeze the upper address bits. The bit select is defined as ((1 << smod[4:0]) - 1) where a resulting 1 in a bit location selects the next state address for the corresponding address bit location and a 0 selects the original register value for the corresponding address bit location. For this application, the soff is typically set to the transfer size to implement post-increment addressing with the smod function constraining the addresses to a 0-modulo-size range. |
| 26–24 | ssize | Source data transfer size. 000 8-bit 001 16-bit 010 32-bit 011 64-bit 100 16-byte (32-bit AHB bus, WRAP4 burst) 100 Reserved (64-bit AHB bus, reserved) 101 32-byte (if supported by the platform) 110 Reserved 111 Reserved The attempted specification of a 64-bit source size in a 32-bit AMBA AHB bus implementation produces a configuration error. Likewise, the attempted specification of a 16-byte source size in a 64-bit AMBA AHB bus implementation generates a configuration error. The attempted specification of a 32-byte burst on platforms that do not support such a transfer type will result in a configuration error. |
| 23–19 | dmod | Destination address modulo. See the smod definition. |
| 18–16 | dsize | Destination data transfer size. See the ssize definition. |
| 15–0 | soff | Source address signed offset. Sign-extended offset applied to the current source address to form the next-state value as each source read is completed. |

When minor loop mapping (DMACR[EMLM] = 1) is enabled, TCD word 2 is redefined as four fields: a source minor loop offset enable, a destination minor loop offset enable, a minor loop offset field and a nbytes field. Figure 25-68 shows the TCD word 2 fields.




Figure 25-68. TCD Word 2 (TCD.{smloe, dmloe, nbytes}) Field

Table 25-55 describes the word 2 fields.

Table 25-55. TCD Word 2 (TCD.{smloe, dmloe, nbytes}) Description

| Bits | Name | Description |
|-------------------------------|-------------------------------|---|
| 31 | smloe | Source minor loop offset enable. This flag selects whether the minor loop offset is applied to the source address upon minor loop completion. 0 The minor loop offset is not applied to the saddr. 1 The minor loop offset is applied to the saddr. |
| 30 | dmloe | Destination minor loop offset enable. This flag selects whether the minor loop offset is applied to the destination address upon minor loop completion. 0 The minor loop offset is not applied to the daddr. 1 The minor loop offset is applied to the daddr. |
| nbytes [29:10] or mloff[19:0] | nbytes [29:10] or mloff[19:0] | Inner minor byte transfer count or minor loop offset. If both smloe and dmloe are cleared, this field is part of the byte transfer count. If either smloe or dmloe are set, this field represents a sign-extended offset applied to the source or destination address to form the next-state value after the minor loop is completed. |
| nbytes[9:0] | nbytes[9:0] | Inner minor byte transfer count. Number of bytes to be transferred in each service request of the channel. As a channel is activated, the contents of the appropriate TCD is loaded into the dma_engine, and the appropriate reads and writes performed until the complete byte transfer count has been transferred. This is an indivisible operation and cannot be stalled or halted. Once the minor count is exhausted, the current values of the saddr and daddr are written back into the local memory, the major iteration count is decremented and restored to the local memory. If the major iteration count is completed, additional processing is performed. This field is extended to 30 bits when both smloe and dmloe are cleared (disabled). |

Figure 25-69 shows the TCD word 3 fields.




Figure 25-69. TCD Word 3 (TCD.slast) Field

Table 25-56 describes the word 3 fields.

Table 25-56. TCD Word 3 (TCD.slast) Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 31-0 | slast | Last source address adjustment. Adjustment value added to the source address at the completion of the outer major iteration count. This value can be applied to 'restore' the source address to the initial value, or adjust the address to reference the next data structure. |

Figure 25-70 shows the TCD word 4 fields.

Time Division Multiplexing (TDM) Interface




Figure 25-70. TCD Word 4 (TCDn.daddr) Field

Table 25-57 describes the word 4 fields.

Table 25-57. TCD Word 4 (TCD.daddr) Field Description

| Bits | Name | Description |
|------|-------|---|
| 31–0 | daddr | Destination address. Memory address pointing to the destination data. |

Figure 25-71 shows the TCD word 5 fields.




Figure 25-71. TCD Word 5 (TCD.{citer, doff}) Fields

Table 25-58 describes the word 5 fields.

Table 25-58. TCD Word 5 (TCD.{citer, doff}) Field Descriptions

| Bits | Name | Description |
|----------------------------------|----------------------------------|--|
| 31 | citer.e_link | Enable channel-to-channel linking on minor loop complete. As the channel completes the inner minor loop, this flag enables the linking to another channel, defined by citer.linkch[5:0]. The link target channel initiates a channel service request via an internal mechanism that sets the TCD.start bit of the specified channel. If channel linking is disabled, the citer value is extended to 15 bits in place of a link channel number. If the major loop is exhausted, this link mechanism is suppressed in favor of the major.e_link channel linking. This bit must be equal to the biter.e_link bit otherwise a configuration error will be reported. 0 The channel-to-channel linking is disabled. 1 The channel-to-channel linking is enabled. |
| citer[14:9] or citer.linkch[5:0] | citer[14:9] or citer.linkch[5:0] | Current major iteration count or link channel number. If (TCD.citer.e_link = 0) then no channel-to-channel linking (or chaining) is performed after the inner minor loop is exhausted. TCD word 5, bits [30:25] are used to form a 15 bit citer field. Otherwise, after the minor loop is exhausted, the dma_engine initiates a channel service request at the channel defined by citer.linkch[5:0] by setting that channel's TCD.start bit. The value contained in citer.linkch[5:0] must not exceed the number of implemented channels. |

Table 25-58. TCD Word 5 (TCD.{citer, doff} Field Descriptions (continued))

| Bits | Name | Description |
|------------|------------|--|
| citer[8:0] | citer[8:0] | Current major iteration count. This 9 or 15-bit count represents the current major loop count for the channel. It is decremented each time the minor loop is completed and updated in the transfer control descriptor memory. Once the major iteration count is exhausted, the channel performs a number of operations (for example, final source and destination address calculations), optionally generating an interrupt to signal channel completion before reloading the citer field from the beginning iteration count (biter) field. When the citer field is initially loaded by software, it must be set to the same value as that contained in the biter field. If the channel is configured to execute a single service request, the initial values of biter and citer should be 0x0001. |
| doff[15:0] | doff[15:0] | Destination address signed offset. Sign-extended offset applied to the current destination address to form the next-state value as each destination write is completed. |

Figure 25-72 shows the TCD word 6 fields.



Figure 25-72. TCD Word 6 (TCD.dlast_sga) Field

Table 25-59 describes the word 6 fields.

Table 25-59. TCD Word 6 (TCD.dlast_sga) Field Descriptions

| Bits | Name | Description |
|------|-----------|--|
| 31-0 | dlast_sga | Last destination address adjustment or the memory address for the next transfer control descriptor to be loaded into this channel (scatter/gather). If (TCD.e_sg = 0), then Adjustment value added to the destination address at the completion of the outer major iteration count. This value can be applied to 'restore' the destination address to the initial value, or adjust the address to reference the next data structure. else, This address points to the beginning of a 0-modulo-32 region containing the next transfer control descriptor to be loaded into this channel. This channel reload is performed as the major iteration count completes. The scatter/gather address must be 0-modulo-32, else a configuration error is reported. |

Figure 25-73 shows the TCD word 7 fields.




Figure 25-73. TCD Word 7 (TCD.{biter, control/status}) Fields

Table 25-60 describes the word 7 fields.

Table 25-60. TCD Word 7 (TCD.{biter, control/status}) Field Descriptions

| Bits | Name | Description |
|-------|-----------------------------------|--|
| 31 | biter.e_link | Enable channel-to-channel linking on minor loop complete. This is the initial value copied into the citer.e_link field when the major loop is completed. The citer.e_link field controls channel linking during channel execution. This bit must be equal to the citer.e_link bit otherwise a configuration error will be reported. 0 The channel-to-channel linking is disabled. 1 The channel-to-channel linking is enabled. |
| 30–25 | biter[14:9] or biter.linkch [5:0] | Beginning major iteration count or beginning link channel number. This is the initial value copied into the citer field or citer.linkch field when the major loop is completed. The citer fields control the iteration count and linking during channel execution. If (TCD.biter.e_link = 0) then No channel-to-channel linking (or chaining) is performed after the inner minor loop is exhausted. TCD word 5, bits [30:25] are used to form a 15-bit biter field. Otherwise, after the minor loop is exhausted, the dma_engine initiates a channel service request at the channel defined by biter.linkch[5:0] by setting that channel's TCD.start bit. The value contained in biter.linkch[5:0] must not exceed the number of implemented channels. |
| 24–16 | biter[8:0] | Beginning major iteration count. This is the initial value copied into the citer field or citer.linkch field when the major loop is completed. The citer fields controls the iteration count and linking during channel execution. This 9 or 15-bit counter presents the beginning major loop count for the channel. As the major iteration count is exhausted, the contents of the entire 16 bit biter entry is reloaded into the 16 bit citer entry. When the biter field is initially loaded by software, it must be set to the same value as that contained in the citer field. If the channel is configured to execute a single service request, the initial values of biter and citer should be 0x0001. |
| 15–14 | bwc[1:0] | Bandwidth control. This two-bit field provides a mechanism to effectively throttle the amount of bus bandwidth consumed by the DMA. In general, as the DMA processes the inner minor loop, it continuously generates read/write, read/write, ..., sequences until the minor count is exhausted. This field forces the DMA to stall after the completion of each read/write access to control the bus request bandwidth seen by the platform's cross-bar arbitration switch. To minimize start-up latency, bandwidth control stalls are suppressed for the first two AHB bus cycles and after the last write of each minor loop. The dynamic priority elevation setting elevates the priority of the DMA as seen by the cross-bar arbitration switch for the executing channel. Dynamic priority elevation is suppressed during the first two AHB bus cycles. 00 No dma_engine stalls 01 dynamic priority elevation ten dma_engine stalls for four cycles after each R/W 11 dma_engine stalls for eight cycles after each R/W |

Table 25-60. TCD Word 7 (TCD.{biter, control/status}) Field Descriptions (continued)

| Bits | Name | Description |
|------|--------------------|--|
| 13–8 | major.linkch [5:0] | Link channel number. If (TCD.major.e_link = 0) then No channel-to-channel linking (or chaining) is performed after the outer major loop counter is exhausted. Otherwise, after the major loop counter is exhausted, the dma_engine initiates a channel service request at the channel defined by major.linkch[5:0] by setting that channel's TCD.start bit. The value contained in major.linkch[5:0] must not exceed the number of implemented channels. |
| 7 | done | Channel done. This flag indicates the DMA has completed the outer major loop. It is set by the dma_engine as the citer count reaches zero; it is cleared by software, or the hardware when the channel is activated. This bit must be cleared in order to write the major.e_link or e_sg bits. |
| 6 | active | Channel active. This flag signals the channel is currently in execution. It is set when channel service begins, and is cleared by the dma_engine as the inner minor loop completes or if any error condition is detected. |
| 5 | major.e_link | Enable channel-to-channel linking on major loop complete. As the channel completes the outer major loop, this flag enables the linking to another channel, defined by major.linkch[5:0]. The link target channel initiates a channel service request via an internal mechanism that sets the TCD.start bit of the specified channel. To support the dynamic linking coherency model, this field is forced to zero when written to while the TCD.done bit is set. 0 The channel-to-channel linking is disabled. 1 The channel-to-channel linking is enabled. |
| 4 | e_sg | Enable scatter/gather processing. As the channel completes the outer major loop, this flag enables scatter/gather processing in the current channel. If enabled, the dma_engine uses dlast_sga as a memory pointer to a 0-modulo-32 address containing a 32-byte data structure which is loaded as the transfer control descriptor into the local memory. To support the dynamic scatter/gather coherency model, this field is forced to zero when written to while the TCD.done bit is set. 0 The current channel's TCD is "normal" format. 1 The current channel's TCD specifies a scatter gather format. The dlast_sga field provides a memory pointer to the next TCD to be loaded into this channel after the outer major loop completes its execution. |
| 3 | d_req | Disable request. If this flag is set, the DMA hardware automatically clears the corresponding DMAERQ bit when the current major iteration count reaches zero. 0 The channel's DMAERQ bit is not affected. 1 The channel's DMAERQ bit is cleared when the outer major loop is complete. |
| 2 | int_half | Enable an interrupt when major counter is half complete. If this flag is set, the channel generates an interrupt request by setting the appropriate bit in the DMAINT register when the current major iteration count reaches the halfway point. Specifically, the comparison performed by the dma_engine is (citer == (biter >> 1)). This halfway point interrupt request is provided to support double-buffered schemes or other types of data movement where the processor needs an early indication of the transfer's progress. The halfway complete interrupt is disabled when biter values are less than two. 0 The half-point interrupt is disabled 1 The half-point interrupt is enabled. |

Table 25-60. TCD Word 7 (TCD.{biter, control/status}) Field Descriptions (continued)

| Bits | Name | Description |
|------|---------|---|
| 1 | int_maj | Enable an interrupt when major iteration count completes. If this flag is set, the channel generates an interrupt request by setting the appropriate bit in the DMAINT register when the current major iteration count reaches zero. 0 The end-of-major loop interrupt is disabled. 1 The end-of-major loop interrupt is enabled. |
| 0 | start | Channel start. If this flag is set, the channel is requesting service. The DMA hardware automatically clears this flag after the channel begins execution. 0 The channel is not explicitly started. 1 The channel is explicitly started via a software initiated service request. |

25.10.4 Functional Description

This section provides an overview of the microarchitecture and functional operation of the DMA module.

25.10.4.1 DMA Microarchitecture

The DMA module is partitioned into two major modules: the `dma_engine` and the transfer control descriptor local memory. Additionally, the `dma_engine` is further partitioned into four submodules, which are detailed below.

- `dma_engine`
 - *addr_path*: This module implements registered versions of two channel transfer control descriptors: channel x and channel y, and is responsible for all the master bus address calculations. All the implemented channels provide the exact same functionality. This hardware structure allows the data transfers associated with one channel to be preempted after the completion of a read/write sequence if a higher priority channel service request is asserted while the first channel is active. Once a channel is activated, it runs until the minor loop is completed unless preempted by a higher priority channel. This capability provides a mechanism (optionally enabled by DCHPRIn[ECP]) where a large data move operation can be preempted to minimize the time another channel is blocked from execution. When any other channel is activated, the contents of its transfer control descriptor is read from the local memory and loaded into the registers of the other `addr_path.channel_{x,y}`. Once the inner minor loop completes execution, the `addr_path` hardware writes the new values for the `TCD.{saddr, daddr, citer}` back into the local memory. If the major iteration count is exhausted, additional processing is performed, including the final address pointer updates, reloading the `TCD.citer` field, and a possible fetch of the next TCD from memory as part of a scatter/gather operation.
 - *data_path*: This module implements the actual bus master read/write datapath. It includes 32 bytes of register storage (matching the maximum transfer size) and the necessary mux logic to support any required data alignment. The AMBA-AHB read data bus is the primary input, and the AHB write data bus is the primary output. The `addr_` and `data_path` modules directly support the two-stage pipelined AMBA-AHB bus. The `addr_path` module represents the first stage of the bus pipeline (the address phase), while the `data_path` module implements the second stage of the pipeline (the data phase).

- *pmodel_charb*: This module implements the first section of DMA programming model as well as the channel arbitration logic. The programming model registers are connected to the IPS bus (not shown). The `ipd_req[n]` inputs and `dma_ipi_int[n]` outputs are also connected to this module (via the control logic).
- *control*: This module provides all the control functions for the `dma_engine`. For data transfers where the source and destination sizes are equal, the `dma_engine` performs a series of source read, destination write operations until the number of bytes specified in the inner minor loop byte count has been moved. For descriptors where the sizes are not equal, multiple access of the smaller size data are required for each reference of the larger size. As an example, if the source size references 16-bit data and the destination is 32-bit data, two reads are performed, then one 32-bit write.
- `transfer_control_descriptor` local memory
 - memory controller: This logic implements the required dual-ported controller, handling accesses from both the `dma_engine` as well as references from the IPS bus. As noted earlier, in the event of simultaneous accesses, the `dma_engine` is given priority and the IPS transaction is stalled. The hooks to a BIST controller for the local TCD memory are included in this module.
 - memory array: The TCD is implemented using a single-ported, synchronous compiled RAM memory array.

25.10.4.2 DMA Basic Data Flow

The basic flow of a data transfer can be partitioned into three segments. As shown in [Figure 25-74](#), the first segment involves the channel service request. In the diagram, this example uses the assertion of the `ipd_req[n]` signal to request service for channel `n`. Channel service request via software and the `TCD.start` bit follows the same basic flow as an `ipd_req`. The `ipd_req[n]` input signal is registered internally and then routed to through the `dma_engine`, first through the control module, then into the programming model/channel arbitration (`pmodel_charb`) module. In the next cycle, the channel arbitration is performed, either using the fixed-priority or round-robin algorithm. After the arbitration is complete, the activated channel number is sent through the address path (`addr_path`) and converted into the required address to access the TCD local memory. Next, the TCD memory is accessed and the required descriptor read from the local memory and loaded into the `dma_engine.addr_path.channel_{x,y}` registers. The TCD memory

is organized 64-bits in width to minimize the time needed to fetch the activated channel's descriptor and load it into the `dma_engine.addr_path.channel_{x,y}` registers.

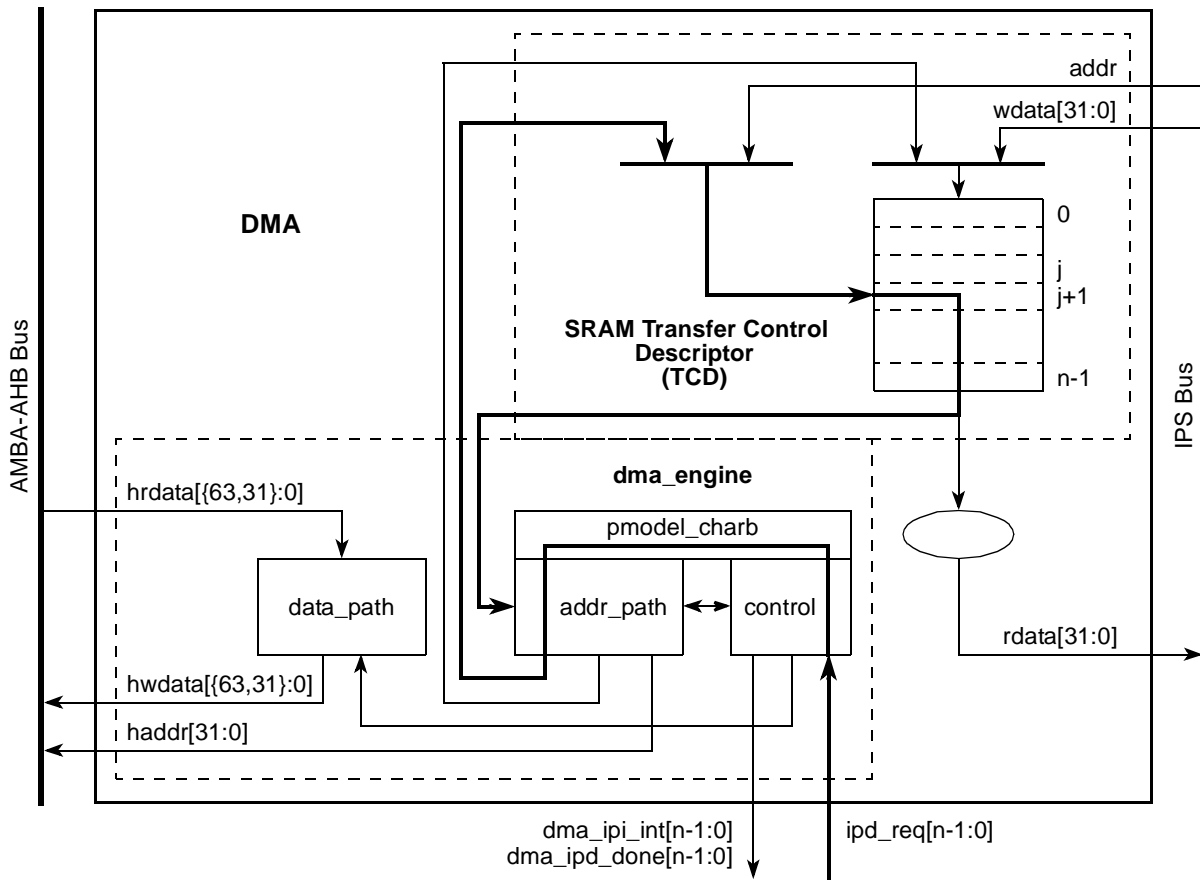


Figure 25-74. DMA Operation—Part 1

In the second part of the basic data flow as shown in [Figure 25-75](#), the modules associated with the data transfer (`addr_path`, `data_path` and `control`) sequence through the required source reads and destination writes to perform the actual data movement. The source reads are initiated and the fetched data is temporarily stored in the `data_path` module until it is gated onto the AMBA-AHB bus during the destination write. This source read/destination write processing continues until the inner minor byte count has been transferred. The `dma_ipd_done[n]` signal is asserted at the end of the minor byte count transfer.

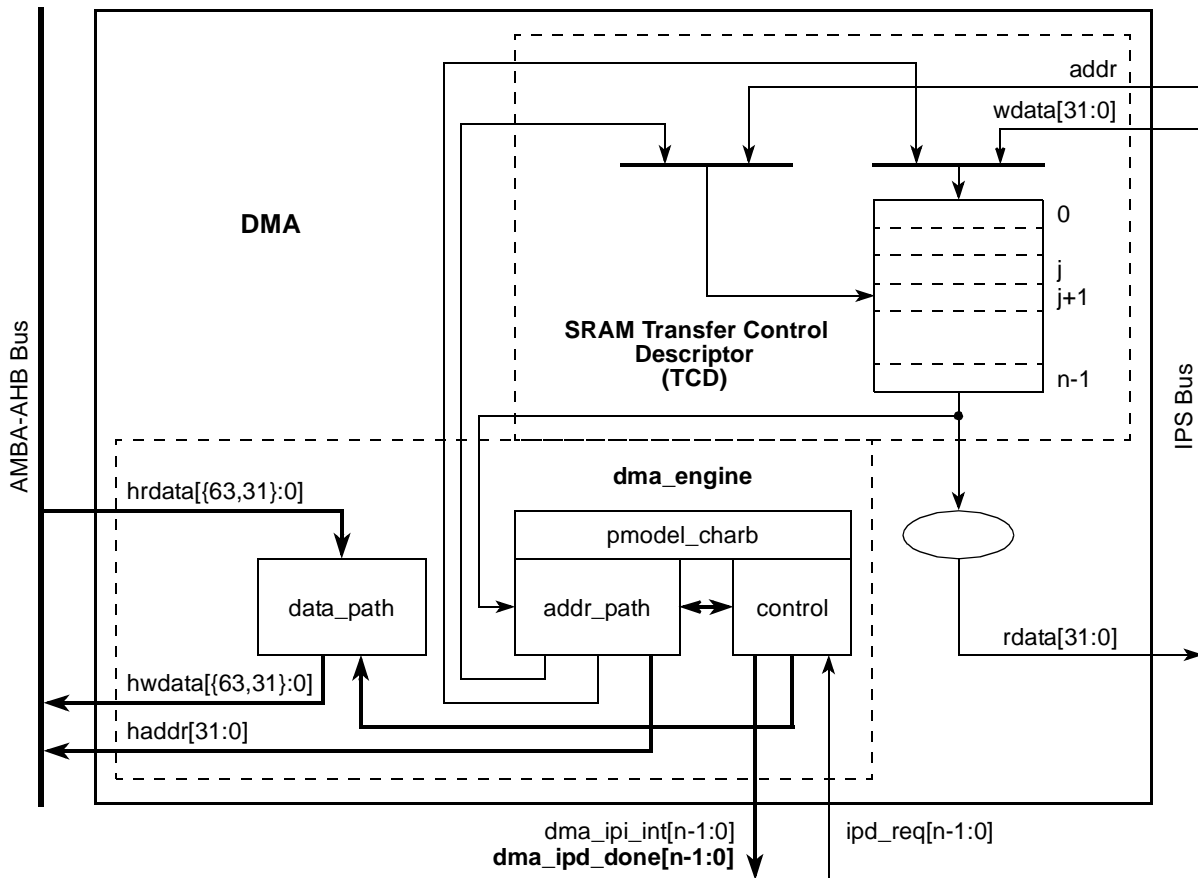


Figure 25-75. DMA Operation—Part 2

Once the inner minor byte count has been moved, the final phase of the basic data flow is performed. In this segment, the `addr_path` logic performs the required updates to certain fields in the channel's TCD, for example, `saddr`, `daddr`, `citer`. If the outer major iteration count is exhausted, then there are additional operations which are performed. These include the final address adjustments and reloading of the biter field into the `citer`. Additionally, assertion of an optional interrupt request occurs at this time, as does a possible fetch of a new TCD from memory using the scatter/gather address pointer included in the

descriptor. The updates to the TCD memory and the assertion of an interrupt request are shown in Figure 25-76.

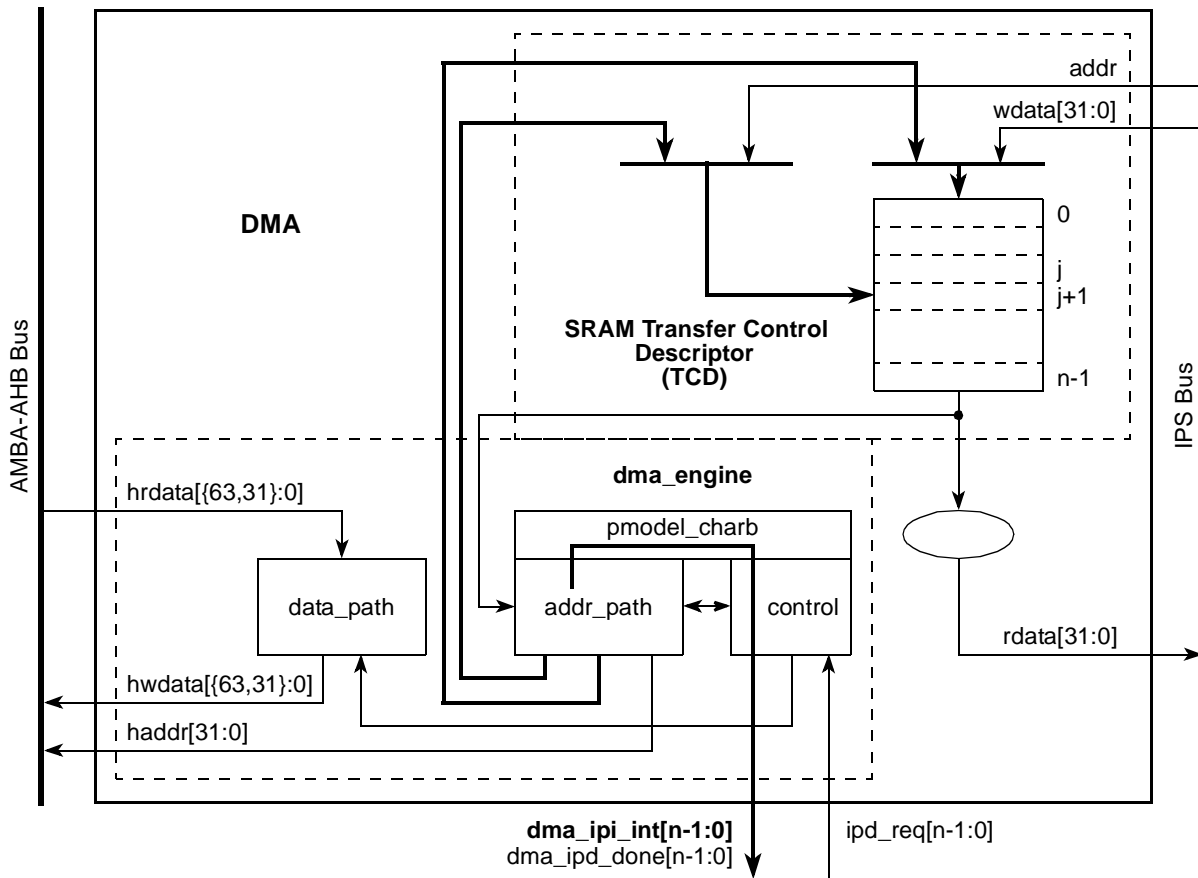


Figure 25-76. DMA Operation—Part 3

25.10.5 Initialization/Application Information

This section discusses DMA initialization, and programming errors.

25.10.5.1 DMA Initialization

The following sequence shows a typical initialization of the DMA:

1. Write the DMACR register if a configuration other than the default is desired
2. Write the channel priority levels into the DCHPRIn registers if a configuration other than the default is desired
3. Enable error interrupts in the DMAEEI registers if so desired
4. Write the 32 byte TCD for each channel that may request service
5. Enable any hardware service requests via the DMAERQ register
6. Request channel service by either software (setting the TCD.start bit) or by hardware (slave device asserting its ipd_req signal).

Once any channel requests service, a channel is selected for execution based on the arbitration and priority levels written into the programmer's model. The `dma_engine` will read the entire TCD for the selected channel into its internal address path module. As the TCD is being read, the first transfer is initiated on the AHB bus unless a configuration error is detected. Transfers from the source (as defined by the source address, `TCD.saddr`) to the destination (as defined by the destination address, `TCD.daddr`) continue until the specified number of bytes (`TCD.nbytes`) have been transferred. When the transfer is complete, the `dma_engine`'s local `TCD.saddr`, `TCD.daddr`, and `TCD.citer` are written back to the main TCD memory and any minor loop channel linking is performed, if enabled. If the major loop is exhausted, further post processing is executed, that is, interrupts, major loop channel linking, and scatter/gather operations, if enabled.

25.10.5.2 DMA Programming Errors

The DMA performs various tests on the transfer control descriptor to verify consistency in the descriptor data. Most programming errors are reported on a per channel basis with the exception of two errors; group priority error and channel priority error, GPE and CPE in the DMAES register, respectively.

For all error types other than group or channel priority errors, the channel number causing the error is recorded in the DMAES register. If the error source is not removed before the next activation of the problem channel, the error will be detected and recorded again.

The sequence listed below is correct. For item 2, the `dma_ipd_ack{done}` lines will assert only if the selected channel is requesting service via the `ipd_req` signal. I think the typical application will enable error interrupts for all channels. So the user will get an error interrupt, but the channel number for the DMAERR register and the error interrupt request line may be wrong because they reflect the selected channel.

Channel priority errors are identified within a group once that group has been selected as the active group. For example:

1. The DMA is configured for fixed group and fixed channel arbitration modes.
2. Group3 is the highest priority and all channels are unique in that group.
3. Group2 is the next highest priority and has two channels with the same priority level.
4. If Group3 has any service requests, those requests will be executed.
5. Once all of Group3 requests have completed, Group2 will be the next active group.
6. If Group2 has a service request, then an undefined channel in Group2 will be selected and a channel priority error will occur.
7. This will repeat until the all of Group2 requests have been removed or a higher priority Group3 request comes in.

A group priority error is global and any request in any group will cause a group priority error.

In general, if priority levels are not unique, the highest (channel/group) priority that has an active request will be selected, but the lowest numbered (channel/group) with that priority will be selected by arbitration and executed by the `dma_engine`. The hardware service request handshake signals, error interrupts and error reporting will be associated with the selected channel.

25.10.6 DMA Transfer

This section discusses DMA transfers for single and multiple requests.

25.10.6.1 Single Request

To perform a simply transfer of n bytes of data with one activation, set the major loop to one (TCD.citer = TCD.biter = 1). The data transfer will begin after the channel service request is acknowledged and the channel is selected to execute. Once the transfer is complete, the TCD.done bit will be set and an interrupt will be generated if properly enabled.

For example, the following TCD entry is configured to transfer 16 bytes of data. The DMA is programmed for one iteration of the major loop transferring 16 bytes per iteration. The source memory has a byte wide memory port located at 0x1000. The destination memory has a word wide port located at 0x2000. The address offsets are programmed in increments to match the size of the transfer; one byte for the source and four bytes for the destination.

The final source and destination addresses are adjusted to return to their beginning values:

- TCD.citer = TCD.biter = 1
- TCD.nbytes = 16
- TCD.saddr = 0x1000
- TCD.soff = 1
- TCD.ssize = 0
- TCD.slast = -16
- TCD.daddr = 0x2000
- TCD.doff = 4
- TCD.dsize = 2
- TCD.dlast_sga= -16
- TCD.int_maj = 1
- TCD.start = 1 (TCD.word7 should be written last after all other fields have been initialized)
- All other TCD fields = 0

This would generate the following sequence of events:

1. IPS write to the TCD.start bit requests channel service
2. The channel is selected by arbitration for servicing
3. dma_engine writes: TCD.done = 0, TCD.start = 0, TCD.active = 1
4. dma_engine reads: channel TCD data from local memory to internal register file
5. The source to destination transfers are executed as follows:
 - a) read_byte(0x1000), read_byte (0x1001), read_byte (0x1002), read_byte (0x1003)
 - b) write_word(0x2000) → first iteration of the minor loop
 - c) read_byte(0x1004), read_byte (0x1005), read_byte (0x1006), read_byte (0x1007)
 - d) write_word(0x2004) → second iteration of the minor loop

- e) read_byte(0x1008), read_byte (0x1009), read_byte (0x100A), read_byte (0x100B)
- f) write_word(0x2008) → third iteration of the minor loop
- g) read_byte(0x100C), read_byte (0x100D), read_byte (0x100E), read_byte (0x100F)
- h) write_word(0x200c) → last iteration of the minor loop → major loop complete
6. dma_engine writes: TCD.saddr = 0x1000, TCD.daddr = 0x2000, TCD.citer = 1 (TCD.biter)
7. dma_engine writes: TCD.active = 0, TCD.done = 1, DMAINT[n] = 1
8. The channel retires

The DMA goes idle or services next channel.

25.10.6.2 Multiple Requests

The next example is the same as previous with the exception of transferring 32 bytes via two hardware requests. The only fields that change are the major loop iteration count and the final address offsets. The DMA is programmed for two iterations of the major loop transferring 16 bytes per iteration. After the channel's hardware requests is enabled in the DMAERQ register, channel service requests are initiated by the slave device.

- TCD.citer = TCD.biter = 2
- TCD.slast = -32
- TCD.dlast_sga = -32

This would generate the following sequence of events:

1. First hardware (ipd_req) request for channel service
2. The channel is selected by arbitration for servicing
3. dma_engine writes: TCD.done = 0, TCD.start = 0, TCD.active = 1
4. dma_engine reads: channel TCD data from local memory to internal register file
5. The source to destination transfers are executed as follows:
 - a) read_byte(0x1000), read_byte (0x1001), read_byte (0x1002), read_byte (0x1003)
 - b) write_word(0x2000) → first iteration of the minor loop
 - c) read_byte(0x1004), read_byte (0x1005), read_byte (0x1006), read_byte (0x1007)
 - d) write_word(0x2004) → second iteration of the minor loop
 - e) read_byte(0x1008), read_byte (0x1009), read_byte (0x100A), read_byte (0x100B)
 - f) write_word(0x2008) → third iteration of the minor loop
 - g) read_byte(0x100c), read_byte (0x100D), read_byte (0x100E), read_byte (0x100F)
 - h) write_word(0x200c) → last iteration of the minor loop
6. dma_engine writes: TCD.saddr = 0x1010, TCD.daddr = 0x2010, TCD.citer = 1
7. dma_engine writes: TCD.active = 0
8. The channel retires → one iteration of the major loop

The DMA goes idle or services next channel.

9. Second hardware (ipd_req) requests channel service

10. The channel is selected by arbitration for servicing
11. dma_engine writes: TCD.done = 0, TCD.start = 0, TCD.active = 1
12. dma_engine reads: channel TCD data from local memory to internal register file
13. The source to destination transfers are executed as follows:
 - a) read_byte(0x1010), read_byte (0x1011), read_byte (0x1012), read_byte (0x1013)
 - b) write_word(0x2010) → first iteration of the minor loop
 - c) read_byte(0x1014), read_byte (0x1015), read_byte (0x1016), read_byte (0x1017)
 - d) write_word(0x2014) → second iteration of the minor loop
 - e) read_byte(0x1018), read_byte (0x1019), read_byte (0x101A), read_byte (0x101B)
 - f) write_word(0x2018) → third iteration of the minor loop
 - g) read_byte(0x101C), read_byte (0x101D), read_byte (0x101E), read_byte (0x101F)
 - h) write_word(0x201C) → last iteration of the minor loop → major loop complete
14. dma_engine writes: TCD.saddr = 0x1000, TCD.daddr = 0x2000, TCD.citer = 2 (TCD.biter)
15. dma_engine writes: TCD.active = 0, TCD.done = 1, DMAINT[n] = 1
16. The channel retires → major loop complete

The DMA goes idle or services the next channel.

25.10.7 TCD Status

This section discusses minor loop complete and active channel TCD reads.

25.10.7.1 Minor Loop Complete

There are two methods to test for minor loop completion when using software initiated service requests. The first method is to read the TCD.citer field and test for a change. Another method may be extracted from the sequence shown below. The second method is to test the TCD.start bit AND the TCD.active bit. The minor loop complete condition is indicated by both bits reading zero after the TCD.start was written to a one. Polling the TCD.active bit may be inconclusive because the active status may be missed if the channel execution is short in duration. The TCD status bits execute the following sequence for a software activated channel:

1. TCD.start = 1, TCD.active = 0, TCD.done = 0 (channel service request via software)
2. TCD.start = 0, TCD.active = 1, TCD.done = 0 (channel is executing)
3. TCD.start = 0, TCD.active = 0, TCD.done = 0 (channel has completed the minor loop and is idle) or
4. TCD.start = 0, TCD.active = 0, TCD.done = 1 (channel has completed the major loop and is idle)

The best method to test for minor loop completion when using hardware initiated service requests is to read the TCD.citer field and test for a change. The hardware request and acknowledge handshakes signals are not visible in the programmer's model.

The TCD status bits execute the following sequence for a hardware activated channel:

1. ipd_req asserts (channel service request via hardware)

2. TCD.start = 0, TCD.active = 1, TCD.done = 0 (channel is executing)
3. TCD.start = 0, TCD.active = 0, TCD.done = 0 (channel has completed the minor loop and is idle) or
4. TCD.start = 0, TCD.active = 0, TCD.done = 1 (channel has completed the major loop and is idle)

For both activation types, the major loop complete status is explicitly indicated via the TCD.done bit.

The TCD.start bit is cleared automatically when the channel begins execution regardless of how the channel was activated.

25.10.7.2 Active Channel TCD Reads

The DMA will read back the 'true' TCD.saddr, TCD.daddr, and TCD.nbytes values if read while a channel is executing. The 'true' values of the saddr, daddr, and nbytes are the values the dma_engine is currently using in its internal register file and not the values in the TCD local memory for that channel. The addresses (saddr and daddr) and nbytes (decrements to zero as the transfer progresses) can give an indication of the progress of the transfer. All other values are read back from the TCD local memory.

25.10.8 Hardware Request Release Timing

This section provides a timing diagram for deasserting the ipd_req hardware request signal. [Figure 25-77](#) shows an encapsulating write (that is, 2 word reads → 1 longword write) with grey indicating the release of the ipd_req hardware request signal.

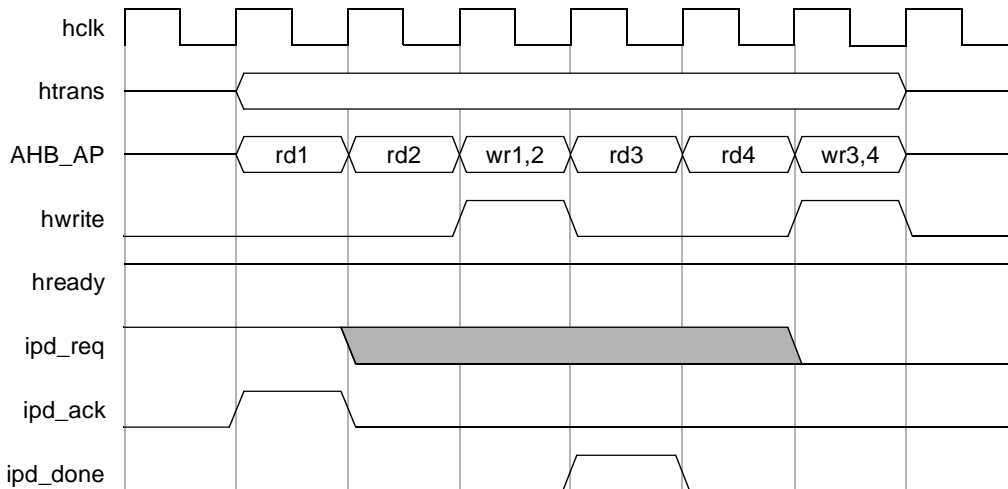


Figure 25-77. ipd_req Removal



Appendix A

Complete List of Configuration, Control, and Status Registers

A.1 Local Access Windows

Table A-1. Local Access Window Registers

| Local Access Window—Block Base Address 0x0_0000 | | | | |
|---|--|--------|--------------------------|------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | Internal memory map base address register (IMMRBAR) | R/W | 0xFF40_0000 | 5.2.4.1/55-6 |
| 0x004 | Reserved | — | — | — |
| 0x008 | Alternate configuration base address register (ALTCBAR) | R/W | 0x0000_0000 | 5.2.4.2/5-7 |
| 0x00C– 0x01C | Reserved | — | — | — |
| 0x020 | eLBC local access window 0 base address register (LBLAWBAR0) | R/W | 0x0000_0000 ¹ | 5.2.4.3/5-8 |
| 0x024 | eLBC local access window 0 attribute register (LBLAWAR0) | R/W | 0x0000_0000 ² | 5.2.4.4/5-9 |
| 0x028 | eLBC local access window 1 base address register (LBLAWBAR1) | R/W | 0x0000_0000 | 5.2.4.3/5-8 |
| 0x02C | eLBC local access window 1 attribute register (LBLAWAR1) | R/W | 0x0000_0000 | 5.2.4.4/5-9 |
| 0x030 | eLBC local access window 2 base address register (LBLAWBAR2) | R/W | 0x0000_0000 | 5.2.4.3/5-8 |
| 0x034 | eLBC local access window 2 attribute register (LBLAWAR2) | R/W | 0x0000_0000 | 5.2.4.4/5-9 |
| 0x038 | eLBC local access window 3 base address register (LBLAWBAR3) | R/W | 0x0000_0000 | 5.2.4.3/5-8 |
| 0x03C | eLBC local access window 3 attribute register (LBLAWAR3) | R/W | 0x0000_0000 | 5.2.4.4/5-9 |
| 0x040– 0x05C | Reserved | — | — | — |
| 0x060 | PCI local access window 0 base address register (PCILAWBAR0) | R/W | 0x0000_0000 ³ | 5.2.4.5/5-10 |
| 0x064 | PCI local access window 0 attribute register (PCILAWAR0) | R/W | 0x0000_0000 ⁴ | 5.2.4.6/5-11 |
| 0x068 | PCI local access window 1 base address register (PCILAWBAR1) | R/W | 0x0000_0000 ⁵ | 5.2.4.5/5-10 |
| 0x06C | PCI local access window 1 attribute register (PCILAWAR1) | R/W | 0x0000_0000 | 5.2.4.6/5-11 |

Table A-1. Local Access Window Registers (continued)

| Local Access Window—Block Base Address 0x0_0000 | | | | |
|---|---|--------|--------------------------|-------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x070–0x07C | Reserved | — | — | — |
| 0x080 | PCI Express 1 local access window base address register (PCIEXP1LAWBAR) | R/W | 0x0000_0000 | 5.2.4.7/5-12 |
| 0x084 | PCI Express 1 local access window attribute register (PCIEXP1LAWAR) | R/W | 0x0000_0000 | 5.2.4.8/5-12 |
| 0x088 | PCI Express 2 local access window base address register (PCIEXP2LAWBAR) | R/W | 0x0000_0000 | 5.2.4.9/5-13 |
| 0x08C | PCI Express 2 local access window attribute register (PCIEXP2LAWAR) | R/W | 0x0000_0000 | 5.2.4.10/5-14 |
| 0x890–0x09C | Reserved | — | — | — |
| 0x0A0 | DDR local access window 0 base address register (DDRLAWBAR0) | R/W | 0x0000_0000 ⁶ | 5.2.4.11/5-14 |
| 0x0A4 | DDR local access window 0 attribute register (DDRLAWAR0) | R/W | 0x0000_0000 ⁷ | 5.2.4.12/5-15 |
| 0x0A8 | DDR local access window 1 base address register (DDRLAWBAR1) | R/W | 0x0000_0000 | 5.2.4.11/5-14 |
| 0x0AC | DDR local access window 1 attribute register (DDRLAWAR1) | R/W | 0x0000_0000 | 5.2.4.12/5-15 |
| 0x0B0–0x0FC | Reserved | — | — | — |

- ¹ Depends on reset configuration word high values. See [Section 5.2.4.3.1, “LBLAWBAR0\[BASE_ADDR\] Reset Value,”](#) for details.
- ² Depends on reset configuration word high values. See [Section 5.2.4.4.1, “LBLAWAR0\[EN\] and LBLAWAR0\[SIZE\] Reset Value,”](#) for details.
- ³ Depends on reset configuration word high values. See [Section 5.2.4.5.1, “PCILAWBAR0\[BASE_ADDR\] Reset Value,”](#) for details.
- ⁴ Depends on reset configuration word high values. See [Section 5.2.4.11.1, “DDRLAWBAR0\[BASE_ADDR\] Reset Value,”](#) for details.
- ⁵ Depends on reset configuration word high values. See [Section 5.2.4.6.1, “PCILAWAR0\[EN\] and PCILAWAR0\[SIZE\] Reset Value,”](#) for details.
- ⁶ Depends on reset configuration word high values. See [Section 5.2.4.11.1, “DDRLAWBAR0\[BASE_ADDR\] Reset Value,”](#) for details.
- ⁷ Depends on reset configuration word high values. See [Section 5.2.4.12.1, “DDRLAWAR0\[EN\] and DDRLAWAR0\[SIZE\] Reset Value,”](#) for details.

A.2 System Configuration Registers

Table A-2. System Configuration Registers

| System Configuration—Block Base Address 0x0_0000 | | | | |
|--|--|--------|-------------------------|----------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x100 | System general purpose register low (SGPRL) | R/W | 0x0000_0000 | 5.3.2.1/5-19 |
| 0x104 | System general purpose register high (SGPRH) | R/W | 0x0000_0000 | 5.3.2.2/5-20 |
| 0x108 | System part and revision ID register (SPRIDR) | R | 0x80B4_0010 | 5.3.2.3/5-20 |
| 0x10C | Reserved | — | — | — |
| 0x110 | System priority configuration register (SPCR) | R/W | 0x0000_0000 | 5.3.2.4/5-21 |
| 0x114 | System I/O configuration register low (SICRL) | R/W | 0xF00F00_n0n0 | 5.3.2.5/5-23 |
| 0x118 | System I/O configuration register high (SICRH) | R/W | 0xFFFF_C0n ³ | 5.3.2.6/5-26 |
| 0x11C– 0x1FC | Reserved | — | — | — |
| 0x128 | DDR control driver register (DDRCDR) | R/W | 0x7B84_0001 | 5.3.2.8/5-29 |
| 0x12C | DDR debug status register (DDRDSR) | R | 0x3B80_0000 | 5.3.2.9/5-31 |
| 0x140 | PCI Express control register 1 (PECR1) | R/W | All zeros | 5.3.2.10/55-31 |
| 0x144 | PCI Express control register 2 (PECR2) | R/W | All zeros | 5.3.2.10/55-31 |
| 0x148– 0x1FC | Reserved | — | — | — |

A.3 Watchdog Timer (WDT)

Table A-3. Watchdog Timer (WDT) Registers

| Watchdog Timer (WDT)—Block Base Address 0x0_0200 | | | | |
|--|--|--------|---|--------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000– 0x003 | Reserved | — | — | — |
| 0x004 | System watchdog control register (SWCRR) | R/W | 0xFFFF_0003 or 0xFFFF_0007 ¹ | 5.4.4.1/5-35 |
| 0x008 | System watchdog count register (SWCNR) | R | 0x0000_FFFF | 5.4.4.2/5-36 |
| 0x00C– 0x00D | Reserved | — | — | — |
| 0x00E | System watchdog service register (SWSRR) | R/W | 0x0000 | 5.4.4.3/5-36 |

¹ SWCRR[SWEN] reset value directly depends on RCWHR[SWEN] (reset configuration word high).

A.4 Real Time Clock (RTC)

Table A-4. Real Time Clock (RTC) Registers

| Real Time Clock (RTC)—Block Base Address 0x0_0300 | | | | |
|---|---|--------|-------------|------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x00 | Real time counter control register (RTCNR) | R/W | 0x0000_0000 | 5.5.5.1/5-42 |
| 0x04 | Real time counter load register (RTLDR) | R/W | 0x0000_0000 | 5.5.5.2/5-43 |
| 0x08 | Real time counter prescale register (RTPSR) | R/W | 0x0000_0000 | 5.5.5.3/5-43 |
| 0x0C | Real time counter register (RTCTR) | R | 0x0000_0000 | 5.5.5.4/5-43 |
| 0x10 | Real time counter event register (RTEVR) | w1c | 0x0000_0000 | 5.5.5.5/5-44 |
| 0x14 | Real time counter alarm register (RTALR) | R/W | 0xFFFF_FFFF | 5.5.5.6/5-45 |
| 0x18–0x1F | Reserved | — | — | — |

A.5 Periodic Interval Timer (PIT)

Table A-5. Periodic Interval Timer (PIT) Registers

| Periodic Interval Timer (PIT)—Block Base Address 0x0_0400 | | | | |
|---|---|--------|-------------|------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x00 | Periodic interval timer control register (PTCNR) | R/W | 0x0000_0000 | 5.6.5.1/5-49 |
| 0x04 | Periodic interval timer load register (PTLDR) | R/W | 0x0000_0000 | 5.6.5.2/5-50 |
| 0x08 | Periodic interval timer prescale register (PTPSR) | R/W | 0x0000_0000 | 5.6.5.3/5-50 |
| 0x0C | Periodic interval timer counter register (PTCTR) | R | 0x0000_0000 | 5.6.5.4/5-51 |
| 0x10 | Periodic interval timer event register (PTEVR) | w1c | 0x0000_0000 | 5.6.5.5/5-51 |
| 0x14–0x1F | Reserved | — | — | — |

A.6 General Purpose (Global) Timers (GTMs)

Table A-6. General Purpose (Global) Timers (GTMs) Registers

| General Purpose (Global) Timer Module 1—Block Base Address 0x0_0500 General Purpose (Global) Timer Module 2—Block Base Address 0x0_0600 | | | | |
|--|---|--------|--------|------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x00 | Timer 1 and 2 global timers configuration register (GTCFR1) | R/W | 0x00 | 5.7.5.1/5-59 |
| 0x01–0x03 | Reserved | — | — | — |
| 0x04 | Timer 3 and 4 global timers configuration register (GTCFR2) | R/W | 0x00 | 5.7.5.1/5-59 |
| 0x05–0x0F | Reserved | — | — | — |
| 0x10 | Timer 1 global timers mode register (GTMDR1) | R/W | 0x0000 | 5.7.5.2/5-62 |
| 0x12 | Timer 2 global timers mode register (GTMDR2) | | | |

Table A-6. General Purpose (Global) Timers (GTMs) Registers (continued)

| General Purpose (Global) Timer Module 1—Block Base Address 0x0_0500 General Purpose (Global) Timer Module 2—Block Base Address 0x0_0600 | | | | |
|--|---|--------|--------|--------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x14 | Timer 1 global timers reference register (GTRFR1) | R/W | 0xFFFF | 5.7.5.3/5-63 |
| 0x16 | Timer 2 global timers reference register (GTRFR2) | | | |
| 0x18 | Timer 1 global timers capture register (GTCPR1) | R/W | 0x0000 | 5.7.5.4/5-63 |
| 0x1A | Timer 2 global timers capture register (GTCPR2) | | | |
| 0x1C | Timer 1 global timers counter register (GTCNR1) | R/W | 0x0000 | 5.7.5.5/5-64 |
| 0x1E | Timer 2 global timers counter register (GTCNR2) | | | |
| 0x20 | Timer 3 global timers mode register (GTMDR3) | R/W | 0x0000 | 5.7.5.2/5-62 |
| 0x22 | Timer 4 global timers mode register (GTMDR4) | | | |
| 0x24 | Timer 3 global timers reference register (GTRFR3) | R/W | 0xFFFF | 5.7.5.3/5-63 |
| 0x26 | Timer 4 global timers reference register (GTRFR4) | | | |
| 0x28 | Timer 3 global timers capture register (GTCPR3) | R | 0x0000 | 5.7.5.4/5-63 |
| 0x2A | Timer 4 global timers capture register (GTCPR4) | | | |
| 0x2C | Timer 3 global timers counter register (GTCNR3) | R/W | 0x0000 | 5.7.5.5/5-64 |
| 0x2E | Timer 4 global timers counter register (GTCNR4) | | | |
| 0x30 | Timer 1 global timers event register (GTEVR1) | w1c | 0x0000 | 5.7.5.6/5-64 |
| 0x32 | Timer 2 global timers event register (GTEVR2) | | | |
| 0x34 | Timer 3 global timers event register (GTEVR3) | | | |
| 0x36 | Timer 4 global timers event register (GTEVR4) | | | |
| 0x38 | Timer 1 global timers prescale register (GTPSR1) | R/W | 0x0003 | 5.7.5.7/5-65 |
| 0x3A | Timer 2 global timers prescale register (GTPSR2) | | | |
| 0x3C | Timer 3 global timers prescale register (GTPSR3) | | | |
| 0x3E | Timer 4 global timers prescale register (GTPSR4) | | | |

A.7 Integrated Programmable Interrupt Controller (IPIC)

Table A-7. Integrated Programmable Interrupt Controller Registers

| Integrated Programmable Interrupt Controller—Block Base Address 0x0_0700 | | | | |
|--|--|--------|-------------|--------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x00 | System global interrupt configuration register (SICFR) | R/W | 0x0000_0000 | 8.5.1/8-8 |
| 0x04 | System regular interrupt vector register (SIVCR) | R | 0x0000_0000 | 8.5.2/8-9 |
| 0x08 | System internal interrupt pending register (SIPNR_H) | R | 0x0000_0000 | 8.5.3/8-12 |

Table A-7. Integrated Programmable Interrupt Controller Registers (continued)

| | | | | |
|-----------|---|-----|-------------|-------------|
| 0x0C | System internal interrupt pending register (SIPNR_L) | R | 0x0000_0000 | 8.5.3/8-12 |
| 0x10 | System internal interrupt group A priority register (SIPRR_A) | R/W | 0x0530_9770 | 8.5.4/8-15 |
| 0x14 | System internal interrupt group B priority register (SIPRR_B) | R/W | 0x0530_9770 | 8.5.5/8-16 |
| 0x18 | System internal interrupt group C priority register (SIPRR_C) | R/W | 0x0530_9770 | 8.5.6/8-16 |
| 0x1C | System internal interrupt group D priority register (SIPRR_D) | R/W | 0x0530_9770 | 8.5.7/8-17 |
| 0x20 | System internal interrupt mask register (SIMSR_H) | R/W | 0x0000_0000 | 8.5.8/8-18 |
| 0x24 | System internal interrupt mask register (SIMSR_L) | R/W | 0x0000_0000 | 8.5.8/8-18 |
| 0x28 | System internal interrupt control register (SICNR) | R/W | 0x0000_0000 | 8.5.9/8-19 |
| 0x2C | System external interrupt pending register (SEPNR) | R/W | Special | 8.5.10/8-21 |
| 0x30 | System mixed interrupt group A priority register (SMPRR_A) | R/W | 0x0530_9770 | 8.5.11/8-22 |
| 0x34 | System mixed interrupt group B priority register (SMPRR_B) | R/W | 0x0530_9770 | 8.5.12/8-23 |
| 0x38 | System external interrupt mask register (SEMSR) | R/W | 0x0000_0000 | 8.5.13/8-23 |
| 0x3C | System external interrupt control register (SECNR) | R/W | 0x0000_0000 | 8.5.14/8-24 |
| 0x40 | System error status register (SERSR) | R/W | 0x0000_0000 | 8.5.15/8-26 |
| 0x44 | System error mask register (SERMR) | R/W | 0xFFFE0000 | 8.5.16/8-27 |
| 0x48 | System error control register (SERCR) | R/W | 0x0000_0000 | 8.5.17/8-27 |
| 0x4C | System external interrupt polarity control register (SEPCR) | R/W | 0x0000_0000 | 8.5.18/8-28 |
| 0x4F | Reserved | — | — | — |
| 0x50 | System internal interrupt force register (SIFCR_H) | R/W | 0x0000_0000 | 8.5.19/8-29 |
| 0x54 | System internal interrupt force register (SIFCR_L) | R/W | 0x0000_0000 | 8.5.19/8-29 |
| 0x58 | System external interrupt force register (SEFCR) | R/W | 0x0000_0000 | 8.5.20/8-30 |
| 0x5C | System error force register (SERFR) | R/W | 0x0000_0000 | 8.5.21/8-30 |
| 0x60 | System critical interrupt vector register (SCVCR) | R | 0x0000_0000 | 8.5.22/8-31 |
| 0x64 | System management interrupt vector register (SMVCR) | R | 0x0000_0000 | 8.5.23/8-31 |
| 0x68–0xBF | Reserved | — | — | — |

A.8 System Arbiter

Table A-8. System Arbiter Registers

| System Arbiter—Block Base Address 0x0_0800 | | | | |
|--|--------------------------------------|--------|--|--------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x00 | Arbiter configuration register (ACR) | R/W | 0x0000_0000/ 0x0010_0000 ¹ | 6.2.1/6-3 |
| 0x04 | Arbiter timers register (ATR) | R/W | FFFF_FFFF | 6-1/6-2 |
| 0x08 | Arbiter Event Enable Register (AEER) | R/W | 0x0000_007F | 6.2.3/6-5 |

Table A-8. System Arbiter Registers (continued)

| System Arbiter—Block Base Address 0x0_0800 | | | | |
|--|--|--------|--------------------------|--------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x0C | Arbiter event register (AER) | w1c | 0x0000_0000 | 6-1/6-2 |
| 0x10 | Arbiter interrupt definition register (AIDR) | R/W | 0x0000_0000 | 6-1/6-2 |
| 0x14 | Arbiter mask register (AMR) | R/W | 0x0000_0000 | 6-1/6-2 |
| 0x18 | Arbiter event attributes register (AEATR) | R | 0x0000_0000 ² | 6-1/6-2 |
| 0x1C | Arbiter event address register (AEADR) | R | 0x0000_0000 ² | 6-1/6-2 |
| 0x20 | Arbiter event response register (AERR) | R/W | 0x0000_0000 | 6-1/6-2 |

¹ Reset value is determined from the core PLL configuration of the reset configuration word. See [Chapter 4, “Reset, Clocking, and Initialization,”](#) for details.

² The registers AEATR and AEADR are affected only by the assertion of $\overline{\text{PORESET}}$

A.9 Reset Configuration

Table A-9. Reset Configuration Registers

| Reset Configuration—Block Base Address 0x0_0900 | | | | |
|---|--|--------|-------------|--------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | Reset configuration word low register (RCWLR) | R | 0x0000_0000 | 4.5.1/4-33 |
| 0x004 | Reset configuration word high register (RCWHR) | R | 0x0000_0000 | 4.5.1/4-33 |
| 0x008 | Reserved, should be cleared | — | — | — |
| 0x00C | Reserved, should be cleared | — | — | — |
| 0x010 | Reset status register (RSR) | R/W | 0x0000_0000 | 4.5.1/4-33 |
| 0x014 | Reset mode register (RMR) | R/W | 0x0000_0000 | 4.5.1/4-33 |
| 0x018 | Reset protection register (RPR) | R/W | 0x0000_0000 | 4.5.1/4-33 |
| 0x01C | Reset control register (RCR) | R/W | 0x0000_0000 | 4.5.1/4-33 |
| 0x020 | Reset control enable register (RCER) | R/W | 0x0000_0000 | 4.5.1/4-33 |
| 0x024– 0x0FC | Reserved, should be cleared. | — | — | — |

A.10 Clock Configuration

Table A-10. Clock Configuration Registers

| Reset Configuration—Block Base Address 0x0_0A00 | | | | |
|---|--------------------------------------|--------|--------------------------------------|--------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | System PLL mode register (SPMR) | R | 0xn _{nnn} _n _{nnn} | 4.5.1/4-33 |
| 0x004 | Output clock control register (OCCR) | R/W | 0x0000_C0C0 | 4.5.1/4-33 |

Table A-10. Clock Configuration Registers (continued)

| Reset Configuration—Block Base Address 0x0_0A00 | | | | |
|---|--------------------------------------|--------|-------------|----------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x008 | System clock control register (SCCR) | R/W | 0x5155_1410 | 4.5.1/4-33 |
| 0x00C–0x0FC | Reserved, should be cleared | — | — | — |

A.11 Power Management Controller (PMC)

Table A-11. Power Management Controller (PMC) Registers

| Power Management Controller—Block Base Address 0x0_0B00 | | | | |
|---|---|--------|-------------|------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | Power management controller configuration register (PMCCR) | R/W | 0x0000_0000 | 5.8.2.1/5-71 |
| 0x004 | Power management controller event register (PMCER) | R/W | 0x0000_0000 | 5.8.2.2/5-72 |
| 0x008 | Power management controller mask register (PMCMR) | R/W | 0x0000_0000 | 5.8.2.3/5-74 |
| 0x00C | Power management controller configuration register 1 (PMCCR1) | R/W | 0x0000_0000 | 5.8.2.4/5-75 |
| 0x010 | Power management controller configuration register 2 (PMCCR2) | R/W | 0x0002_0002 | 5.8.2.5/5-77 |
| 0x014–0x0FC | Reserved | — | — | — |

A.12 General Purpose I/O (GPIO)

Table A-12. General Purpose I/O (GPIO) Registers

| General Purpose I/O (GPIO)—Block Base Address 0x0_0C00 | | | | |
|--|--|--------|-------------|-----------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | GPIO direction register (GPDIR) | R/W | 0x0000_0000 | 24.3.1/24-3 |
| 0xC04 | GPIO open drain register (GPODR) | R/W | 0x0000_0000 | 24.3.2/24-3 |
| 0x008 | GPIO data register (GPDAT) | R/W | 0x0000_0000 | 24.3.3/24-4 |
| 0x00C | GPIO interrupt event register (GPIER) | w1c | Undefined | 24.3.4/24-4 |
| 0x010 | GPIO interrupt mask register (GPIMR) | R/W | 0x0000_0000 | 24.3.5/24-4 |
| 0x014 | GPIO external interrupt control register (GPICR) | R/W | 0x0000_0000 | 24.3.6/24-5 |

A.13 DDR Memory Controller

Table A-13. DDR Memory Controller Registers

| DDR Memory Controller—Block Base Address 0x0_2000 | | | | |
|---|---|--------|--|-------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | CS0_BNDS—Chip select 0 memory bounds | R/W | 0x0000_0000 | 9.4.1.1/9-9 |
| 0x008 | CS1_BNDS—Chip select 1 memory bounds | R/W | 0x0000_0000 | 9.4.1.1/9-9 |
| 0x080 | CS0_CONFIG—Chip select 0 configuration | R/W | 0x0000_0000 | 9.4.1.2/9-10 |
| 0x084 | CS1_CONFIG—Chip select 1 configuration | R/W | 0x0000_0000 | 9.4.1.2/9-10 |
| 0x100 | TIMING_CFG_3—DDR SDRAM timing configuration 3 | R/W | 0x0000_0000 | 9.4.1.3/9-12 |
| 0x104 | TIMING_CFG_0—DDR SDRAM timing configuration 0 | R/W | 0x0011_0105 | 9.4.1.4/9-12 |
| 0x108 | TIMING_CFG_1—DDR SDRAM timing configuration 1 | R/W | 0x0000_0000 | 9.4.1.5/9-14 |
| 0x10C | TIMING_CFG_2—DDR SDRAM timing configuration 2 | R/W | 0x0000_0000 | 9.4.1.6/9-16 |
| 0x110 | DDR_SDRAM_CFG—DDR SDRAM control configuration | R/W | 0x0200_0000 | 9.4.1.7/9-18 |
| 0x114 | DDR_SDRAM_CFG_2—DDR SDRAM control configuration 2 | R/W | 0x0000_0000 | 9.4.1.8/9-21 |
| 0x118 | DDR_SDRAM_MODE—DDR SDRAM mode configuration | R/W | 0x0000_0000 | 9.4.1.9/9-22 |
| 0x11C | DDR_SDRAM_MODE_2—DDR SDRAM mode configuration 2 | R/W | 0x0000_0000 | 9.4.1.10/9-23 |
| 0x120 | DDR_SDRAM_MD_CNTL—DDR SDRAM mode control | R/W | 0x0000_0000 | 9.4.1.11/9-24 |
| 0x124 | DDR_SDRAM_INTERVAL—DDR SDRAM interval configuration | R/W | 0x0000_0000 | 9.4.1.12/9-27 |
| 0x128 | DDR_DATA_INIT—DDR SDRAM data initialization | R/W | 0x0000_0000 | 9.4.1.13/9-27 |
| 0x130 | DDR_SDRAM_CLK_CNTL—DDR SDRAM clock control | R/W | 0x0200_0000 | 9.4.1.14/9-28 |
| 0x140–0x144 | Reserved | — | — | — |
| 0x148 | DDR_INIT_ADDR—DDR training initialization address | R/W | 0x0000_0000 | 9.4.1.15/9-28 |
| 0x150–0xBF4 | Reserved | — | — | — |
| 0xBF8 | DDR_IP_REV1—DDR IP block revision 1 | R | 0xn ¹ nnn ¹ _nnn ¹ | 9.4.1.16/9-29 |
| 0xBFC | DDR_IP_REV2—DDR IP block revision 2 | R | 0x00 ¹ nn ¹ _00 ¹ nn ¹ | 9.4.1.17/9-29 |
| 0xE40 | ERR_DETECT—Memory error detect | w1c | All zeros | 9.4.1.18/9-30 |
| 0xE44 | ERR_DISABLE—Memory error disable | R/W | All zeros | 9.4.1.19/9-30 |
| 0xE48 | ERR_INT_EN—Memory error interrupt enable | R/W | All zeros | 9.4.1.20/9-31 |

¹ Implementation-dependent reset values are listed in specified section/page.

A.14 I²C Controller

Table A-14. I²C Controller Registers

| I ² C Controller—Block Base Address 0x0_3000 | | | | |
|---|---|--------|-------|-------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | I2CADR—I ² C address register | R/W | 0x00 | 20.3.1.1/20-4 |
| 0x004 | I2CFDR—I ² C frequency divider register | R/W | 0x00 | 20.3.1.2/20-5 |
| 0x008 | I2CCR—I ² C control register | R/W | 0x00 | 20.3.1.3/20-6 |
| 0x00C | I2CSR—I ² C status register | R/W | 0x81 | 20.3.1.4/20-7 |
| 0x010 | I2CDR—I ² C data register | R/W | 0x00 | 20.3.1.5/20-9 |
| 0x014 | I2CDFSRR—I ² C digital filter sampling rate register | R/W | 0x10 | 20.3.1.6/20-9 |
| 0x01C–0x1FF | Reserved, should be cleared | — | — | — |

A.15 DUART

Table A-15. DUART Registers

| UART 1—Block Base Address 0x0_4000 UART 2—Block Base Address 0x0_4100 | | | | |
|--|---|--------|-------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x500 | URBR—ULCR[DLAB] = 0 UART1 receiver buffer register | R | 0x00 | 21.3.1.1/21-6 |
| | UTHR—ULCR[DLAB] = 0 UART1 transmitter holding register | W | 0x00 | 21.3.1.2/21-6 |
| | UDLB—ULCR[DLAB] = 1 UART1 divisor least significant byte register | R/W | 0x00 | 21.3.1.3/21-7 |
| 0x501 | UIER—ULCR[DLAB] = 0 UART1 interrupt enable register | R/W | 0x00 | 21.3.1.4/21-8 |
| | UDMB—ULCR[DLAB] = 1 UART1 divisor most significant byte register | R/W | 0x00 | 21.3.1.3/21-7 |
| 0x502 | UIIR—ULCR[DLAB] = 0 UART1 interrupt ID register | R | 0x01 | 21.3.1.5/21-9 |
| | UFCR—ULCR[DLAB] = 0 UART1 FIFO control register | W | 0x00 | 21.3.1.6/21-10 |
| | UAFR—ULCR[DLAB] = 1 UART1 alternate function register | R/W | 0x00 | 21.3.1.7/21-11 |
| 0x503 | ULCR—ULCR[DLAB] = x UART1 line control register | R/W | 0x00 | 21.3.1.8/21-12 |
| 0x504 | UMCR—ULCR[DLAB] = x UART1 MODEM control register | R/W | 0x00 | 21.3.1.9/21-14 |
| 0x505 | ULSR—ULCR[DLAB] = x UART1 line status register | R | 0x60 | 21.3.1.10/21-15 |
| 0x506 | UMSR—ULCR[DLAB] = x UART1 MODEM status register | R | 0x00 | 21.3.1.11/21-16 |
| 0x507 | USCR—ULCR[DLAB] = x UART1 scratch register | R/W | 0x00 | 21.3.1.12/21-17 |
| 0x510 | UDSR—ULCR[DLAB] = x UART1 DMA status register | R | 0x01 | 21.3.1.13/21-17 |
| 0x600 | URBR—ULCR[DLAB] = 0 UART2 receiver buffer register | R | 0x00 | 21.3.1.1/21-6 |
| | UTHR—ULCR[DLAB] = 0 UART2 transmitter holding register | W | 0x00 | 21.3.1.2/21-6 |
| | UDLB—ULCR[DLAB] = 1 UART2 divisor least significant byte register | R/W | 0x00 | 21.3.1.3/21-7 |

Table A-15. DUART Registers (continued)

| UART 1—Block Base Address 0x0_4000 UART 2—Block Base Address 0x0_4100 | | | | |
|--|--|--------|-------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x601 | UIER—ULCR[DLAB] = 0 UART2 interrupt enable register | R/W | 0x00 | 21.3.1.4/21-8 |
| | UDMB—ULCR[DLAB] = 1 UART2 divisor most significant byte register | R/W | 0x00 | 21.3.1.3/21-7 |
| 0x602 | UIIR—ULCR[DLAB] = 0 UART2 interrupt ID register | R | 0x01 | 21.3.1.5/21-9 |
| | UFGR—ULCR[DLAB] = 0 UART2 FIFO control register | W | 0x00 | 21.3.1.6/21-10 |
| | UAFR—ULCR[DLAB] = 1 UART2 alternate function register | R/W | 0x00 | 21.3.1.7/21-11 |
| 0x603 | ULCR—ULCR[DLAB] = x UART2 line control register | R/W | 0x00 | 21.3.1.8/21-12 |
| 0x604 | UMCR—ULCR[DLAB] = x UART2 MODEM control register | R/W | 0x00 | 21.3.1.9/21-14 |
| 0x605 | ULSR—ULCR[DLAB] = x UART2 line status register | R | 0x60 | 21.3.1.10/21-15 |
| 0x606 | UMSR—ULCR[DLAB] = x UART2 MODEM status register | R | 0x00 | 21.3.1.11/21-16 |
| 0x607 | USCR—ULCR[DLAB] = x UART2 scratch register | R/W | 0x00 | 21.3.1.12/21-17 |
| 0x610 | UDSR—ULCR[DLAB] = x UART2 DMA status register | R | 0x01 | 21.3.1.13/21-17 |

A.16 Enhanced Local Bus Controller (eLBC)

Table A-16. Enhanced Local Bus Controller (eLBC) Registers

| Enhanced Local Bus Controller (eLBC)—Block Base Address 0x0_5000 | | | | |
|--|--------------------------|--------|--------------|--------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | BR0—Base register 0 | R/W | 0x0000_ nnnn | 10.3.1.1/10-9 |
| 0x008 | BR1—Base register 1 | R/W | 0x0000_0000 | 10.3.1.1/10-9 |
| 0x010 | BR2—Base register 2 | R/W | 0x0000_0000 | 10.3.1.1/10-9 |
| 0x018 | BR3—Base register 3 | R/W | 0x0000_0000 | 10.3.1.1/10-9 |
| 0x020–0x038 | Reserved | R/W | 0x0000_0000 | 10.3.1.1/10-9 |
| 0x004 | OR0—Options register 0 | R/W | 0x0000_0FF7 | 10.3.1.2/10-11 |
| 0x00C | OR1—Options register 1 | R/W | 0x0000_0000 | 10.3.1.2/10-11 |
| 0x014 | OR2—Options register 2 | R/W | 0x0000_0000 | 10.3.1.2/10-11 |
| 0x01C | OR3—Options register 3 | R/W | 0x0000_0000 | 10.3.1.2/10-11 |
| 0x024–0x064 | Reserved | — | — | — |
| 0x068 | MAR—UPM address register | R/W | 0x0000_0000 | 10.3.1.3/10-19 |
| 0x06C | Reserved | — | — | — |
| 0x070 | MAMR—UPMA mode register | R/W | 0x0000_0000 | 10.3.1.4/10-20 |
| 0x074 | MBMR—UPMB mode register | R/W | 0x0000_0000 | 10.3.1.4/10-20 |

Table A-16. Enhanced Local Bus Controller (eLBC) Registers (continued)

| Enhanced Local Bus Controller (eLBC)—Block Base Address 0x0_5000 | | | | |
|--|---|--------|-------------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x078 | MCMR—UPMC mode register | R/W | 0x0000_0000 | 10.3.1.4/10-20 |
| 0x07C–0x080 | Reserved | — | — | — |
| 0x084 | MRTPR—Memory refresh timer prescaler register | R/W | 0x0000_0000 | 10.3.1.5/10-22 |
| 0x088 | MDR—UPM/FCM data register | R/W | 0x0000_0000 | 10.3.1.6/10-22 |
| 0x08C | Reserved | — | — | — |
| 0x090 | LSOR—Special operation initiation register | R/W | 0x0000_0000 | 10.3.1.7/10-23 |
| 0x094–0x09C | Reserved | — | — | — |
| 0x0A0 | LURT—UPM refresh timer | R/W | 0x0000_0000 | 10.3.1.4/10-20 |
| 0x0A4–0x0AC | Reserved | — | — | — |
| 0x0B0 | LTESR—Transfer error status register | w1c | 0x0000_0000 | 10.3.1.9/10-25 |
| 0x0B4 | LTEDR—Transfer error disable register | R/W | 0x0000_0000 | 10.3.1.10/10-27 |
| 0x0B8 | LTEIR—Transfer error interrupt register | R/W | 0x0000_0000 | 10.3.1.11/10-28 |
| 0x0BC | LTEATR—Transfer error attributes register | R/W | 0x0000_0000 | 10.3.1.12/10-29 |
| 0x0C0 | LTEAR—Transfer error address register | R/W | 0x0000_0000 | 10.3.1.13/10-30 |
| 0x0C4 | LTECCR—Transfer error ECC register | w1c | 0x0000_0000 | 10.3.1.14/10-30 |
| 0x0C48–0x0CC | Reserved | — | — | — |
| 0x0D0 | LBCR—Configuration register | R/W | 0x0004_0000 | 10.3.1.14/10-30 |
| 0x0D4 | LCRR—Clock ratio register | R/W | 0x8000_0008 | 10.3.1.15/10-32 |
| 0x0D8–0x0DC | Reserved | — | — | — |
| 0x0E0 | FMR—Flash mode register | R/W | 0x0000_0000 | 10.3.1.16/10-33 |
| 0x0E4 | FIR—Flash instruction register | R/W | 0x0000_0000 | 10.3.1.17/10-34 |
| 0x0E8 | FCR—Flash command register | R/W | 0x0000_0000 | 10.3.1.18/10-35 |
| 0x0EC | FBAR—Flash block address register | R/W | 0x0000_0000 | 10.3.1.19/10-36 |
| 0x0F0 | FPAR—Flash page address register | R/W | 0x0000_0000 | 10.3.1.20/10-36 |
| 0x0F4 | FBCR—Flash byte count register | R/W | 0x0000_0000 | 10.3.1.21/10-38 |
| 0x0F8–0x0FC | Reserved | — | — | — |
| 0x100 | FECC0—Flash ECC block 0 register | R/O | 0x0000_0000 | 10.3.1/10-9 |
| 0x104 | FECC1—Flash ECC block 1 register | R/O | 0x0000_0000 | |
| 0x108 | FECC2—Flash ECC block 2 register | R/O | 0x0000_0000 | |
| 0x10C | FECC3—Flash ECC block 3 register | R/O | 0x0000_0000 | |

Table A-16. Enhanced Local Bus Controller (eLBC) Registers (continued)

| Enhanced Local Bus Controller (eLBC)—Block Base Address 0x0_5000 | | | | |
|--|---------------------|--------|-------------|-------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x110–0xFFC | Reserved | — | — | — |
| 0x000 | BR0—Base register 0 | R/W | 0x0000_0000 | 10.3.1.1/10-9 |

A.17 Serial Peripheral Interface (SPI)

Table A-17. Serial Peripheral Interface (SPI) Registers

| Serial Peripheral Interface (SPI)—Block Base Address 0x0_7000 | | | | |
|---|-------------------------------|--------|-------------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000–0x01F | Reserved | — | — | — |
| 0x020 | SPI mode register (SPMODE) | R/W | 0x0000_0000 | 22.4.1.1/2222-9 |
| 0x024 | SPI event register (SPIE) | Mixed | 0x0000_0000 | 22.4.1.2/2222-12 |
| 0x028 | SPI mask register (SPIM) | R/W | 0x0000_0000 | 22.4.1.3/2222-13 |
| 0x02C | SPI command register (SPCOM) | W | 0x0000_0000 | 22.4.1.4/2222-14 |
| 0x030 | SPI transmit register (SPITD) | W | 0x0000_0000 | 22.4.1.5/2222-14 |
| 0x034 | SPI receive register (SPIRD) | R | 0xFFFF_FFFF | 22.4.1.6/2222-15 |
| 0x038–0xFFFF | Reserved | — | — | — |

A.18 DMA Controller

Table A-18. DMA Controller Registers

| DMA 0—Block Base Address 0x0_8100 DMA 1—Block Base Address 0x0_8180 DMA 2—Block Base Address 0x0_8200 DMA 3—Block Base Address 0x0_8280 | | | | |
|--|--|--------|-------------|-----------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x030 | OMISR—Outbound message interrupt status register | Mixed | 0x0000_0000 | 12.4.1/12-4 |
| 0x034 | OMIMR—Outbound message interrupt mask register | R/W | 0x0000_0000 | 12.4.2/12-5 |
| 0x050 | IMR0—Inbound message register 0 | R/W | 0x0000_0000 | 12.4.3/12-6 |
| 0x054 | IMR1—Inbound message register 1 | R/W | 0x0000_0000 | 12.4.3/12-6 |
| 0x058 | OMR0—Outbound message register 0 | R/W | 0x0000_0000 | 12.4.4/12-6 |
| 0x05C | OMR1—Outbound message register 1 | R/W | 0x0000_0000 | 12.4.4/12-6 |
| 0x060 | ODR—Outbound doorbell register | R/W | 0x0000_0000 | 12.4.5/12-6 |
| 0x068 | IDR—Inbound doorbell register | R/W | 0x0000_0000 | 12.4.5/12-6 |
| 0x080 | IMISR—Inbound message interrupt status register | Mixed | 0x0000_0000 | 12.4.6/12-8 |

Table A-18. DMA Controller Registers (continued)

| DMA 0—Block Base Address 0x0_8100 DMA 1—Block Base Address 0x0_8180 DMA 2—Block Base Address 0x0_8200 DMA 3—Block Base Address 0x0_8280 | | | | |
|--|--|--------|-------------|--------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x084 | IMIMR—Inbound message interrupt mask register | R/W | 0x0000_0000 | 12.4.7/12-9 |
| 0x100 | DMAMR0—DMA 0 mode register | R/W | 0x0000_0000 | 12.4.8.1/12-10 |
| 0x104 | DMASR0—DMA 0 status register | R/W | 0x0000_0000 | 12.4.8.2/12-12 |
| 0x108 | DMACDAR0—DMA 0 current descriptor address register | R/W | 0x0000_0000 | 12.4.8.3/12-13 |
| 0x110 | DMASAR0—DMA 0 source address register | R/W | 0x0000_0000 | 12.4.8.4/12-14 |
| 0x118 | DMADAR0—DMA 0 destination address register | R/W | 0x0000_0000 | 12.4.8.5/12-14 |
| 0x120 | DMABCR0—DMA 0 byte count register | R/W | 0x0000_0000 | 12.4.8.6/12-15 |
| 0x124 | DMANDAR0—DMA 0 next descriptor address register | R/W | 0x0000_0000 | 12.4.8.7/12-15 |
| 0x180 | DMAMR1—DMA 1 mode register | R/W | 0x0000_0000 | 12.4.8.1/12-10 |
| 0x184 | DMASR1—DMA 1 status register | R/W | 0x0000_0000 | 12.4.8.2/12-12 |
| 0x188 | DMACDAR1—DMA 1 current descriptor address register | R/W | 0x0000_0000 | 12.4.8.3/12-13 |
| 0x190 | DMASAR1—DMA 1 source address register | R/W | 0x0000_0000 | 12.4.8.4/12-14 |
| 0x198 | DMADAR1—DMA 1 destination address register | R/W | 0x0000_0000 | 12.4.8.5/12-14 |
| 0x1A0 | DMABCR1—DMA 1 byte count register | R/W | 0x0000_0000 | 12.4.8.6/12-15 |
| 0x1A4 | DMANDAR1—DMA 1 next descriptor address register | R/W | 0x0000_0000 | 12.4.8.7/12-15 |
| 0x200 | DMAMR2—DMA 2 mode register | R/W | 0x0000_0000 | 12.4.8.1/12-10 |
| 0x204 | DMASR2—DMA 2 status register | R/W | 0x0000_0000 | 12.4.8.2/12-12 |
| 0x208 | DMACDAR2—DMA 2 current descriptor address register | R/W | 0x0000_0000 | 12.4.8.3/12-13 |
| 0x210 | DMASAR2—DMA 2 source address register | R/W | 0x0000_0000 | 12.4.8.4/12-14 |
| 0x218 | DMADAR2—DMA 2 destination address register | R/W | 0x0000_0000 | 12.4.8.5/12-14 |
| 0x220 | DMABCR2—DMA 2 byte count register | R/W | 0x0000_0000 | 12.4.8.6/12-15 |
| 0x224 | DMANDAR2—DMA 2 next descriptor address register | R/W | 0x0000_0000 | 12.4.8.7/12-15 |
| 0x280 | DMAMR3—DMA 3 mode register | R/W | 0x0000_0000 | 12.4.8.1/12-10 |
| 0x284 | DMASR3—DMA 3 status register | R/W | 0x0000_0000 | 12.4.8.2/12-12 |
| 0x288 | DMACDAR3—DMA 3 current descriptor address register | R/W | 0x0000_0000 | 12.4.8.3/12-13 |
| 0x290 | DMASAR3—DMA 3 source address register | R/W | 0x0000_0000 | 12.4.8.4/12-14 |
| 0x298 | DMADAR3—DMA 3 destination address register | R/W | 0x0000_0000 | 12.4.8.5/12-14 |
| 0x2A0 | DMABCR3—DMA 3 byte count register | R/W | 0x0000_0000 | 12.4.8.6/12-15 |
| 0x2A4 | DMANDAR3—DMA 3 next descriptor address register | R/W | 0x0000_0000 | 12.4.8.7/12-15 |

Table A-18. DMA Controller Registers (continued)

| DMA 0—Block Base Address 0x0_8100 DMA 1—Block Base Address 0x0_8180 DMA 2—Block Base Address 0x0_8200 DMA 3—Block Base Address 0x0_8280 | | | | |
|--|------------------------------------|--------|-------------|--------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x2A8 | DMAGSR—DMA general status register | R | 0x0000_0000 | 12.4.8.8/12-16 |
| 0x2B0– 0x2FF | Reserved | — | — | — |

A.19 PCI Configuration Access

Table A-19. PCI Configuration Access Registers

| PCI Configuration Access—Block Base Address 0x0_8300 | | | | |
|--|--------------------|--------|-------------|--------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x0 | PCI_CONFIG_ADDRESS | W | 0x0000_0000 | 13.3.1.1/13-13 |
| 0x4 | PCI_CONFIG_DATA | R/W | 0x0000_0000 | 13.3.1.2/13-14 |
| 0x8 | PCI_INT_ACK | R | | 13.3.1.3/13-15 |

A.20 I/O Sequencer (IOS)

Table A-21. I/O Sequencer (IOS) Registers

| I/O Sequencer (IOS)—Block Base Address 0x0_8400 | | | | |
|---|--|--------|-------------|-----------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x00 | POTAR0—PCI outbound translation address register 0 | R/W | 0x0000_0000 | 11.4.1/11-3 |
| 0x08 | POBAR0—PCI outbound base address register 0 | R/W | 0x0000_0000 | 11.4.2/11-3 |
| 0x10 | POCMR0—PCI outbound comparison mask register 0 | R/W | 0x0000_0000 | 11.4.3/11-4 |
| 0x18 | POTAR1—PCI outbound translation address register 1 | R/W | 0x0000_0000 | 11.4.1/11-3 |
| 0x20 | POBAR1—PCI outbound base address register 1 | R/W | 0x0000_0000 | 11.4.2/11-3 |
| 0x28 | POCMR1—PCI outbound comparison mask register 1 | R/W | 0x0000_0000 | 11.4.3/11-4 |
| 0x30 | POTAR2—PCI outbound translation address register 2 | R/W | 0x0000_0000 | 11.4.1/11-3 |
| 0x38 | POBAR2—PCI outbound base address register 2 | R/W | 0x0000_0000 | 11.4.2/11-3 |
| 0x40 | POCMR2—PCI outbound comparison mask register 2 | R/W | 0x0000_0000 | 11.4.3/11-4 |
| 0x48 | POTAR3—PCI outbound translation address register 3 | R/W | 0x0000_0000 | 11.4.1/11-3 |
| 0x50 | POBAR3—PCI outbound base address register 3 | R/W | 0x0000_0000 | 11.4.2/11-3 |
| 0x58 | POCMR3—PCI outbound comparison mask register 3 | R/W | 0x0000_0000 | 11.4.3/11-4 |

Table A-21. I/O Sequencer (IOS) Registers (continued)

| I/O Sequencer (IOS)—Block Base Address 0x0_8400 | | | | |
|---|--|--------|-------------|-----------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x60 | POTAR4—PCI outbound translation address register 4 | R/W | 0x0000_0000 | 11.4.1/11-3 |
| 0x68 | POBAR4—PCI outbound base address register 4 | R/W | 0x0000_0000 | 11.4.2/11-3 |
| 0x70 | POCMR4—PCI outbound comparison mask register 4 | R/W | 0x0000_0000 | 11.4.3/11-4 |
| 0x78 | POTAR5—PCI outbound translation address register 5 | R/W | 0x0000_0000 | 11.4.1/11-3 |
| 0x80 | POBAR5—PCI outbound base address register 5 | R/W | 0x0000_0000 | 11.4.2/11-3 |
| 0x88 | POCMR5—PCI outbound comparison mask register 5 | R/W | 0x0000_0000 | 11.4.3/11-4 |
| 0xF0 | PMCR—Power management control register | R/W | 0x0000_0000 | 11.4.4/11-5 |
| 0xF8 | DTCR—Discard timer control register | R/W | 0x0000_0000 | 11.4.5/11-6 |

A.21 PCI Controller

Table A-22. PCI Controller Registers

| PCI Controller—Block Base Address 0x0_8500 | | | | |
|--|---|--------|-------------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| PCI Error Management Registers | | | | |
| 0x00 | PCI error status register (PCI_ESR) | w1c | 0x0000_0000 | 13.3.2.1/13-15 |
| 0x04 | PCI error capture disable register (PCI_ECDR) | R/W | 0x0000_0000 | 13.3.2.2/13-16 |
| 0x08 | PCI error enable register (PCI_EER) | R/W | 0x0000_0000 | 13.3.2.3/13-17 |
| 0x0C | PCI error attributes capture register (PCI_EATCR) | R/W | 0x0000_0000 | 13.3.2.4/13-18 |
| 0x10 | PCI error address capture register (PCI_EACR) | R | 0x0000_0000 | 13.3.2.5/13-19 |
| 0x14 | PCI error extended address capture register (PCI_EEACR) | R | 0x0000_0000 | 13.3.2.6/13-20 |
| 0x18 | PCI error data capture register (PCI_EDCR) | R/W | 0x0000_0000 | 13.3.2.7/13-20 |
| PCI Control and Status Registers | | | | |
| 0x20 | PCI general control register (PCI_GCR) | R/W | 0x0000_0000 | 13.3.2.8/13-20 |
| 0x24 | PCI error control register (PCI_ECR) | R/W | 0x0000_0000 | 13.3.2.9/13-21 |
| 0x28 | PCI general status register (PCI_GSR) | R | 0x0000_0000 | 13.3.2.10/13-22 |
| PCI Inbound ATU Registers | | | | |
| 0x38 | PCI inbound translation address register 2 (PITAR2) | R/W | 0x0000_0000 | 13.3.2.11/13-22 |
| 0x3C | Reserved | — | — | — |
| 0x40 | PCI inbound base address register 2 (PIBAR2) | R/W | 0x0000_0000 | 13.3.2.12/13-23 |
| 0x44 | PCI inbound extended base address register 2 (PIEBAR2) | R/W | 0x0000_0000 | 13.3.2.13/13-24 |
| 0x48 | PCI inbound window attributes register 2 (PIWAR2) | R/W | 0x0000_0000 | 13.3.2.14/13-24 |

Table A-22. PCI Controller Registers (continued)

| PCI Controller—Block Base Address 0x0_8500 | | | | |
|--|--|--------|-------------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x50 | PCI inbound translation address register 1 (PITAR1) | R/W | 0x0000_0000 | 13.3.2.11/13-22 |
| 0x54 | Reserved | — | — | — |
| 0x58 | PCI inbound base address register 1 (PIBAR1) | R/W | 0x0000_0000 | 13.3.2.12/13-23 |
| 0x5C | PCI inbound extended base address register 1 (PIEBAR1) | R/W | 0x0000_0000 | 13.3.2.13/13-24 |
| 0x60 | PCI inbound window attributes register 1 (PIWAR1) | R/W | 0x0000_0000 | 13.3.2.14/13-24 |
| 0x68 | PCI inbound translation address register 0 (PITAR0) | R/W | 0x0000_0000 | 13.3.2.11/13-22 |
| 0x6C | Reserved | — | — | — |
| 0x70 | PCI inbound base address register 0 (PIBAR0) | R/W | 0x0000_0000 | 13.3.2.12/13-23 |
| 0x78 | PCI inbound window attributes register 0 (PIWAR0) | R/W | 0x0000_0000 | 13.3.2.13/13-24 |
| 0x7C– 0xFF | Reserved | — | — | — |

A.22 PCI Express Controller

Table A-23. PCI Express Controller Registers

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|--------------------------------------|--------|--|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| PCI Express Controller 1 Registers | | | | |
| PCI Express1 Core Configuration Header Registers | | | | |
| 0x000 | PCI Express Vendor ID Register | R | 0x1957 | 14.4.1.1/14-15 |
| 0x002 | PCI Express Device ID Register | R | Device-specific | 14.4.1.2/14-16 |
| 0x004 | PCI Express Command Register | Mixed | All zeros | 14.4.1.3/14-16 |
| 0x006 | PCI Express Status Register | Mixed | 0x0010 | 14.4.1.4/14-17 |
| 0x008 | PCI Express Revision ID Register | R | Revision-specific | 14.4.1.5/14-18 |
| 0x009 | PCI Express Class Code Register | Mixed | 0x0B20 | 14.4.1.6/14-19 |
| 0x00C | PCI Express Cache Line Size Register | R/W | All zeros | 14.4.1.7/14-19 |
| 0x00D | PCI Express Latency Timer Register | R | All zeros | 14.4.1.8/14-20 |
| 0x00E | PCI Express Header Type Register | R | 0x0000 (EP mode) 0x0001 (RC mode) | 14.4.1.9/14-20 |
| 0x00F | PCI Express BIST Register | R | All zeros | 14.4.1.10/14-21 |

Table A-23. PCI Express Controller Registers (continued)

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|---|---------|-------------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x010–0x014 | Base Address Registers 0 and 1 (BAR0/BAR1) (EP mode only) | Mixed | 0x0008 | 14.4.2.1.1/14-22 |
| 0x018–0x020 | Base Address Registers 2 and 4 (BAR2/BAR4) (EP mode only) | Mixed | 0x0000_000C | 14.4.2.1.2/14-23 |
| 0x01C–0x024 | Base Address Registers 3 and 5 (BAR3/BAR5) (EP mode only) | R/W | All zeros | 14.4.2.1.3/14-23 |
| 0x02C | PCI Express Subsystem Vendor ID Register (EP mode only) | Special | All zeros | 14.4.2.2/14-24 |
| 0x02E | PCI Express Subsystem ID Register (EP mode only) | Special | All zeros | 14.4.2.3/14-24 |
| 0x034 | PCI Express Capabilities Pointer Register | R | 0x0044 | 14.4.2.4/14-25 |
| 0x03C | PCI Express Interrupt Line Register (EP mode only) | R/W | All zeros | 14.4.2.5/14-25 |
| 0x03D | PCI Express Interrupt Pin Register | R | 0x0001 | 14.4.2.6/14-26 |
| 0x03E | PCI Express Minimum Grant Register (EP mode only) | R | All zeros | 14.4.2.7/14-26 |
| 0x03F | PCI Express Maximum Latency Register (EP mode only) | R | All zeros | 14.4.2.8/14-26 |
| 0x018 | PCI Express Primary Bus Number Register (RC mode only) | R/W | All zeros | 14.4.3.1/14-27 |
| 0x019 | PCI Express Secondary Bus Number Register (RC mode only) | R/W | All zeros | 14.4.3.2/14-28 |
| 0x01A | PCI Express Subordinate Bus Number Register (RC mode only) | R/W | All zeros | 14.4.3.3/14-28 |
| 0x01B | Secondary Latency Timer Register 2 (RC mode only) | | All zeros | |
| 0x01C | PCI Express I/O Base Register (RC mode only) | R | All zeros | 14.4.3.5/14-29 |
| 0x01D | PCI Express I/O Limit Register (RC mode only) | R | All zeros | 14.4.3.6/14-29 |
| 0x01E | PCI Express Secondary Status Register (RC mode only) | Mixed | All zeros | 14.4.3.7/14-30 |
| 0x020 | PCI Express Memory Base Register (RC mode only) | R/W | All zeros | 14.4.3.8/14-31 |
| 0x022 | PCI Express Memory Limit Register (RC mode only) | R/W | All zeros | 14.4.3.9/14-31 |
| 0x024 | PCI Express Prefetchable Memory Base Register (RC mode only) | R/W | All zeros | 14.4.3.10/14-32 |
| 0x026 | PCI Express Prefetchable Memory Limit Register (RC mode only) | R/W | All zeros | 14.4.3.11/14-32 |
| 0x028 | PCI Express Prefetchable Base Upper 32-Bit Register (RC mode only) | R/W | All zeros | 14.4.3.12/14-33 |
| 0x02C | PCI Express Prefetchable Limit Upper 32-Bit Register (RC mode only) | R/W | All zeros | 14.4.3.13/14-33 |
| 0x030 | PCI Express I/O Base Upper 16-Bit Register (RC mode only) | R | All zeros | 14.4.3.14/14-33 |
| 0x032 | PCI Express I/O Limit Upper 16-Bit Register (RC mode only) | R' | All zeros | 14.4.3.15/14-34 |
| 0x034 | PCI Express Capabilities Pointer Register | R | 0x044 | 14.4.3.16/14-34 |
| 0x03C | PCI Express Interrupt Line Register | R/W | All zeros | 14.4.3.17/14-35 |

Table A-23. PCI Express Controller Registers (continued)

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|---|--------|-------------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x03D | PCI Express Interrupt Pin Register | R | 0x0001 | 14.4.3.18/14-35 |
| 0x03E | PCI Express Bridge Control Register (RC mode only) | R/W | All zeros | 14.4.3.19/14-35 |
| 0x044 | PCI Express Power Management Capability ID Register | R | 0x01 | 14.4.4.1/14-38 |
| 0x045 | PCI Express Power Management Next Capabilities Pointer Register | R | 0x4C | 14.4.4.2/14-38 |
| 0x046 | PCI Express Power Management Capabilities Register | R | 0x7E02 | 14.4.4.3/14-39 |
| 0x048 | PCI Express Power Management Status and Control Register | Mixed | All zeros | 14.4.4.4/14-39 |
| 0x04B | PCI Express Power Management Data Register | R | All zeros | 14.4.4.5/14-40 |
| 0x04C | PCI Express Capability ID Register | R | 0x10 | 14.4.4.6/14-40 |
| 0x04D | PCI Express Next Capabilities Pointer Register | R | 0x70 | 14.4.4.7/14-41 |
| 0x04E | PCI Express Capabilities Register | R | 0x00n1 | 14.4.4.8/14-41 |
| 0x050 | PCI Express Device Capabilities Register | R | 0x003F_C001 | 14.4.4.9/14-42 |
| 0x054 | PCI Express Device Control Register | R/W | 0x2810 | 14.4.4.10/14-43 |
| 0x056 | PCI Express Device Status Register | Mixed | All zeros | 14.4.4.11/14-44 |
| 0x058 | PCI Express Link Capabilities Register | R | 0x0003_B481 | 14.4.4.12/14-44 |
| 0x05C | PCI Express Link Control Register | R/W | All zeros | 14.4.4.13/14-45 |
| 0x05E | PCI Express Link Status Register | R | 0x0011 | 14.4.4.14/14-46 |
| 0x060 | PCI Express Slot Capabilities Register | R | All zeros | 14.4.4.15/14-46 |
| 0x064 | PCI Express Slot Control Register | R/W | 0x0000 | 14.4.4.16/14-47 |
| 0x066 | PCI Express Slot Status Register | Mixed | 0x0040 | 14.4.4.17/14-48 |
| 0x068 | PCI Express Root Control Register (RC mode only) | R/W | 0x0000 | 14.4.4.18/14-48 |
| 0x06C | PCI Express Root Status Register (RC mode only) | Mixed | All zeros | 14.4.4.19/14-49 |
| 0x070 | PCI Express MSI Message Capability ID Register (EP mode only) | R | 0x05 | 14.4.4.20/14-49 |
| 0x072 | PCI Express MSI Message Control Register (EP mode only) | Mixed | 0x0088 | 14.4.4.21/14-50 |
| 0x074 | PCI Express MSI Message Address Register (EP mode only) | R/W | All zeros | 14.4.4.22/14-50 |
| 0x078 | PCI Express MSI Message Upper Address Register (EP mode only) | R/W | All zeros | 14.4.4.23/14-51 |
| 0x07C | PCI Express MSI Message Data Register (EP mode only) | R/W | 0x0000 | 14.4.4.24/14-51 |
| 0x100 | PCI Express Advanced Error Reporting Capability ID Register | R | 0x0001 | 14.4.5.1/14-53 |
| 0x104 | PCI Express Uncorrectable Error Status Register | R/W | All zeros | 14.4.5.2/14-53 |
| 0x108 | PCI Express Uncorrectable Error Mask Register | R/W | All zeros | 14.4.5.3/14-54 |

Table A-23. PCI Express Controller Registers (continued)

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|--|--------|-------------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x10C | PCI Express Uncorrectable Error Severity Register | R/W | 0x0006_2010 | 14.4.5.4/14-55 |
| 0x110 | PCI Express Correctable Error Status Register | w1c | All zeros | 14.4.5.5/14-56 |
| 0x114 | PCI Express Correctable Error Mask Register | R/W | All zeros | 14.4.5.6/14-57 |
| 0x118 | PCI Express Advanced Error Capabilities and Control Register | R/W | 0x0000_00A0 | 14.4.5.7/14-57 |
| 0x11C | PCI Express Header Log Register | R | All zeros | 14.4.5.8/14-58 |
| 0x120 | PCI Express Header Log Register | R | All zeros | |
| 0x124 | PCI Express Header Log Register | R | All zeros | |
| 0x128 | PCI Express Header Log Register | R | All zeros | |
| 0x12C | PCI Express Root Error Command Register | R/W | All zeros | 14.4.5.9/14-59 |
| 0x130 | PCI Express Root Error Status Register | Mixed | All zeros | 14.4.5.10/14-59 |
| 0x134 | PCI Express Error Source Identification Register | R | All zeros | 14.4.5.11/14-60 |
| PCI Express Core Control and Status Registers (CSRs) | | | | |
| 0x404 | PCI Express LTSSM State Status Register (PEX_LTSSM_STAT) | R | All zeros | 14.4.6.1/14-61 |
| 0x41C | PCI Express N_FTS Control Register (PEX_NFTS_CTRL) | R/W | 0x0000_4040 | 14.4.6.2/14-62 |
| 0x438 | PCI Express ACK Replay Timeout Register (PEX_ACKRPLY_TO) | R/W | 0x005B_2090 | 14.4.6.3/14-63 |
| 0x440 | PCI Express IP Block Core Clock Ratio Register (PEX_GCLK_RATIO) | Mixed | 0x0000_0010 | 14.4.6.4/14-64 |
| 0x450 | PCI Express Power Management Timer Register (PEX_PM_TIMER) | Mixed | 0x000A_63E4 | 14.4.6.5/14-65 |
| 0x454 | PCI Express PME Time-Out Register (PEX_PME_TIMEOUT) | Mixed | 0x00FD_4BC0 | 14.4.6.6/14-66 |
| 0x45C | PCI Express ASPM Request Timer Register (PEX_ASPM_REQTMR) (RC mode only) | R/W | 0x0000_0629 | 14.4.6.7/14-66 |
| 0x478 | PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE) | R/W | All zeros | 14.4.6.8/14-67 |
| 0x47C | PCI Express Device Capabilities Update Register (PEX_DEVCAP_UPDATE) | R/W | All zeros | 14.4.6.9/14-68 |
| 0x480 | PCI Express Link Capabilities Update Register (PEX_LINKCAP_UPDATE) | R/W | 0x0000_3D41 | 14.4.6.10/14-69 |
| 0x490 | PCI Express Slot Capabilities Update Register (PEX_SLCAP_UPDATE) | R/W | 0x0000_07C0 | 14.4.6.11/14-70 |
| 0x4B0 | PCI Express Configuration Ready Register (PEX_CFG_READY) | Mixed | All zeros | 14.4.6.12/14-71 |
| PCI Express BAR Configuration Registers (EP Mode) | | | | |

Table A-23. PCI Express Controller Registers (continued)

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|--|--------|-------------|--------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x4D4 | PCI Express BAR Enable Register (PEX_BAR_ENABLE) | R/W | 0x0000_000F | 14.4.7.1/14-71 |
| 0x4D8 | PCI Express BAR Size Low Configuration Register (PEX_BAR_SIZE_L) | R/W | 0xFC00_0000 | 14.4.7.2/14-72 |
| 0x4DC | Reserved | R/W | 0xFFFF_FFFF | 14.4.7.3/14-73 |
| 0x4E0 | PCI Express BAR Select Configuration Register (PEX_BAR_SEL) | R/W | All zeros | 14.4.7.3/14-73 |
| 0x504 | PCI Express BAR Prefetch Configuration Register (PEX_BAR_PF) | R/W | All zeros | 14.4.7.4/14-73 |
| PCI Express Extended Status and Control Register | | | | |
| 0x590 | PCI Express PME_To_Ack Timeout Register (PEX_PME_TO_ACK_TOR) | R/W | 0x00FD_4BC0 | 14.4.8.1/14-74 |
| 0x594 | PME_To_Ack Status Register (PEX_PME_TO_ACK_SR) | w1c | All zeros | 14.4.8.2/14-75 |
| 0x5A0 | PCI Express PCI Interrupt Mask Register (PEX_SS_INTR_MASK) | R/W | 0x0000_003F | 14.4.8.3/14-76 |
| PCI Express CSB Bridge Registers | | | | |
| Global Registers | | | | |
| 0x800 | Reserved | RO | 0x0110_1010 | — |
| 0x804 | Reserved | RO | 0x0003_249F | — |
| 0x808 | PCI Express CSB Bridge Control register (PEX_CSB_CTRL) | R/W | 0x0000_0130 | 14.5.2.1/14-77 |
| 0x80C | Reserved | RO | All zeros | — |
| 0x814 | PCI Express DMA Descriptor Timer Register (PEX_DMA_DSTMR) | R/W | All zeros | 14.5.2.2/14-78 |
| 0x818 | Reserved | RO | All zeros | — |
| 0x81C | PCI Express CSB Bridge Status register (PEX_CSB_STAT) | RO | All zeros | 14.5.2.3/14-78 |
| 0x820 | Reserved | RO | All zeros | — |
| PCI Express Outbound PIO Registers | | | | |
| 0x840 | PCI Express Outbound PIO Control Register (PEX_CSB_OBCTRL) | R/W | All zeros | 14.5.3.1/14-79 |
| 0x844 | PCI Express Outbound PIO Status Register (PEX_CSB_OBSTAT) | w1c | All zeros | 14.5.3.2/14-80 |
| 0x848 | Reserved | RO | All zeros | — |
| PCI Express Inbound PIO Registers | | | | |
| 0x8E0 | PCI Express Inbound PIO Control Register (PEX_CSB_IBCTRL) | R/W | All zeros | 14.5.4.1/14-81 |

Table A-23. PCI Express Controller Registers (continued)

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|--|--------|-----------|--------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x8E4 | PCI Express Inbound PIO Status Register (PEX_CSB_IBSTAT) | w1c | All zeros | 14.5.4.2/14-81 |
| 0x8E8 | Reserved | RO | All zeros | — |
| PCI Express DMA Registers | | | | |
| 0x990 | Reserved | RO | All zeros | — |
| 0x9A0 | PCI Express Write DMA Control Register (PEX_WDMA_CTRL) | R/W | All zeros | 14.5.5.1/14-82 |
| 0x9A4 | PCI Express Write DMA first Address Register (PEX_WDMA_ADDR) | R/W | All zeros | 14.5.5.2/14-83 |
| 0x9A8 | PCI Express Write DMA Status Register (PEX_WDMA_STAT) | w1c | All zeros | 14.5.5.3/14-83 |
| 0x9AC | Reserved | RO | All zeros | — |
| 0xA40 | PCI Express Read DMA Control Register (PEX_RDMA_CTRL) | R/W | All zeros | 14.5.5.4/14-84 |
| 0xA44 | PCI Express Read DMA first Address Register (PEX_RDMA_ADDR) | R/W | All zeros | 14.5.5.5/14-85 |
| 0xA48 | PCI Express Read DMA Status Register (PEX_RDMA_STAT) | w1c | All zeros | 14.5.5.6/14-85 |
| Mailbox Registers | | | | |
| 0xB20 | PCI Express Outbound Mailbox Control Register (PEX_OMBCR) | R/W | All zeros | 14.5.6.1/14-86 |
| 0xB24 | PCI Express Outbound Mailbox Data Register (PEX_OMBDR) | R/W | All zeros | 14.5.6.2/14-87 |
| 0xB60 | PCI Express Inbound Mailbox Control Register (PEX_IMBCR) | R/W | All zeros | 14.5.6.3/14-87 |
| 0xB64 | PCI Express Inbound Mailbox Data Register (PEX_IMBDR) | R/W | All zeros | 14.5.6.4/14-88 |
| PCI Express Host Interrupts Registers | | | | |
| 0xBA0 | PCI Express Host Interrupt Enable Register (PEX_HIER) | R/W | All zeros | 14.5.7.1/14-89 |
| 0xBA4 | PCI Express Host Interrupt Status Register (PEX_HISR) | w1c | All zeros | 14.5.7.2/14-90 |
| 0xBA8 | PCI Express Host Outbound PIO Interrupt Vector Register (PEX_HOPIVR) | R/W | All zeros | 14.5.7.3/14-91 |
| 0xBC0 | PCI Express Host Inbound PIO Interrupt Vector Register (PEX_HIPIVR) | R/W | All zeros | 14.5.7.4/14-91 |
| 0xBC8 | PCI Express Host Write DMA Interrupt Vector Register (PEX_HWDIVR) | R/W | All zeros | 14.5.7.5/14-92 |
| 0xBD0 | PCI Express Host Read DMA Interrupt Vector Register (PEX_HRDIVR) | R/W | All zeros | 14.5.7.6/14-92 |
| 0xBD8 | PCI Express Host Miscellaneous Interrupt Vector Register (PEX_HMIVR) | R/W | All zeros | 14.5.7.7/14-93 |

Table A-23. PCI Express Controller Registers (continued)

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|---|--------|-------------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| CSB System Interrupts Registers | | | | |
| 0xBE0 | CSB System PIO Interrupt Enable Register (PEX_CSPIER) | R/W | All zeros | 14.5.8.1/14-93 |
| 0xBE4 | CSB System Write DMA Interrupt Enable Register (PEX_CSWDIER) | R/W | All zeros | 14.5.8.2/14-94 |
| 0xBE8 | CSB System Read DMA Interrupt Enable Register (PEX_CSRDIER) | R/W | All zeros | 14.5.8.3/14-95 |
| 0xBEC | CSB System Miscellaneous Interrupt Enable Register (PEX_CSMIER) | R/W | 0x0000_0002 | 14.5.8.4/14-95 |
| 0xBF0 | CSB System PIO Interrupt Status Register (PEX_CSPISR) | w1c | All zeros | 14.5.8.5/14-97 |
| 0xBF4 | CSB System Write DMA Interrupt Status Register (PEX_CSWDISR) | w1c | All zeros | 14.5.8.6/14-97 |
| 0xBF8 | CSB System Read DMA Interrupt Status Register (PEX_CSRDISR) | w1c | All zeros | 14.5.8.7/14-98 |
| 0xBFC | CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR) | w1c | All zeros | 14.5.8.8/14-99 |
| Power Management Registers | | | | |
| 0xC80 | PCI Express PM Control Register (PEX_PM_CTRL) | R/W | All zeros | 14.5.9.1/14-100 |
| 0xC88 | PCI Express slot control misc register | R/W | All zeros | 14.5.9.2/14-101 |
| PCI Express Outbound Address Mapping Registers | | | | |
| 0xCA0 | PCI Express Outbound Window Attributes Register 0 (PEX_OWAR0) | R/W | All zeros | 14.5.10.1/14-102 |
| 0xCA4 | PCI Express Outbound Window Base Address Register 0 (PEX_OWBAR0) | R/W | All zeros | 14.5.10.2/14-103 |
| 0xCA8 | PCI Express Outbound Window Translation Address Register Low 0 (PEX_OWTARL0) | R/W | All zeros | 14.5.10.3/14-103 |
| 0xCAC | PCI Express Outbound Window Translation Address Register High 0 (PEX_OWTARH0) | R/W | All zeros | 14.5.10.4/14-104 |
| 0xCB0 | PCI Express Outbound Window Attributes Register 1 (PEX_OWAR1) | R/W | All zeros | 14.5.10.1/14-102 |
| 0xCB4 | PCI Express Outbound Window Base Address Register 1 (PEX_OWBAR1) | R/W | All zeros | 14.5.10.2/14-103 |
| 0xCB8 | PCI Express Outbound Window Translation Address Register Low 1 (PEX_OWTARL1) | R/W | All zeros | 14.5.10.3/14-103 |
| 0xCBC | PCI Express Outbound Window Translation Address Register High 1 (PEX_OWTARH1) | R/W | All zeros | 14.5.10.4/14-104 |
| 0xCC0 | PCI Express Outbound Window Attributes Register 2 (PEX_OWAR2) | R/W | All zeros | 14.5.10.1/14-102 |

Table A-23. PCI Express Controller Registers (continued)

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|--|--------|-----------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0xCC4 | PCI Express Outbound Window Base Address Register 2 (PEX_OWBAR2) | R/W | All zeros | 14.5.10.2/14-103 |
| 0xCC8 | PCI Express Outbound Window Translation Address Register Low 2 (PEX_OWTLARL2) | R/W | All zeros | 14.5.10.3/14-103 |
| 0xCCC | PCI Express Outbound Window Translation Address Register High 2 (PEX_OWTLARH2) | R/W | All zeros | 14.5.10.4/14-104 |
| 0xCD0 | PCI Express Outbound Window Attributes Register 3 (PEX_OWAR3) | R/W | All zeros | 14.5.10.1/14-102 |
| 0xCD4 | PCI Express Outbound Window Base Address Register 3 (PEX_OWBAR3) | R/W | All zeros | 14.5.10.2/14-103 |
| 0xCD8 | PCI Express Outbound Window Translation Address Register Low 3 (PEX_OWTLARL3) | R/W | All zeros | 14.5.10.3/14-103 |
| 0xCDC | PCI Express Outbound Window Translation Address Register High 3 (PEX_OWTLARH3) | R/W | All zeros | 14.5.10.4/14-104 |
| PCI Express EP Inbound Address Translation Registers | | | | |
| 0xDE0 | PCI Express EP Inbound Window Translation Address Register 0 (PEX_EPIWTAR0) | R/W | All zeros | 14.5.11.1/14-105 |
| 0xDE4 | PCI Express EP Inbound Window Translation Address Register 1 (PEX_EPIWTAR1) | R/W | All zeros | 14.5.11.1/14-105 |
| 0xDE8 | PCI Express EP Inbound Window Translation Address Register 2 (PEX_EPIWTAR2) | R/W | All zeros | 14.5.11.1/14-105 |
| 0xDEC | PCI Express EP Inbound Window Translation Address Register 3 (PEX_EPIWTAR3) | R/W | All zeros | 14.5.11.1/14-105 |
| PCI Express RC Inbound Address Mapping Registers | | | | |
| 0xE60 | PCI Express RC Inbound Window Attributes Register 0 (PEX_RCIWAR0) | R/W | All zeros | 14.5.12.1/14-106 |
| 0xE64 | PCI Express RC Inbound Window Translation Address Register 0 (PEX_RCIWTAR0) | R/W | All zeros | 14.5.12.2/14-107 |
| 0xE68 | PCI Express RC Inbound Window Base Address Register Low 0 (PEX_RCIWBARL0) | R/W | All zeros | 14.5.12.3/14-108 |
| 0xE6C | PCI Express RC Inbound Window Base Address Register High 0 (PEX_RCIWBARH0) | R/W | All zeros | 14.5.12.4/14-108 |
| 0xE70 | PCI Express RC Inbound Window Attributes Register 1 (PEX_RCIWAR1) | R/W | All zeros | 14.5.12.1/14-106 |
| 0xE74 | PCI Express RC Inbound Window Translation Address Register 1 (PEX_RCIWTAR1) | R/W | All zeros | 14.5.12.2/14-107 |
| 0xE78 | PCI Express RC Inbound Window Base Address Register Low 1 (PEX_RCIWBARL1) | R/W | All zeros | 14.5.12.3/14-108 |

Table A-23. PCI Express Controller Registers (continued)

| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
|--|--|--------|-----------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0xE7C | PCI Express RC Inbound Window Base Address Register High 1 (PEX_RCIWBARH1) | R/W | All zeros | 14.5.12.4/14-108 |
| 0xE80 | PCI Express RC Inbound Window Attributes Register 2 (PEX_RCIWAR2) | R/W | All zeros | 14.5.12.1/14-106 |
| 0xE84 | PCI Express RC Inbound Window Translation Address Register 2 (PEX_RCIWTAR2) | R/W | All zeros | 14.5.12.2/14-107 |
| 0xE88 | PCI Express RC Inbound Window Base Address Register Low 2 (PEX_RCIWBARL2) | R/W | All zeros | 14.5.12.3/14-108 |
| 0xE8C | PCI Express RC Inbound Window Base Address Register High 2 (PEX_RCIWBARH2) | R/W | All zeros | 14.5.12.4/14-108 |
| 0xE90 | PCI Express RC Inbound Window Attributes Register 3 (PEX_RCIWAR3) | R/W | All zeros | 14.5.12.1/14-106 |
| 0xE94 | PCI Express RC Inbound Window Translation Address Register 3 (PEX_RCIWTAR3) | R/W | All zeros | 14.5.12.2/14-107 |
| 0xE98 | PCI Express RC Inbound Window Base Address Register Low 3 (PEX_RCIWBARL3) | R/W | All zeros | 14.5.12.3/14-108 |
| 0xE9C | PCI Express RC Inbound Window Base Address Register High 3 (PEX_RCIWBARH3) | R/W | All zeros | 14.5.12.4/14-108 |
| PCI Express Controller 2 Memory-Mapped Registers | | | | |
| 0x000–0xFFC | PCI Express Controller 2 registers Note: All registers defined for PCI Express controller 1 are also defined for PCI Express controller 2; the offsets of the PCI Express controller 2 registers are the same except they have a different block base address of 0x0_A000. | | | |

A.23 Time Division Multiplexing (TDM) Interface

Table A-24. Time Division Multiplexing (TDM) Interface Registers

| TDM Configuration—Block Base Address 0x1_6000 TDM Data—Block Base Address 0x1_6100 TDM Clock Control—Block Base Address 0x1_6180 | | | | |
|--|--|--------|-------------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| TDM Configuration | | | | |
| 0x000 | TDM general interface register (TDMGIR) | R/W | 0x0000_0000 | 25.4.2.1.1/25-8 |
| 0x004 | TDM receive interface register (TDMRIR) | R/W | 0x0000_0000 | 25.4.2.1.2/25-9 |
| 0x008 | TDM transmit interface register (TDMTIR) | R/W | 0x0001_0000 | 25.4.2.1.3/25-12 |
| 0x00C | TDM receive frame parameters (TDMRFP) | R/W | 0x0000_0000 | 25.4.2.1.4/25-14 |
| 0x010 | TDM transmit frame parameters (TDMTFP) | R/W | 0x0000_0000 | 25.4.2.1.5/25-16 |

Table A-24. Time Division Multiplexing (TDM) Interface Registers (continued)

| TDM Configuration—Block Base Address 0x1_6000 TDM Data—Block Base Address 0x1_6100 TDM Clock Control—Block Base Address 0x1_6180 | | | | |
|--|---|--------|-------------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| TDM Channel, Interrupt, and Rx/Tx Enables | | | | |
| 0x020–0x02C | TDM receive channel enable registers (TDMRCEN0–3) | R/W | 0x0000_0000 | 25.4.2.2.1/25-17 |
| 0x040–0x04C | TDM transmit channel enable registers (TDMTCEN0–3) | R/W | 0x0000_0000 | 25.4.2.2.2/25-17 |
| 0x060–0x06C | TDM transmit channel mask registers (TDMTCMA0–3) | R/W | 0x0000_0000 | 25.4.2.2.3/25-18 |
| 0x080 | TDM receive control register (TDMRCR) | R/W | 0x0000_0000 | 25.4.2.2.4/25-19 |
| 0x084 | TDM transmit control register (TDMTCR) | R/W | 0x0000_0000 | 25.4.2.2.5/25-19 |
| 0x088 | TDM receive interrupt enable register (TDMRIER) | R/W | 0x0000_0000 | 25.4.2.2.6/25-20 |
| 0x08C | TDM transmit interrupt enable register (TDMTIER) | R/W | 0x0000_0000 | 25.4.2.2.7/25-21 |
| TDM Status | | | | |
| 0x0A0 | TDM receive event register (TDMRER) | R/W | 0x0000_0000 | 25.4.2.3.1/25-22 |
| 0x0A4 | TDM transmit event register (TDMTER) | R/W | 0x0000_0000 | 25.4.2.3.2/25-24 |
| 0x0A8 | TDM receive status register (TDMRSR) | RO | 0x0000_0000 | 25.4.2.3.3/25-26 |
| 0x0AC | TDM transmit status register (TDMTSR) | RO | 0x0000_0000 | 25.4.2.3.4/25-26 |
| TDM Data | | | | |
| 0x000 | TDM receive data registers (TDMRDREG) | RO | 0x0000_0000 | 25.4.4.1/2525-28 |
| 0x008 | TDM transmit data registers (TDMTDREG) | WO | 0x0000_0000 | 25.4.4.1/2525-29 |
| TDM Clock Control | | | | |
| 0x000 | TDM Division Value Rx Serial Clock Register (TDMCLK_DIV_VAL_RX) | R/W | 0x0000_0000 | 25.9.3.1/25-52 |
| 0x004 | TDM Division Value for Tx Serial Clock Register (TDMCLK_DIV_VAL_TX) | R/W | 0x0000_0000 | 25.9.3.2/25-53 |

A.24 Serial ATA (SATA) Controller

Table A-25. Serial ATA (SATA) Controller Registers

| SATA 1 Controller—Block Base Address 0x1_8000 SATA 2 Controller—Block Base Address 0x1_9000 | | | | |
|--|--|--------|-------------|---------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| SATA Command Registers | | | | |
| 0x000 | CQR—Command queue register | R/W | 0x0000_0000 | 15.3.2.1/15-5 |
| 0x008 | CAR—Command active register | R | 0x0000_0000 | 15.3.2.2/15-6 |
| 0x010 | CCR—Command completed register | w1c | 0x0000_0000 | 15.3.2.3/15-6 |
| 0x018 | CER—Command error register | w1c | 0x0000_0000 | 15.3.2.4/15-7 |
| 0x020 | DE—Device error register | w1c | 0x0000_0000 | 15.3.2.5/15-8 |
| 0x024 | CHBA—Command header base address | R/W | 0x0000_0000 | 15.3.2.6/15-8 |
| 0x028 | HStatus—Host status register | w1c | 0x2000_0000 | 15.3.2.7/15-9 |
| 0x02C | HControl—Host control register | Mixed | 0x0000_0100 | 15.3.2.8/15-12 |
| 0x030 | CQPMP—Port number queue register | R/W | 0x0000_0000 | 15.3.2.9/15-13 |
| 0x034 | SIG—Signature register | R | 0xFFFF_FFFF | 15.3.2.10/15-14 |
| 0x038 | ICC—Interrupt coalescing control register | R/W | 0x0100_0000 | 15.3.2.11/15-14 |
| SATA1 Superset Registers | | | | |
| 0x100 | SStatus—SATA interface status register | R | 0x0000_0000 | 15.3.3.1/15-15 |
| 0x104 | SError—SATA interface error register | w1c | 0x0000_0000 | 15.3.3.2/15-16 |
| 0x108 | SControl—SATA interface control register | R/W | 0x0000_0300 | 15.3.3.3/15-18 |
| 0x10C | SNotification—SATA interface notification register | w1c | 0x0000_0000 | 15.3.3.4/15-19 |
| SATA1 Control Status Registers | | | | |
| 0x140 | TransCfg—Transport layer configuration | R/W | 0x0800_0016 | 15.3.4.1/15-20 |
| 0x144 | TransStatus—Transport layer status | R | 0x0000_0000 | 15.3.4.2/15-21 |
| 0x148 | LinkCfg—Link layer configuration | R/W | 0x0000_FF34 | 15.3.4.3/15-21 |
| 0x14C | LinkCfg1—Link layer configuration1 | R/W | 0x0000_0000 | 15.3.4.4/15-22 |
| 0x150 | LinkCfg2—Link layer configuration2 | R/W | 0x0000_0000 | 15.3.4.5/15-23 |
| 0x154 | LinkStatus—Link layer status | R | 0x0000_0000 | 15.3.4.6/15-23 |
| 0x158 | LinkStatus1—Link layer status1 | R | 0x0000_0000 | 15.3.4.7/15-24 |
| 0x15C | PhyCtrlCfg1—PHY control configuration1 | R/W | 0x0000_3800 | 15.3.4.8/15-26 |
| 0x160 | CommandStatus—Link layer command status | R | 0x0000_0000 | 15.3.4.9/15-27 |
| 0x164–0x3FC | Reserved | — | — | — |
| 0x400 | PhyCtrlCfg2—PHY control configuration2 | R/W | 0x0000_8000 | 15.3.4.10/15-29 |

Table A-25. Serial ATA (SATA) Controller Registers (continued)

| SATA 1 Controller—Block Base Address 0x1_8000 SATA 2 Controller—Block Base Address 0x1_9000 | | | | |
|--|--------------------------------|--------|-------------|--------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| SATA1 System Control Registers | | | | |
| 0x410 | SYSPR—System priority register | R/W | 0x0000_0000 | 15.3.5.1/15-30 |
| 0x40C–0xFFFF | Reserved | — | — | — |
| SATA2—Block Base Address: 0x1_9000 | | | | |

Note: SATA2 has the same memory-mapped registers that are described for SATA1 from 0x1_8000 to 0x1_8FFF except the offsets are from 0x1_9000 to 0x1_9FFF.

A.25 USB DR Controller

Table A-26. USB DR Controller Registers

| USB DR Controller—Block Base Address 0x2_3000 | | | | |
|---|--|--------|-------------|-----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000–0x0FF | Reserved, should be cleared | — | — | — |
| 0x100 | CAPLENGTH—Capability register length | R | 0x40 | 17.3.1.1/1717-9 |
| 0x102 | HCIVERSION—Host interface version number ¹ | R | 0x0100 | 17.3.1.2/1717-9 |
| 0x104 | HCSPARAMS—Host ctrl. structural parameters ¹ | R | 0x1100_0100 | 17.3.1.3/1717-10 |
| 0x108 | HCCPARAMS—Host ctrl. capability parameters ¹ | R | 0x0000_0006 | 17.3.1.4/1717-10 |
| 0x120 | DCIVERSION—Device interface version number | R | 0x0001 | 17.3.1.5/1717-11 |
| 0x124 | DCCPARAMS—Device controller parameters | R | 0x8301_0000 | 17.3.1.6/1717-12 |
| 0x140 | USBCMD—USB command | Mixed | 0x0000_8000 | 17.3.2.1/1717-13 |
| 0x144 | USBSTS—USB status | Mixed | 0x0000_0000 | 17.3.2.2/1717-15 |
| 0x148 | USBINTR—USB interrupt enable | R/W | 0x0000_0000 | 17.3.2.3/1717-17 |
| 0x14C | FRINDEX—USB frame index | R/W | 0x0000_ffff | 17.3.2.4/1717-19 |
| 0x154 | PERIODICLISTBASE—Frame list base address ¹ | R/W | 0xffff_0000 | 17.3.2.6/1717-20 |
| | DEVICEADDR—USB device address | R/W | 0x0000_0000 | 17.3.2.7/1717-21 |
| 0x158 | ASYNCLISTADDR—Next asynchronous list addr (host mode) ¹ | R/W | 0x0000_0000 | 17.3.2.8/1717-22 |
| | ENDPOINTLISTADDR—Address at endpoint list (device mode) | R/W | 0x0000_0000 | 17.3.2.9/1717-22 |
| 0x160 | BURSTSIZE—Programmable burst size | R/W | 0x0000_1010 | 17.3.2.10/1717-23 |
| 0x164 | TXFILLTUNING—Host TT transmit pre-buffer packet tuning | R/W | 0x0000_0000 | 17.3.2.11/1717-24 |
| 0x170 | ULPI VIEWPORT—ULPI Register Access | Mixed | 0x0000_0000 | 17.3.2.12/1717-25 |

Table A-26. USB DR Controller Registers (continued)

| USB DR Controller—Block Base Address 0x2_3000 | | | | |
|---|---|--------|-------------|-----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x180 | CONFIGFLAG—Configured flag register | R | 0x0000_0001 | 17.3.2.13/1717-27 |
| 0x184 | PORTSC—Port status/control | Mixed | 0x0000_0010 | 17.3.2.14/1717-27 |
| 0x1A4 | OTGSC—On-The-Go status and control ¹ | Mixed | 0x200C_0000 | 17.3.2.15/1717-32 |
| 0x1A8 | USBMODE—USB device mode | R/W | 0x0000_0000 | 17.3.2.16/1717-35 |
| 0x1AC | ENDPTSETUPSTAT—Endpoint setup status | R/W | 0x0000_0000 | 17.3.2.17/1717-36 |
| 0x1B0 | ENDPOINTPRIME—Endpoint initialization | R/W | 0x0000_0000 | 17.3.2.18/1717-36 |
| 0x1B4 | ENDPTFLUSH—Endpoint de-initialize | R/W | 0x0000_0000 | 17.3.2.19/1717-37 |
| 0x1B8 | ENDPTSTATUS—Endpoint status | R | 0x0000_0000 | 17.3.2.20/1717-38 |
| 0x1BC | ENDPTCOMPLETE—Endpoint complete | w1c | 0x0000_0000 | 17.3.2.21/1717-38 |
| 0x1C0 | ENDPTCTRL0—Endpoint control 0 | Mixed | 0x0080_0080 | 17.3.2.22/1717-39 |
| 0x1C4 | ENDPTCTRL1—Endpoint control 1 | R/W | 0x0000_0000 | 17.3.2.23/1717-40 |
| 0x1C8 | ENDPTCTRL2—Endpoint control 2 | R/W | 0x0000_0000 | 17.3.2.23/1717-40 |
| 0x1CA– 0x1D4 | Reserved | — | — | — |
| 0x400 | SNOOP1—Snoop 1 | R/W | 0x0000_0000 | 17.3.2.24/1717-42 |
| 0x404 | SNOOP2—Snoop 2 | R/W | 0x0000_0000 | 17.3.2.24/1717-42 |
| 0x408 | AGE_CNT_THRESH—Age count threshold | R/W | 0x0000_0000 | 17.3.2.25/1717-43 |

¹ This register has separate functions for the host and device operation; the host function is listed first in the table.

A.26 Enhanced Three-Speed Ethernet Controllers (eTSECs)

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|----------------------------------|--------|-------------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| eTSEC General Control and Status Registers | | | | |
| 0x000 | TSEC_ID*—Controller ID register | R | 0x0124_0000 | 19.5.3.1.1/19-22 |
| 0x004 | TSEC_ID2*—Controller ID register | R | 0x0030_00F0 | 19.5.3.1.2/19-23 |
| 0x008– 0x00C | Reserved | — | — | — |
| 0x010 | IEVENT—Interrupt event register | w1c | 0x0000_0000 | 19.5.3.1.3/19-24 |
| 0x014 | IMASK—Interrupt mask register | R/W | 0x0000_0000 | 19.5.3.1.4/19-28 |
| 0x018 | EDIS—Error disabled register | R/W | 0x0000_0000 | 19.5.3.1.5/19-30 |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|---|--------|-------------|----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x01C | Reserved | — | — | — |
| 0x020 | ECNTRL—Ethernet control register | R/W | 0x0000_0000 | 19.5.3.1.6/19-31 |
| 0x024 | Reserved | — | — | — |
| 0x028 | PTV—Pause time value register | R/W | 0x0000_0000 | 19.5.3.1.7/19-33 |
| 0x02C | DMACTRL—DMA control register | R/W | 0x0000_0000 | 19.5.3.1.8/19-34 |
| 0x030 | Reserved | — | — | — |
| 0x034– 0x054 | Reserved | — | — | — |
| eTSEC Transmit Control and Status Registers | | | | |
| 0x100 | TCTRL—Transmit control register | R/W | 0x0000_0000 | 19.5.3.2.1/19-36 |
| 0x104 | TSTAT—Transmit status register | w1c | 0x0000_0000 | 19.5.3.2.2/19-38 |
| 0x108 | DFVLAN*—Default VLAN control word | R/W | 0x8100_0000 | 19.5.3.2.3/19-42 |
| 0x10C | Reserved | — | — | — |
| 0x110 | TXIC—Transmit interrupt coalescing register | R/W | 0x0000_0000 | 19.5.3.2.4/19-43 |
| 0x114 | TQUEUE*—Transmit queue control register | R/W | 0x0000_8000 | 19.5.3.2.5/19-44 |
| 0x118– 0x13C | Reserved | — | — | — |
| 0x140 | TR03WT*—TxBD Rings 0–3 round-robin weightings | R/W | 0x0000_0000 | 19.5.3.2.6/19-44 |
| 0x144 | TR47WT*—TxBD Rings 4–7 round-robin weightings | R/W | 0x0000_0000 | 19.5.3.2.7/19-45 |
| 0x148– 0x180 | Reserved | — | — | — |
| 0x184 | TBPTR0—TxBD pointer for ring 0 | R/W | 0x0000_0000 | 19.5.3.2.8/19-46 |
| 0x188 | Reserved | — | — | — |
| 0x18C | TBPTR1*—TxBD pointer for ring 1 | R/W | 0x0000_0000 | 19.5.3.2.8/19-46 |
| 0x190 | Reserved | — | — | — |
| 0x194 | TBPTR2*—TxBD pointer for ring 2 | R/W | 0x0000_0000 | 19.5.3.2.8/19-46 |
| 0x198 | Reserved | — | — | — |
| 0x19C | TBPTR3*—TxBD pointer for ring 3 | R/W | 0x0000_0000 | 19.5.3.2.8/19-46 |
| 0x1A0 | Reserved | — | — | — |
| 0x1A4 | TBPTR4*—TxBD pointer for ring 4 | R/W | 0x0000_0000 | 19.5.3.2.8/19-46 |
| 0x1A8 | Reserved | — | — | — |
| 0x1AC | TBPTR5*—TxBD pointer for ring 5 | R/W | 0x0000_0000 | 19.5.3.2.8/19-46 |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|--|--------|-------------|-----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x1B0 | Reserved | — | — | — |
| 0x1B4 | TBPTR6*—TxBD pointer for ring 6 | R/W | 0x0000_0000 | 19.5.3.2.8/19-46 |
| 0x1B8 | Reserved | — | — | — |
| 0x1BC | TBPTR7*—TxBD pointer for ring 7 | R/W | 0x0000_0000 | 19.5.3.2.8/19-46 |
| 0x1C0– 0x200 | Reserved | — | — | — |
| 0x204 | TBASE0—TxBD base address of ring 0 | R/W | 0x0000_0000 | 19.5.3.2.9/19-47 |
| 0x208 | Reserved | — | — | — |
| 0x20C | TBASE1*—TxBD base address of ring 1 | R/W | 0x0000_0000 | 19.5.3.2.9/19-47 |
| 0x210 | Reserved | — | — | — |
| 0x214 | TBASE2*—TxBD base address of ring 2 | R/W | 0x0000_0000 | 19.5.3.2.9/19-47 |
| 0x218 | Reserved | — | — | — |
| 0x21C | TBASE3*—TxBD base address of ring 3 | R/W | 0x0000_0000 | 19.5.3.2.9/19-47 |
| 0x220 | Reserved | — | — | — |
| 0x224 | TBASE4*—TxBD base address of ring 4 | R/W | 0x0000_0000 | 19.5.3.2.9/19-47 |
| 0x228 | Reserved | — | — | — |
| 0x22C | TBASE5*—TxBD base address of ring 5 | R/W | 0x0000_0000 | 19.5.3.2.9/19-47 |
| 0x230 | Reserved | — | — | — |
| 0x234 | TBASE6*—TxBD base address of ring 6 | R/W | 0x0000_0000 | 19.5.3.2.9/19-47 |
| 0x238 | Reserved | — | — | — |
| 0x23C | TBASE7*—TxBD base address of ring 7 | R/W | 0x0000_0000 | 19.5.3.2.9/19-47 |
| 0x240– 0x27C | Reserved | — | — | — |
| 0x280 | TMR_TXTS1_ID* - Tx time stamp identification tag (set 1) | R/W | 0x0000_0000 | 19.5.3.2.10/19-47 |
| 0x284 | TMR_TXTS2_ID* - Tx time stamp identification tag (set 2) | R/W | 0x0000_0000 | 19.5.3.2.10/19-47 |
| 0x288– 0x2BC | Reserved | — | — | — |
| 0x2C0 | TMR_TXTS1_H* - Tx time stamp high (set 1) | R/W | 0x0000_0000 | 19.5.3.2.11/19-48 |
| 0x2C4 | TMR_TXTS1_L* - Tx time stamp high (set 1) | R/W | 0x0000_0000 | 19.5.3.2.11/19-48 |
| 0x2C8 | TMR_TXTS2_H* - Tx time stamp high (set 2) | R/W | 0x0000_0000 | 19.5.3.2.11/19-48 |
| 0x2CC | TMR_TXTS2_L* - Tx time stamp high (set 2) | R/W | 0x0000_0000 | 19.5.3.2.11/19-48 |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|---|--------|-------------|-----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x2D0–0x2FC | Reserved | — | — | — |
| eTSEC Receive Control and Status Registers | | | | |
| 0x300 | RCTRL—Receive control register | R/W | 0x0000_0000 | 19.5.3.3.1/19-48 |
| 0x304 | RSTAT—Receive status register | w1c | 0x0000_0000 | 19.5.3.3.2/19-50 |
| 0x308–0x30C | Reserved | — | — | — |
| 0x310 | RXIC—Receive interrupt coalescing register | R/W | 0x0000_0000 | 19.5.3.3.3/19-52 |
| 0x314 | RQUEUE*—Receive queue control register. | R/W | 0x0080_0080 | 19.5.3.3.4/19-53 |
| 0x318–0x32C | Reserved | — | — | — |
| 0x330 | RBIFX*—Receive bit field extract control register | R/W | 0x0000_0000 | 19.5.3.3.5/19-54 |
| 0x334 | RQFAR*—Receive queue filing table address register | R/W | 0x0000_0000 | 19.5.3.3.6/19-56 |
| 0x338 | RQFCR*—Receive queue filing table control register | R/W | 0x0000_0000 | 19.5.3.3.7/19-56 |
| 0x33C | RQFPR*—Receive queue filing table property register | R/W | 0x0000_0000 | 19.5.3.3.8/19-57 |
| 0x340 | MRBLR—Maximum receive buffer length register | R/W | 0x0000_0000 | 19.5.3.3.9/19-61 |
| 0x344–0x380 | Reserved | — | — | — |
| 0x384 | RBPTR0—RxBd pointer for ring 0 | R/W | 0x0000_0000 | 19.5.3.3.10/19-61 |
| 0x388 | Reserved | — | — | — |
| 0x38C | RBPTR1*—RxBd pointer for ring 1 | R/W | 0x0000_0000 | 19.5.3.3.10/19-61 |
| 0x390 | Reserved | — | — | — |
| 0x394 | RBPTR2*—RxBd pointer for ring 2 | R/W | 0x0000_0000 | 19.5.3.3.10/19-61 |
| 0x398 | Reserved | — | — | — |
| 0x39C | RBPTR3*—RxBd pointer for ring 3 | R/W | 0x0000_0000 | 19.5.3.3.10/19-61 |
| 0x3A0 | Reserved | — | — | — |
| 0x3A4 | RBPTR4*—RxBd pointer for ring 4 | R/W | 0x0000_0000 | 19.5.3.3.10/19-61 |
| 0x3A8 | Reserved | — | — | — |
| 0x3AC | RBPTR5*—RxBd pointer for ring 5 | R/W | 0x0000_0000 | 19.5.3.3.10/19-61 |
| 0x3B0 | Reserved | — | — | — |
| 0x3B4 | RBPTR6*—RxBd pointer for ring 6 | R/W | 0x0000_0000 | 19.5.3.3.10/19-61 |
| 0x3B8 | Reserved | — | — | — |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|---|--------|-------------|-----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x3BC | RBPTR7*—RxBd pointer for ring 7 | R/W | 0x0000_0000 | 19.5.3.3.10/19-61 |
| 0x3C0– 0x400 | Reserved | — | — | — |
| 0x404 | RBASE0—RxBd base address of ring 0 | R/W | 0x0000_0000 | 19.5.3.3.11/19-62 |
| 0x408 | Reserved | — | — | — |
| 0x40C | RBASE1*—RxBd base address of ring 1 | R/W | 0x0000_0000 | 19.5.3.3.11/19-62 |
| 0x410 | Reserved | — | — | — |
| 0x414 | RBASE2*—RxBd base address of ring 2 | R/W | 0x0000_0000 | 19.5.3.3.11/19-62 |
| 0x418 | Reserved | — | — | — |
| 0x41C | RBASE3*—RxBd base address of ring 3 | R/W | 0x0000_0000 | 19.5.3.3.11/19-62 |
| 0x420 | Reserved | — | — | — |
| 0x424 | RBASE4*—RxBd base address of ring 4 | R/W | 0x0000_0000 | 19.5.3.3.11/19-62 |
| 0x428 | Reserved | — | — | — |
| 0x42C | RBASE5*—RxBd base address of ring 5 | R/W | 0x0000_0000 | 19.5.3.3.11/19-62 |
| 0x430 | Reserved | — | — | — |
| 0x434 | RBASE6*—RxBd base address of ring 6 | R/W | 0x0000_0000 | 19.5.3.3.11/19-62 |
| 0x438 | Reserved | — | — | — |
| 0x43C | RBASE7*—RxBd base address of ring 7 | R/W | 0x0000_0000 | 19.5.3.3.11/19-62 |
| 0x440– 0x4BC | Reserved | — | — | — |
| 0x4C0 | TMR_RXTS_H* - Rx timer time stamp register high | R/W | 0x0000_0000 | 19.5.3.3.12/19-62 |
| 0x4C4 | TMR_RXTS_L* - Rx timer time stamp register low | R/W | 0x0000_0000 | 19.5.3.3.12/19-62 |
| 0x4C8– 0x4FC | Reserved | — | — | — |
| eTSEC MAC Registers | | | | |
| 0x500 | MACCFG1—MAC configuration register 1 | R/W | 0x0000_0000 | 19.5.3.5.1/19-66 |
| 0x504 | MACCFG2—MAC configuration register 2 | R/W | 0x0000_7000 | 19.5.3.5.2/19-67 |
| 0x508 | IPGIFG—Inter-packet/inter-frame gap register | R/W | 0x4060_5060 | 19.5.3.5.3/19-69 |
| 0x50C | HAFDUP—Half-duplex control | R/W | 0x00A1_F037 | 19.5.3.5.4/19-70 |
| 0x510 | MAXFRM—Maximum frame length | R/W | 0x0000_0600 | 19.5.3.5.5/19-71 |
| 0x514– 0x51C | Reserved | — | — | — |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|--|--------|-------------|--|
| Offset | Register | Access | Reset | Section/Page |
| 0x520 | MIIMCFG—MII management configuration | R/W | 0x0000_0007 | 19.5.3.5.6/19-71 |
| 0x524 | MIIMCOM—MII management command | R/W | 0x0000_0000 | 19.5.3.5.7/19-72 |
| 0x528 | MIIMADD—MII management address | R/W | 0x0000_0000 | 19.5.3.5.8/19-73 |
| 0x52C | MIIMCON—MII management control | WO | 0x0000_0000 | 19.5.3.5.9/19-74 |
| 0x530 | MIIMSTAT—MII management status | R | 0x0000_0000 | 19.5.3.5.10/19-74 |
| 0x534 | MIIMIND—MII management indicator | R | 0x0000_0000 | 19.5.3.5.11/19-75 |
| 0x538 | Reserved | — | — | — |
| 0x53C | IFSTAT—Interface status | R | 0x0000_0000 | 19.5.3.5.12/19-75 |
| 0x540 | MACSTNADDR1—MAC station address register 1 | R/W | 0x0000_0000 | 19.5.3.5.13/19-76 |
| 0x544 | MACSTNADDR2—MAC station address register 2 | R/W | 0x0000_0000 | 19.5.3.5.14/19-77 |
| 0x548 | MAC01ADDR1*—MAC exact match address 1, part 1 | R/W | 0x0000_0000 | 19.5.3.5.15/19-77 19.5.3.5.16/19-78 |
| 0x54C | MAC01ADDR2*—MAC exact match address 1, part 2 | R/W | 0x0000_0000 | |
| 0x550 | MAC02ADDR1*—MAC exact match address 2, part 1 | R/W | 0x0000_0000 | |
| 0x554 | MAC02ADDR2*—MAC exact match address 2, part 2 | R/W | 0x0000_0000 | |
| 0x558 | MAC03ADDR1*—MAC exact match address 3, part 1 | R/W | 0x0000_0000 | |
| 0x55C | MAC03ADDR2*—MAC exact match address 3, part 2 | R/W | 0x0000_0000 | |
| 0x560 | MAC04ADDR1*—MAC exact match address 4, part 1 | R/W | 0x0000_0000 | |
| 0x564 | MAC04ADDR2*—MAC exact match address 4, part 2 | R/W | 0x0000_0000 | |
| 0x568 | MAC05ADDR1*—MAC exact match address 5, part 1 | R/W | 0x0000_0000 | |
| 0x56C | MAC05ADDR2*—MAC exact match address 5, part 2 | R/W | 0x0000_0000 | |
| 0x570 | MAC06ADDR1*—MAC exact match address 6, part 1 | R/W | 0x0000_0000 | |
| 0x574 | MAC06ADDR2*—MAC exact match address 6, part 2 | R/W | 0x0000_0000 | |
| 0x578 | MAC07ADDR1*—MAC exact match address 7, part 1 | R/W | 0x0000_0000 | |
| 0x57C | MAC07ADDR2*—MAC exact match address 7, part 2 | R/W | 0x0000_0000 | |
| 0x580 | MAC08ADDR1*—MAC exact match address 8, part 1 | R/W | 0x0000_0000 | |
| 0x584 | MAC08ADDR2*—MAC exact match address 8, part 2 | R/W | 0x0000_0000 | |
| 0x588 | MAC09ADDR1*—MAC exact match address 9, part 1 | R/W | 0x0000_0000 | |
| 0x58C | MAC09ADDR2*—MAC exact match address 9, part 2 | R/W | 0x0000_0000 | |
| 0x590 | MAC10ADDR1*—MAC exact match address 10, part 1 | R/W | 0x0000_0000 | |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|---|--------|-------------|--|
| Offset | Register | Access | Reset | Section/Page |
| 0x594 | MAC10ADDR2*—MAC exact match address 10, part 2 | R/W | 0x0000_0000 | 19.5.3.5.15/19-77 19.5.3.5.16/19-78 |
| 0x598 | MAC11ADDR1*—MAC exact match address 11, part 1 | R/W | 0x0000_0000 | |
| 0x59C | MAC11ADDR2*—MAC exact match address 11, part 2 | R/W | 0x0000_0000 | |
| 0x5A0 | MAC12ADDR1*—MAC exact match address 12, part 1 | R/W | 0x0000_0000 | |
| 0x5A4 | MAC12ADDR2*—MAC exact match address 12, part 2 | R/W | 0x0000_0000 | |
| 0x5A8 | MAC13ADDR1*—MAC exact match address 13, part 1 | R/W | 0x0000_0000 | |
| 0x5AC | MAC13ADDR2*—MAC exact match address 13, part 2 | R/W | 0x0000_0000 | |
| 0x5B0 | MAC14ADDR1*—MAC exact match address 14, part 1 | R/W | 0x0000_0000 | |
| 0x5B4 | MAC14ADDR2*—MAC exact match address 14, part 2 | R/W | 0x0000_0000 | |
| 0x5B8 | MAC15ADDR1*—MAC exact match address 15, part 1 | R/W | 0x0000_0000 | |
| 0x5BC | MAC15ADDR2*—MAC exact match address 15, part 2 | R/W | 0x0000_0000 | |
| 0x5C0– 0x67C | Reserved | — | — | |
| eTSEC Transmit and Receive Counters | | | | |
| 0x680 | TR64—Transmit and receive 64-byte frame counter | R/W | 0x0000_0000 | 19.5.3.6.1/19-79 |
| 0x684 | TR127—Transmit and receive 65- to 127-byte frame counter | R/W | 0x0000_0000 | 19.5.3.6.2/19-80 |
| 0x688 | TR255—Transmit and receive 128- to 255-byte frame counter | R/W | 0x0000_0000 | 19.5.3.6.3/19-80 |
| 0x68C | TR511—Transmit and receive 256- to 511-byte frame counter | R/W | 0x0000_0000 | 19.5.3.6.4/19-81 |
| 0x690 | TR1K—Transmit and receive 512- to 1023-byte frame counter | R/W | 0x0000_0000 | 19.5.3.6.5/19-81 |
| 0x694 | TRMAX—Transmit and receive 1024- to 1518-byte frame counter | R/W | 0x0000_0000 | 19.5.3.6.6/19-82 |
| 0x698 | TRMGV—Transmit and receive 1519- to 1522-byte good VLAN frame count | R/W | 0x0000_0000 | 19.5.3.6.7/19-82 |
| eTSEC Receive Counters | | | | |
| 0x69C | RBYT—Receive byte counter | R/W | 0x0000_0000 | 19.5.3.6.8/19-83 |
| 0x6A0 | RPKT—Receive packet counter | R/W | 0x0000_0000 | 19.5.3.6.9/19-83 |
| 0x6A4 | RFCS—Receive FCS error counter | R/W | 0x0000_0000 | 19.5.3.6.10/19-83 |
| 0x6A8 | RMCA—Receive multicast packet counter | R/W | 0x0000_0000 | 19.5.3.6.11/19-84 |
| 0x6AC | RBCA—Receive broadcast packet counter | R/W | 0x0000_0000 | 19.5.3.6.12/19-84 |
| 0x6B0 | RXCF—Receive control frame packet counter | R/W | 0x0000_0000 | 19.5.3.6.13/19-85 |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|--|--------|-------------|-------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x6B4 | RXPF—Receive PAUSE frame packet counter | R/W | 0x0000_0000 | 19.5.3.6.14/19-85 |
| 0x6B8 | RXUO—Receive unknown OP code counter | R/W | 0x0000_0000 | 19.5.3.6.15/19-86 |
| 0x6BC | RALN—Receive alignment error counter | R/W | 0x0000_0000 | 19.5.3.6.16/19-86 |
| 0x6C0 | RFLR—Receive frame length error counter | R/W | 0x0000_0000 | 19.5.3.6.17/19-87 |
| 0x6C4 | RCDE—Receive code error counter | R/W | 0x0000_0000 | 19.5.3.6.18/19-87 |
| 0x6C8 | RCSE—Receive carrier sense error counter | R/W | 0x0000_0000 | 19.5.3.6.19/19-88 |
| 0x6CC | RUND—Receive undersize packet counter | R/W | 0x0000_0000 | 19.5.3.6.20/19-88 |
| 0x6D0 | ROVR—Receive oversize packet counter | R/W | 0x0000_0000 | 19.5.3.6.21/19-89 |
| 0x6D4 | RFRG—Receive fragments counter | R/W | 0x0000_0000 | 19.5.3.6.22/19-89 |
| 0x6D8 | RJBR—Receive jabber counter | R/W | 0x0000_0000 | 19.5.3.6.23/19-90 |
| 0x6DC | RDRP—Receive drop counter | R/W | 0x0000_0000 | 19.5.3.6.24/19-90 |
| eTSEC Transmit Counters | | | | |
| 0x6E0 | TBYT—Transmit byte counter | R/W | 0x0000_0000 | 19.5.3.6.25/19-91 |
| 0x6E4 | TPKT—Transmit packet counter | R/W | 0x0000_0000 | 19.5.3.6.26/19-91 |
| 0x6E8 | TMCA—Transmit multicast packet counter | R/W | 0x0000_0000 | 19.5.3.6.27/19-92 |
| 0x6EC | TBCA—Transmit broadcast packet counter | R/W | 0x0000_0000 | 19.5.3.6.28/19-92 |
| 0x6F0 | TXPF—Transmit PAUSE control frame counter | R/W | 0x0000_0000 | 19.5.3.6.29/19-93 |
| 0x6F4 | TDFR—Transmit deferral packet counter | R/W | 0x0000_0000 | 19.5.3.6.30/19-93 |
| 0x6F8 | TEDF—Transmit excessive deferral packet counter | R/W | 0x0000_0000 | 19.5.3.6.31/19-94 |
| 0x6FC | TSCL—Transmit single collision packet counter | R/W | 0x0000_0000 | 19.5.3.6.32/19-94 |
| 0x700 | TMCL—Transmit multiple collision packet counter | R/W | 0x0000_0000 | 19.5.3.6.33/19-95 |
| 0x704 | TLCL—Transmit late collision packet counter | R/W | 0x0000_0000 | 19.5.3.6.34/19-95 |
| 0x708 | TXCL—Transmit excessive collision packet counter | R/W | 0x0000_0000 | 19.5.3.6.35/19-96 |
| 0x70C | TNCL—Transmit total collision counter | R/W | 0x0000_0000 | 19.5.3.6.36/19-96 |
| 0x710 | Reserved | — | — | — |
| 0x714 | TDRP—Transmit drop frame counter | R/W | 0x0000_0000 | 19.5.3.6.37/19-97 |
| 0x718 | TJBR—Transmit jabber frame counter | R/W | 0x0000_0000 | 19.5.3.6.38/19-97 |
| 0x71C | TFCS—Transmit FCS error counter | R/W | 0x0000_0000 | 19.5.3.6.39/19-98 |
| 0x720 | TXCF—Transmit control frame counter | R/W | 0x0000_0000 | 19.5.3.6.40/19-98 |
| 0x724 | TOVR—Transmit oversize frame counter | R/W | 0x0000_0000 | 19.5.3.6.41/19-99 |
| 0x728 | TUND—Transmit undersize frame counter | R/W | 0x0000_0000 | 19.5.3.6.42/19-99 |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|---|--------|-------------|------------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x72C | TFRG—Transmit fragments frame counter | R/W | 0x0000_0000 | 19.5.3.6.43/19-100 |
| eTSEC Counter Control and TOE Statistics Registers | | | | |
| 0x730 | CAR1—Carry register one register ¹ | R | 0x0000_0000 | 19.5.3.6.44/19-100 |
| 0x734 | CAR2—Carry register two register ¹ | R | 0x0000_0000 | 19.5.3.6.45/19-102 |
| 0x738 | CAM1—Carry register one mask register | R/W | 0xFE03_FFFF | 19.5.3.6.46/19-103 |
| 0x73C | CAM2—Carry register two mask register | R/W | 0x000F_FFFD | 19.5.3.6.47/19-104 |
| 0x740 | RREJ*—Receive filer rejected packet counter | R/W | 0x0000_0000 | 19.5.3.6.48/19-105 |
| 0x744– 0x7FC | Reserved | — | — | — |
| Hash Function Registers | | | | |
| 0x800 | IGADDR0—Individual/group address register 0 | R/W | 0x0000_0000 | 19.5.3.7.1/19-106 |
| 0x804 | IGADDR1—Individual/group address register 1 | R/W | 0x0000_0000 | |
| 0x808 | IGADDR2—Individual/group address register 2 | R/W | 0x0000_0000 | |
| 0x80C | IGADDR3—Individual/group address register 3 | R/W | 0x0000_0000 | |
| 0x810 | IGADDR4—Individual/group address register 4 | R/W | 0x0000_0000 | |
| 0x814 | IGADDR5—Individual/group address register 5 | R/W | 0x0000_0000 | |
| 0x818 | IGADDR6—Individual/group address register 6 | R/W | 0x0000_0000 | |
| 0x81C | IGADDR7—Individual/group address register 7 | R/W | 0x0000_0000 | |
| 0x820– 0x87C | Reserved | — | — | — |
| 0x880 | GADDR0—Group address register 0 | R/W | 0x0000_0000 | 19.5.3.7.2/19-106 |
| 0x884 | GADDR1—Group address register 1 | R/W | 0x0000_0000 | |
| 0x888 | GADDR2—Group address register 2 | R/W | 0x0000_0000 | |
| 0x88C | GADDR3—Group address register 3 | R/W | 0x0000_0000 | |
| 0x890 | GADDR4—Group address register 4 | R/W | 0x0000_0000 | |
| 0x894 | GADDR5—Group address register 5 | R/W | 0x0000_0000 | |
| 0x898 | GADDR6—Group address register 6 | R/W | 0x0000_0000 | |
| 0x89C | GADDR7—Group address register 7 | R/W | 0x0000_0000 | |
| 0x8A0– 0xAFC | Reserved | — | — | — |
| eTSEC DMA Attribute Registers | | | | |
| 0xB00– 0xBF4 | Reserved | — | — | — |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|---|--------|-------------|-----------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0xBF8 | ATTR—Attribute register | R/W | 0x0000_0000 | 19.5.3.8.1/19-107 |
| eTSEC Lossless Flow Control Registers | | | | |
| 0xC00 | RQPRM0*—Receive Queue Parameters register 0 | R/W | 0x0000_0000 | 19.5.3.9.1/19-108 |
| 0xC04 | RQPRM1*—Receive Queue Parameters register 1 | R/W | 0x0000_0000 | |
| 0xC08 | RQPRM2*—Receive Queue Parameters register 2 | R/W | 0x0000_0000 | |
| 0xC0C | RQPRM3*—Receive Queue Parameters register 3 | R/W | 0x0000_0000 | |
| 0xC10 | RQPRM4*—Receive Queue Parameters register 4 | R/W | 0x0000_0000 | |
| 0xC14 | RQPRM5*—Receive Queue Parameters register 5 | R/W | 0x0000_0000 | |
| 0xC18 | RQPRM6*—Receive Queue Parameters register 6 | R/W | 0x0000_0000 | |
| 0xC1C | RQPRM7*—Receive Queue Parameters register 7 | R/W | 0x0000_0000 | |
| 0xC20–0xC40 | Reserved | — | — | — |
| 0xC44 | RFBPTR0*—Last Free RxBD pointer for ring 0 | R/W | 0x0000_0000 | 19.5.3.9.2/19-109 |
| 0xC48 | Reserved | — | — | — |
| 0xC4C | RFBPTR1*—Last Free RxBD pointer for ring 1 | R/W | 0x0000_0000 | 19.5.3.9.2/19-109 |
| 0xC50 | Reserved | — | — | — |
| 0xC54 | RFBPTR2*—Last Free RxBD pointer for ring 2 | R/W | 0x0000_0000 | 19.5.3.9.2/19-109 |
| 0xC58 | Reserved | — | — | — |
| 0xC5C | RFBPTR3*—Last Free RxBD pointer for ring 3 | R/W | 0x0000_0000 | 19.5.3.9.2/19-109 |
| 0xC60 | Reserved | — | — | — |
| 0xC64 | RFBPTR4*—Last Free RxBD pointer for ring 4 | R/W | 0x0000_0000 | 19.5.3.9.2/19-109 |
| 0xC68 | Reserved | — | — | — |
| 0xC6C | RFBPTR5*—Last Free RxBD pointer for ring 5 | R/W | 0x0000_0000 | 19.5.3.9.2/19-109 |
| 0xC70 | Reserved | — | — | — |
| 0xC74 | RFBPTR6*—Last Free RxBD pointer for ring 6 | R/W | 0x0000_0000 | 19.5.3.9.2/19-109 |
| 0xC78 | Reserved | — | — | — |
| 0xC7C | RFBPTR7*—Last Free RxBD pointer for ring 7 | R/W | 0x0000_0000 | 19.5.3.9.2/19-109 |
| eTSEC Future Expansion Space | | | | |
| 0xCC0–0xD94 | Reserved | — | — | — |
| eTSEC IEEE 1588 Registers | | | | |

Table A-27. Enhanced Three-Speed Ethernet Controllers (eTSEC) Registers (continued)

| eTSEC 1—Block Base Address 0x2_4000 eTSEC 2—Block Base Address 0x2_5000 | | | | |
|--|---|--------|-------------|-------------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0xE00 | TMR_CTRL* - Timer control register | R/W | 0x0001_0000 | 19.5.3.10.1/19-110 |
| 0xE04 | TMR_TEVENT* - time stamp event register | R/W | 0x0000_0000 | 19.5.3.10.2/19-111 |
| 0xE08 | TMR_TEMASK* - Timer event mask register | R/W | 0x0000_0000 | 19.5.3.10.3/19-113 |
| 0xE0C | TMR_PEVENT* - time stamp event register | R/W | 0x0000_0000 | 19.5.3.10.4/19-113 |
| 0xE10 | TMR_PEMASK* - Timer event mask register | R/W | 0x0000_0000 | 19.5.3.10.5/19-114 |
| 0xE14 | TMR_STAT* - time stamp status register | R/W | 0x0000_0000 | 19.5.3.10.6/19-115 |
| 0xE18 | TMR_CNT_H* - timer counter high register | R/W | 0x0000_0000 | 19.5.3.10.7/19-115 |
| 0xE1C | TMR_CNT_L* - timer counter low register | R/W | 0x0000_0000 | 19.5.3.10.7/19-115 |
| 0xE20 | TMR_ADD* - Timer drift compensation addend register | R/W | 0x0000_0000 | 19.5.3.10.8/19-116 |
| 0xE24 | TMR_ACC* - Timer accumulator register | R/W | 0x0000_0000 | 19.5.3.10.9/19-117 |
| 0xE28 | TMR_PRSC* - timer prescale | R/W | 0x0000_0002 | 19.5.3.10.10/19-117 |
| 0xE2C | Reserved | — | — | — |
| 0xE30 | TMROFF_H* - Timer offset high | R/W | 0x0000_0000 | 19.5.3.10.11/19-118 |
| 0xE34 | TMROFF_L* - Timer offset low | R/W | 0x0000_0000 | 19.5.3.10.11/19-118 |
| 0xE40 | TMR_ALARM1_H* - Timer alarm 1 high register | R/W | 0xFFFF_FFFF | 19.5.3.10.12/19-118 |
| 0xE44 | TMR_ALARM1_L* - Timer alarm 1 high register | R/W | 0xFFFF_FFFF | |
| 0xE48 | TMR_ALARM2_H* - Timer alarm 2 high register | R/W | 0xFFFF_FFFF | |
| 0xE4C | TMR_ALARM2_L* - Timer alarm 2 high register | R/W | 0xFFFF_FFFF | |
| 0xE50– 0xE7C | Reserved | — | — | — |
| 0xE80 | TMR_FIPER1* - Timer fixed period interval | R/W | 0xFFFF_FFFF | 19.5.3.10.13/19-119 |
| 0xE84 | TMR_FIPER2* - Timer fixed period interval | R/W | 0xFFFF_FFFF | |
| 0xE88 | TMR_FIPER*3 - Timer fixed period interval | R/W | 0xFFFF_FFFF | |
| 0xEA0 | TMR_ETTS1_H* - Time stamp of general purpose external trigger | R/W | 0x0000_0000 | 19.5.3.10.14/19-120 |
| 0xEA4 | TMR_ETTS1_L* - Time stamp of general purpose external trigger | R/W | 0x0000_0000 | |
| 0xEA8 | TMR_ETTS2_H* - Time stamp of general purpose external trigger | R/W | 0x0000_0000 | |
| 0xEAC | TMR_ETTS2_L* - Time stamp of general purpose external trigger | R/W | 0x0000_0000 | |
| 0xEB0– 0xFFF | Reserved | — | — | — |

¹ Cleared on read.

A.27 TDM DMA Controller (DMAC)

Table A-28. TDM DMA Controller (DMAC) Registers

| TDM DMA Controller (DMAC)—Block Base Address 0x2_C000 | | | | |
|---|--|--------|-------------|------------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | DMACR—DMA Control Register | R/W | 0x0000_E400 | 25.10.2.1/2525-57 |
| 0x004 | DMAES—DMA Error Status Register | RO | 0x0000_0000 | 25.10.2.1/2525-59 |
| 0x008 | Reserved | — | — | — |
| 0x00C | DMAERQ—DMA enable request register | R/W | 0x0000_0000 | 25.10.3.1/2525-62 |
| 0x010 | Reserved | — | — | — |
| 0x014 | DMAEEI—DMA enable error interrupt register | R/W | 0x0000_0000 | 25.10.3.2/2525-63 |
| 0x018 | DMASERQ—DMA set enable request | R/W | 0x00 | 25.10.3.3/2525-64 |
| 0x019 | DMACERQ—DMA clear enable request | R/W | 0x00 | 25.10.3.4/2525-64 |
| 0x01A | DMASEEI—DMA Set Enable Error Interrupt | R/W | 0x00 | 25.10.3.5/2525-65 |
| 0x01B | DMACEEI—DMA Clear Enable Error Interrupt | R/W | 0x00 | 25.10.3.6/2525-65 |
| 0x01C | DMACINT—DMA Clear Interrupt Request | R/W | 0x00 | 25.10.3.7/2525-66 |
| 0x01D | DMACERR—DMA Clear Error | R/W | 0x00 | 25.10.3.8/2525-67 |
| 0x01E | DMASSRT—DMA Set START Bit | R/W | 0x00 | 25.10.3.9/2525-67 |
| 0x01F | DMACDNE—DMA Clear DONE Status Bit | R/W | 0x00 | 25.10.3.10/2525-68 |
| 0x020 | Reserved | — | — | — |
| 0x024 | DMAINT—DMA interrupt request register | w1c | 0x0000_0000 | 25.10.3.11/2525-68 |
| 0x028 | Reserved | — | — | — |
| 0x02C | DMAERR—DMA error register | w1c | 0x0000_0000 | 25.10.3.12/2525-69 |
| 0x030 | Reserved | — | — | — |

Table A-28. TDM DMA Controller (DMAC) Registers (continued)

| TDM DMA Controller (DMAC)—Block Base Address 0x2_C000 | | | | |
|---|---|--------|-------------|------------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x034 | DMAHRS—DMA hardware request status register | R/W | 0x000_0000 | 25.10.3.13/2525-70 |
| 0x038 | DMAGPOR—DMA general purpose output register | R/W | 0x0000_0000 | 25.10.3.14/2525-71 |
| 0x03C–0x0FC | Reserved | — | — | — |

A.28 Security Engine Controller (SEC)

Table A-29. Security Engine Controller (SEC) Registers

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|----------------------------|----------------------------|------------------------|----------|--------------------------------|
| Byte Offset (AD 17–0) | SEC Module | Register | Access | Write By | Section/Page |
| 0x1008 | Controller | Interrupt Enable | R/W | Byte | 18.6.4.2/18-45 |
| 0x1010 | | Interrupt Status | R | — | 18.6.4.3/18-47 |
| 0x1018 | | Interrupt Clear | R/W | Byte | 18.6.4.4/18-48 |
| 0x1020 | | Identification | R | — | 18.6.4.5/18-49 |
| 0x1028 | | EU Assignment Status | R | — | 18.6.4.1/18-44 |
| 0x1030 | | Master Control | R/W | Byte | 18.6.4.7/18-51 |
| 0x1108 | | Channel_1 | Configuration Register | R/W | Word |
| 0x1110 | Pointer Status | | R/W | — | 18.5.4.2/18-36 |
| 0x1140 | Current Descriptor Pointer | | R | — | 18.5.4.3/18-38 |
| 0x1148 | Fetch FIFO | | W | Word | 18.5.4.4/18-39 |
| 0x1180–0x11BF | Descriptor Buffer | | R | — | 18.5.5.1/18-40 |
| 0x11C0–0x11DF | Gather Link Tables | | W | Byte | 18.5.5.2/18-40 |
| 0x11E0–0x11FF | Scatter Link Tables | | W | Byte | 18.5.5.2/18-40 |
| 0x1208 | Channel_2 | Config Register | R/W | Word | 18.5.4.1/18-34 |
| 0x1210 | | Pointer Status | R/W | — | 18.5.4.2/18-36 |
| 0x1240 | | Current Descriptor Pointer | R | — | 18.5.4.3/18-38 |
| 0x1248 | | Fetch FIFO | W | Word | 18.5.4.4/18-39 |
| 0x1280–0x12BF | | Descriptor Buffer | R | — | 18.5.5.1/18-40 |
| 0x12C0–0x12DF | | Gather Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x12E0 - 0x12FF | | Scatter Link Tables | W | Byte | 18.5.5.2/18-40 |

Table A-29. Security Engine Controller (SEC) Registers (continued)

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|------------|----------------------------|------------------|----------|---------------------------------|
| Byte Offset (AD 17–0) | SEC Module | Register | Access | Write By | Section/Page |
| 0x1308 | Channel_3 | Config Register | R/W | Word | 18.5.4.1/18-34 |
| 0x1310 | | Pointer Status | R/W | — | 18.5.4.2/18-36 |
| 0x1340 | | Current Descriptor Pointer | R | — | 18.5.4.3/18-38 |
| 0x1348 | | Fetch FIFO | W | Word | 18.5.4.4/18-39 |
| 0x1380–0x13BF | | Descriptor Buffer | R | — | 18.5.5.1/18-40 |
| 0x13C0–0x13DF | | Gather Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x13E0–0x13FF | | Scatter Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x1408 | Channel_4 | Config Register | R/W | Word | 18.5.4.1/18-34 |
| 0x1410 | | Pointer Status | R/W | — | 18.5.4.2/18-36 |
| 0x1440 | | Current Descriptor Pointer | R | — | 18.5.4.3/18-38 |
| 0x1448 | | Fetch FIFO | W | Word | 18.5.4.4/18-39 |
| 0x1480–0x14BF | | Descriptor Buffer | R | — | 18.5.5.1/18-40 |
| 0x14C0–0x14DF | | Gather Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x14E0–0x14FF | | Scatter Link Tables | W | Byte | 18.5.5.2/18-40 |
| 0x1BF8 | Controller | IP block revision | R | — | 18.6.4.6/18-50 |
| 0x2000 | DEU | Mode Register | R/W | Word | 18.8.3.1/18-89 |
| 0x2008 | | Key Size Register | R/W | Word | 18.8.3.2/18-90 |
| 0x2010 | | Data Size Register | R/W | Word | 18.8.3.3/18-90 |
| 0x2018 | | Reset Control Register | R/W | Word | 18.8.3.4/18-91 |
| 0x2028 | | Status Register | R | — | 18.8.3.5/18-92 |
| 0x2030 | | Interrupt Status Register | R/W | Word | 18.8.3.6/18-92 |
| 0x2038 | | Interrupt Mask Register | R/W | Word | 18.8.3.7/18-94 |
| 0x2050 | | EU-Go | W | Word | 18.8.3.8/18-96 |
| 0x2100 | | IV Register | R/W | Word | 18.8.3.9/18-96 |
| 0x2400 | | Key 1 Register | W | Byte | 18.8.3.10/18-96 |
| 0x2408 | | Key 2 Register | W | Byte | 18.8.3.10/18-96 |
| 0x2410 | | Key 3 Register | W | Byte | 18.8.3.10/18-96 |
| 0x2800–0x2FFF | | Input FIFO / Output FIFO | R/W ¹ | Byte | 18.8.3.11/18-97 |

Table A-29. Security Engine Controller (SEC) Registers (continued)

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|---------------------------|---------------------------|------------------|----------|-----------------------------------|
| Byte Offset (AD 17–0) | SEC Module | Register | Access | Write By | Section/Page |
| 0x4000 | AESU | Mode Register | R/W | Word | 18.8.1.2/18-54 |
| 0x4008 | | Key Size Register | R/W | Word | 18.8.1.3/18-57 |
| 0x4010 | | Data Size Register | R/W | Word | 18.8.1.4/18-57 |
| 0x4018 | | Reset Control Register | R/W | Word | 18.8.1.5/18-58 |
| 0x4028 | | Status Register | R | — | 18.8.1.6/18-58 |
| 0x4030 | | Interrupt Status Register | R/W | Word | 18.8.1.7/18-59 |
| 0x4038 | | Interrupt Mask Register | R/W | Word | 18.8.1.8/18-61 |
| 0x4050 | | End Of Message Register | W | Word | 18.8.1.10/18-63 |
| 0x4100–0x415F | | Context | R/W | Byte | 18.8.1.11/18-64 |
| 0x4400–0x441F | | Key Memory | R/W | Byte | 18.8.1.11.8/18-79 |
| 0x4800–0x4FFF | | Input FIFO / Output FIFO | R/W ¹ | Byte | 18.8.1.11.9/18-79 |
| 0x6000 | | MDEU | Mode Register | R/W | Word |
| 0x6008 | Key Size Register | | R/W | Word | 18.8.4.4/18-102 |
| 0x6010 | Data Size Register | | R/W | Word | 18.8.4.5/18-102 |
| 0x6018 | Reset Control Register | | R/W | Word | 18.8.4.6/18-102 |
| 0x6028 | Status Register | | R | — | 18.8.4.7/18-103 |
| 0x6030 | Interrupt Status Register | | R/W | Word | 18.8.4.8/18-104 |
| 0x6038 | Interrupt Mask Register | | R/W | Word | 18.8.4.9/18-106 |
| 0x6040 | ICV Size Register | | W | Word | 18.8.4.10/18-107 |
| 0x6050 | End_of_message | | W | Word | 18.8.4.11/18-107 |
| 0x6100–0x6140 | Context Registers | | R/W | Byte | 18.8.4.12/18-108 |
| 0x6400–0x647F | Key Registers | | W | Byte | 18.8.4.13/18-111 |
| 0x6800–0x6FFF | Input FIFO | | W ¹ | Byte | 18.8.4.14/18-111 |

Table A-29. Security Engine Controller (SEC) Registers (continued)

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|---------------------------|---------------------------|----------------|----------|----------------------------------|
| Byte Offset (AD 17-0) | SEC Module | Register | Access | Write By | Section/Page |
| 0xA000 | RNGU | Mode Register | R/W | Word | 18.8.6.1/18-120 |
| 0xA010 | | Data Size Register | R/W | Word | 18.8.6.2/18-120 |
| 0xA018 | | Reset Control Register | R/W | Word | 18.8.6.3/18-121 |
| 0xA028 | | Status Register | R | — | 18.8.6.4/18-121 |
| 0xA030 | | Interrupt Status Register | R/W | Word | 18.8.6.5/18-122 |
| 0xA038 | | Interrupt Mask Register | R/W | Word | 18.8.6.6/18-123 |
| 0xA050 | | End_of_message | W | Word | 18.8.6.7/18-124 |
| 0xA400–0xA43F | | Entropy Registers | W | Word | 18.8.6.8/18-125 |
| 0xA800–0xAFFF | | Output FIFO | R ¹ | — | 18.8.6.8/18-125 |
| 0xC000 | | PKEU | Mode Register | R/W | Word |
| 0xC008 | Key Size Register | | R/W | Word | 18.8.5.2/18-112 |
| 0xC010 | Data Size Register | | R/W | Word | 18.8.5.4/18-114 |
| 0xC018 | Reset Control Register | | R/W | Word | 18.8.5.5/18-114 |
| 0xC028 | Status Register | | R | — | 18.8.5.6/18-115 |
| 0xC030 | Interrupt Status Register | | R/W | Word | 18.8.5.7/18-116 |
| 0xC038 | Interrupt Mask Register | | R/W | Word | 18.8.5.8/18-117 |
| 0xC040 | ABSize | | R/W | Word | 18.8.5.3/18-113 |
| 0xC050 | End_of_message | | W | Word | 18.8.5.9/18-118 |
| 0xC200–0xC27F | Parameter Memory A0 | | R/W | Byte | 18.8.5.10/18-119 |
| 0xC280–0xC2FF | Parameter Memory A1 | | R/W | Byte | |
| 0xC300–0xC37F | Parameter Memory A2 | | R/W | Byte | |
| 0xC380–0xC3FF | Parameter Memory A3 | | R/W | Byte | |
| 0xC400–0xC47F | Parameter Memory B0 | | R/W | Byte | |
| 0xC480–0xC4FF | Parameter Memory B1 | | R/W | Byte | |
| 0xC500–0xC57F | Parameter memory B2 | | R/W | Byte | |
| 0xC580–0xC5FF | Parameter Memory B3 | | R/W | Byte | |
| 0xC800–0xC9FF | Parameter Memory N | | R/W | Byte | |
| 0xCA00–0xCBFF | Parameter Memory E | | W | Byte | |

Table A-29. Security Engine Controller (SEC) Registers (continued)

| Security Engine Controller (SEC)—Block Base Address 0x3_0000 | | | | | |
|--|------------|---------------------------|----------------|----------|---------------------------------|
| Byte Offset (AD 17–0) | SEC Module | Register | Access | Write By | Section/Page |
| 0xF000 | CRCU | Mode Register | R/W | Word | 18.8.2.2/18-80 |
| 0xF008 | | Key Size Register | R/W | Word | 18.8.2.3/18-81 |
| 0xF010 | | Data Size Register | R/W | Word | 18.8.2.4/18-81 |
| 0xF018 | | Reset Control Register | R/W | Word | 18.8.2.5/18-82 |
| 0xF020 | | Control | R/W | Word | 18.8.2.6/18-82 |
| 0xF028 | | Status Register | R | — | 18.8.2.7/18-83 |
| 0xF030 | | Interrupt Status Register | R/W | Word | 18.8.2.8/18-84 |
| 0xF038 | | Interrupt Mask Register | R/W | Word | 18.8.2.9/18-86 |
| 0xF048 | | Data Out Register | R | Byte | 18.8.2.10/18-87 |
| 0xF050 | | End Of Message Register | W | Word | 18.8.2.11/18-87 |
| 0xF100 | | Received ICV Register | R/W | Byte | 18.8.2.12/18-87 |
| 0xF108 | | Context Register | R/W | Byte | 18.8.2.13/18-87 |
| 0xF400 | | Key Register | R/W | Byte | 18.8.2.14/18-88 |
| 0xF800–0xFFFF | | Input FIFO | W ¹ | Byte | 18.8.2.15/18-89 |

¹ For the EU FIFOs, write operations anywhere in the address range enqueue to the input FIFO, and read operations anywhere in the address range dequeue from the output FIFO. See the referenced section for more detailed information.

A.29 SerDes PHY

Table A-30. SerDes PHY Registers

| SerDes PHY—Block Base Address 0xE_3000 | | | | |
|--|--|--------|-------------|------------------------------|
| Offset | Register | Access | Reset | Section/Page |
| 0x000 | SRDSCR0—SerDes Control Register 0 | R/W | 0x1100_CC30 | 16.3.1/16-5 |
| 0x004 | SRDSCR1—SerDes Control Register 1 | R/W | 0x0000_0000 | 16.3.2/16-8 |
| 0x008 | SRDSCR2—SerDes Control Register 2 | R/W | 0x0080_0000 | 16.3.3/16-9 |
| 0x00C | SRDSCR3—SerDes Control Register 3 | R/W | 0x0101_0000 | 16.3.4/16-10 |
| 0x010 | SRDSCR4—SerDes Control Register 4 | R/W | 0xnn00_0n0n | 16.3.5/16-11 |
| 0x014–0x01C | Reserved | — | — | — |
| 0x020 | SRDSRSTCTL—SerDes Reset Control Register | R/W | 0x0044_4500 | 16.3.6/16-13 |
| 0x024–0x1FC | Reserved | — | — | — |



Appendix B

Revision History

This appendix provides a list of the major differences between revisions of *MPC8315E PowerQUICC II Pro Integrated Host Processor Reference Manual*.

B.1 Changes From Revision 1 to Revision 2

Major changes to the *MPC8315E PowerQUICC II Pro Integrated Host Processor Reference Manual*, from Revision 1 to Revision 2 are as follows:

| Section, Page | Changes |
|-----------------------|---|
| | Throughout book—Replace all instances of LFR \bar{B} with $\overline{\text{LFRB}}$. |
| | Throughout book—Replace LBIUCM with LBCM. |
| 1.1, 1-9 | Remove the following list item appearing under “Power management controller (PMC)”: “Support for wake-on-PCI Express (in D3hot state) and inbound PCI Express PME events (in D3hot state)” |
| 1.2.12, 1-20 | Remove references to three TDMs by changing the second paragraph as follows: “The TDM interface supports 128 channels running at up to 50 Mbps with 8- and 16-bit word size. The TDM bus connects gluelessly to most T1/E1 frames as well as to common buses such as the H.110, SCAS, and MVIP. The TDM also supports an I2S mode and operates in independent or shared mode when receiving or transmitting data.” |
| Chapter 2, throughout | Relabeled MA[0:14] to MA[14:0]. |
| 2.1, 2-3 | In Figure 2-2, remove USB_PHY_GND signal. |
| 2.1, 2-3 | In Figure 2-2, “MPC8315E Signal Groupings (2 of 2),” change “ $\overline{\text{CORE_RESET_IN}}$ ” to “ $\overline{\text{CORE_SRESET_IN}}$,” and “ $\overline{\text{CORE_SRESET}}$ ” to “ $\overline{\text{CORE_SRESET_IN}}$.” |
| 2.1, 2-13 | In Table 2-1, “MPC8315E Signal Reference by Functional Block,” change “ $\overline{\text{CORE_RESET_IN}}$ ” to “ $\overline{\text{CORE_SRESET_IN}}$,” and add a row for “ $\overline{\text{CORE_SRESET_IN}}$,” as follows: |

Table 2-1. MPC8315E Signal Reference by Functional Block

| Name | Description | Functional Block | No. of Signals | I/O | Table/ Page | Alternate Function(s) | Table/ Page |
|--------------------------------------|-----------------------------|------------------|----------------|-----|-------------|-----------------------|-------------|
| $\overline{\text{CORE_SRESET_IN}}$ | Soft reset to the e300 core | System control | 1 | I | 4-1/4-1 | — | — |

Revision History

- 2.3, 2-31 In Table 2-3, “Output Signal States During System Reset,” update reset value of PCI_SYNC_OUT from “High-Z” to “Driven Low.”
- 4.1.1, 4-2 In Table 4-1, “System Control Signals,” add CORE_SRESET_IN, as follows:

Table 4-1. System Control Signals

| Signal | I/O | Description | | |
|---|--|--|--|--|
| CORE_SRESET_IN | I | The CORE_SRESET_IN input is connected directly to the soft reset input of the e300c3 core. The assertion of the e300 soft reset input, CORE_SRESET_IN, causes a high priority interrupt to the e300 core as described in Section 4.2.1, “Reset Operations.” It does not reset the e300 core or any other portion of the device. The CORE_SRESET_IN input is not registered in the reset status register (RSR). | | |
| | | <table border="1"> <tr> <td>State Meaning</td> <td>Asserted—Indicates that the processor must initiate a system reset interrupt. Negated—Indicates that the interrupt is not being requested.</td> </tr> </table> | State Meaning | Asserted—Indicates that the processor must initiate a system reset interrupt. Negated—Indicates that the interrupt is not being requested. |
| | | State Meaning | Asserted—Indicates that the processor must initiate a system reset interrupt. Negated—Indicates that the interrupt is not being requested. | |
| | | <table border="1"> <tr> <td>Timing</td> <td>Assertion—Occurs at any time, asynchronously to any clock. Negation—Occurs after being serviced. (PCI host mode) or PCI_CLK (PCI agent mode).</td> </tr> </table> | Timing | Assertion—Occurs at any time, asynchronously to any clock. Negation—Occurs after being serviced. (PCI host mode) or PCI_CLK (PCI agent mode). |
| | | Timing | Assertion—Occurs at any time, asynchronously to any clock. Negation—Occurs after being serviced. (PCI host mode) or PCI_CLK (PCI agent mode). | |
| <table border="1"> <tr> <td>Requirements</td> <td>An open-drain signal. An external pull-up is required.</td> </tr> </table> | Requirements | An open-drain signal. An external pull-up is required. | | |
| Requirements | An open-drain signal. An external pull-up is required. | | | |
| <table border="1"> <tr> <td>Reset State</td> <td>Always input.</td> </tr> </table> | Reset State | Always input. | | |
| Reset State | Always input. | | | |

- 4.2.1, 4-4 Add the following to the bulleted list:
 - External soft reset (CORE_SRESET_IN)

- 4.2.1.1, 4-5 In Table 4-3, “Reset Causes,” modify the Soft reset row as follows:

Table 4-3. Reset Causes

| Name | Description |
|------------------------------|--|
| Soft reset CORE_SRESET_IN | Input signal. Connected to the <i>sreset</i> input of the e300 core, and when asserted, causes a high priority interrupt to the e300 core. |

- 4.2.1.2, 4-5 Add the following text regarding soft reset in the first paragraph.

“The CORE_SRESET_IN input is not routed to the reset control logic, but directly to the *sreset* input on the e300 core.”
- 4.2.1.2, 4-6 Remove “External Soft Reset” column from Table 4-4, “Reset Actions” and change “JTAG Reset” and “Hard reset to e300 core” value from “yes” to “no” as follows:

| Action | Reset Source | | |
|---|----------------|---|------------|
| | Power-On Reset | External Hard Reset Software Watchdog Bus Monitor Checkstop Software Hard Reset | JTAG Reset |
| Resets: PLLs, clocks, RTC unit, and error capture registers | Yes | No | No |
| Resets: DDR, LBC, I/O multiplexers, GTM, PIT, GPIO, system configuration, and local access windows | Yes | Yes | No |
| Resets other internal logic | Yes | Yes | Yes |
| Reset configuration words loaded | Yes | Yes | No |
| $\overline{\text{HRESET}}$ driven | Yes | Yes | No |
| Hard reset to e300 core | Yes | Yes | No |
| High priority interrupt to the e300 core | No | No | No |

4.2.4, 4-9

Add Section 4.2.4, “Soft Reset Flow,” as follows:

“The CORE_SRESET_IN signal can be initiated externally by asserting CORE_SRESET_IN. CORE_SRESET_IN assertion asserts the *sreset* input to the e300 core. No other hardware is reset.”

4.3.3.1, 4-22

Add the following text to the end of the first paragraph:

“+ LCS0 is the default for GPCM, so GPCM controlled is used to read the reset configuration word from EEPROM. /LGTA should be high to avoid unintended early termination of the read cycle.”

4.3.3.3, 4-28

In Table 4-25, “Hard-Coded Reset Configuration Word High Field Values,” add CFG_RESET_SOURCE[0:3] = 1011 column and update 1100 TSEC1M and TSEC2M values, as follows:

Table 4-25. Hard-Coded Reset Configuration Word High Field Values

| Bits | Name | Field Values when CFG_RESET_SOURCE[0:3] = 1000–1100 | | | | | Meaning |
|-------|----------|--|------|------|------|------|---|
| | | 1000 | 1001 | 1010 | 1011 | 1100 | |
| 0 | PCIHOST | 0 | | | | | PCI agent mode |
| 1 | Reserved | 1 | | | | | — |
| 2 | PCIARB | 0 | | | | | External arbiter is used |
| 3 | Reserved | 0 | | | | | — |
| 4 | COREDIS | 1 | | | | | e300 core is disabled (boot holdoff) |
| 5 | BMS | 1 | | | | | Boot memory space is 0xFF80_0000–0xFFFF_FFFF. MSR[IP] initial value is 0b1. |
| 6–7 | BOOTSEQ | 00 | | | | | Boot sequencer is disabled. |
| 8 | SWEN | 0 | | | | | Software watchdog disabled. |
| 9–11 | ROMLOC | 000 | | | | | Boot ROM interface location. |
| 12–13 | RLEXT | 00 | | | | | Legacy mode. |
| 14–15 | Reserved | 00 | | | | | — |
| 16–18 | TSEC1M | 101 | 011 | 000 | 001 | 000 | 000 = MII mode 001 = RMII 011 = RGMII mode 101 = RTBI mode |
| 19–21 | TSEC2M | 000 | 101 | 001 | 101 | 011 | |
| 22–27 | Reserved | 000000 | | | | | — |
| 28 | TLE | 0 | | | | | Big-endian mode |
| 29 | LALE | 0 | | | | | Normal timing |
| 30–31 | Reserved | 00 | | | | | — |

- 4.3.3.3, 4-28 Add column 1011 and modify TSEC1 and TSEC2 field values for column 1100 in Table 4-25, “Hard-Coded Reset Configuration Word High Field Values,” as follows. Note that the following rows are only part of Table 4-25:

Table 4-25. Hard-Coded Reset Configuration Word High Field Values

| Bits | Name | Field Values when CFG_RESET_SOURCE[0:3] = 1000–1100 | | | | | Meaning |
|-------|--------|---|------|------|------|------|---|
| | | 1000 | 1001 | 1010 | 1011 | 1100 | |
| 16–18 | TSEC1M | 101 | 011 | 000 | 001 | 000 | 000 = MII mode 001 = RMII 011 = RGMII mode 101 = RTBI mode |
| 19–21 | TSEC2M | 000 | 101 | 001 | 101 | 011 | |

- 4.4.3, 4-32 In Table 4-27, “Configurable Clock Units,” update default frequency of PCI Express, SATA, and TDM. In addition, remove *csb_clk/2* and *csb_clk/3* options from PCI Express, as follows:

Table 4-27. Configurable Clock Units

| Unit | Default Frequency | Options |
|---------------------------------|-------------------|---|
| eTSEC1 and eTSEC2 | <i>csb_clk</i> | Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i> |
| Security core, I ² C | <i>csb_clk</i> | Off, <i>csb_clk/2</i> , <i>csb_clk/3</i> |
| USB DR | <i>csb_clk</i> | Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i> |
| PCI and DMA complex | <i>csb_clk</i> | Off, <i>csb_clk</i> |
| PCI Express 1 and PCI Express 2 | <i>csb_clk</i> | Off, <i>csb_clk</i> |
| SATA1/SATA2 | <i>csb_clk</i> | Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i> |
| TDM | <i>csb_clk</i> | Off, <i>csb_clk</i> , <i>csb_clk/2</i> , <i>csb_clk/3</i> |

- 4.5.1.6, 4-37 Change the first sentence to the following:
“RCR, shown in Figure 4-11, can be used by software to initiate a hard reset sequence.”
- 4.5.2.3, 4-41 In Table 4-37, “SCCR Bit Descriptions,” add the following to bit field description of ENCCM (bits 6–7):
The encryption core must have the same clock ratio as the USB unit, unless one of them has its clock disabled.
- 5.2.4.4.1, 5-10 Update Table 5-10, “LBLAWAR0[EN] Reset Value,” by adding the RCWHR[RLEXT] value in the RCWHR[ROMLOC] column.

| RCWHR[RLEXT]/ RCWHR[ROMLOC] | LBLAWAR0[EN] Reset Value | Description |
|--------------------------------|-----------------------------|--|
| 00 / 000–100 | 0 | e300c3 core boot not performed from a local bus device. |
| 00 / 101–111 | 1 | e300c3 core boot performed from a local bus device. Local bus 8-Mbyte ($2^{(22+1)}$) local access window is enabled. |

Revision History

5.2.4.6.1, 5-12 Update Table 5-14, by adding the RCWHR[RLEXT] value in the RCWHR[ROMLOC] column as follows:

| RCWHR[RLEXT]/ RCWHR[ROMLOC] | PCILAWR0[EN] Reset Value | Description |
|--------------------------------|-----------------------------|--|
| 00 / 000, 011–111 | 0 | e300c3 core boot not performed from a PCI device. |
| 01–11 / 000–111 | 0 | e300c3 core boot not performed from a PCI device. |
| 00 / 001 | 1 | e300c3 core boot performed from a PCI device. PCI 8-Mbyte ($2^{(22+1)}$) local access window is enabled. |

5.2.4.12.1, 5-16 Update Table 5-22, “DDRLAWAR0[EN] Reset Value,” by adding the RCWHR[RLEXT] value in the RCWHR[ROMLOC] column.

| RCWHR[RLEXT]/ RCWHR[ROMLOC] | DDRLAWAR0[EN] Reset Value | Description |
|--------------------------------|------------------------------|--|
| 00 / 000 | 1 | e300c3 core boot performed from a DDR SDRAM device. DDR 8-Mbyte ($2^{(22+1)}$) local access window is enabled. |
| Else | 0 | e300c3 core boot not performed from a DDR SDRAM device. |

5.3.1, 5-19 In Table 5-24, “System Configuration Register Memory Map,” update SICRL reset value to “0xF00F00_n0n0.”

In addition, update DDRCDR and DDRDSR reset values to “0x7B84_0001,” and “0x3B80_0000,” respectively.

5.3.2.4, 5-22 In Figure 5-17, “System Priority Configuration Register (SPCR),” and Table 5-30, “SPCR Bit Settings,” make bits TSECDP and TSECBDP (bits 20–23) Reserved.

5.3.2.5, 5-24 In Figure 5-18, “System I/O Configuration Register Low (SICRL),” change reset value, as follows:

Offset 0x00114 Access: Read/Write

| | | | | | | | | |
|-------|--------------------------|------------------------------|----------|----|------------------------|----|--------------------------|--------------------------|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| R | DMA_CH0 | | SPI | | UART | | $\overline{\text{IRQ4}}$ | $\overline{\text{IRQ5}}$ |
| W | | | | | | | | |
| Reset | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 |
| | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
| R | $\overline{\text{IRQ5}}$ | $\overline{\text{IRQ}}[6-7]$ | IIC1 | | TDM | | TDM_SHARED | |
| W | | | | | | | | |
| Reset | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 |
| | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 |
| R | PCI_A | | ELBC_A | | — | | | |
| W | | | | | | | | |
| Reset | All zeros | | | | | | | |
| | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 |
| R | ETSEC1_A ¹ | | ETSEC1_B | | ETSEC1_C ¹ | | — | TSEXPOBI |
| W | | | | | | | | |
| Reset | depends on RCWH[16-18] | | 0 | 0 | depends on RCWH[19-21] | | 0 | 0 |

¹ The SICRL[ETSEC1_A] depends on the value of RCWH[16-18] and SICRL[ETSEC1_C] reset value depends on the value of RCWH[19-21] which is loaded into the reset configuration word.

Figure 5-18. System I/O Configuration Register Low (SICRL)

In addition, correct the corresponding figure footnote from:
 “The SICRL[ETSEC1_A] and SICRL[ETSEC1_C] reset value depends on the value of RCWH[30] which is loaded into the reset configuration word.”
 to:
 “The SICRL[ETSEC1_A] depends on the value of RCWH[16-18] and SICRL[ETSEC1_C] reset value depends on the value of RCWH[19-21] which is loaded into the reset configuration word.”

5.3.2.5, 5-25 In Table 5-31 “SICRL Bit Settings,” correct duplicate instance of “UART_CTS2”—change first instance (muxed with LDVAL and MDVAL) to “UART_SIN2” as follows:

| | | |
|-----------|-------|-------|
| UART_SIN2 | MDVAL | LDVAL |
|-----------|-------|-------|

5.3.2.5, 5-26 Modify the Table 5-73, “SICRL Bit Settings” by replacing “bits 26–27” with “bit 27” and removing the footnote corresponding to bit 26 (against E_TSEC1B) to attach it to E_TSEC1B corresponding to bit 27 as follows:

| SICRL[Bits] Value | | 0b00 | 0b01 | 0b10 | 0b11 | Reset |
|-------------------|-----------------------|----------------|----------------|----------------|----------------|-------|
| Bits | Group | Pin Function 0 | Pin Function 1 | Pin Function 2 | Pin Function 3 | |
| 27 | ETSEC1_B ¹ | TSEC1_TX_CLK | USBDR_CLK | | | 00 |
| | | TSEC1_RXD3 | USBDR_TXDRXD5 | TSEC_TMR_CLK | | |
| | | TSEC1_RXD2 | USBDR_TXDRXD6 | TSEC_TMR_TRIG1 | | |
| | | TSEC1_RXD1 | USBDR_TXDRXD7 | TSEC_TMR_TRIG2 | | |
| | | TSEC1_RXD0 | USBDR_NXT | | | |
| | | TSEC1_RX_ER | USBDR_DIR | | | |

¹ TSEC1_TX_CLK is selected when SICRL[27] is logic 0 irrespective of SICRL[26].

5.3.2.6, 5-27 In second paragraph, change “A value of 0b11 is illegal for all groups,” to “A value of 0b11 selects GPIO mode of the appropriate pin.”

5.3.2.6, 5-27 Modify Table 5-32, “SICRH Bit Settings,” as follows:

Table 5-32. SICRH Bit Settings

| SICRH[Bits] Value | | 0b00 | 0b01 | 0b10 | 0b11 | Reset Value |
|-------------------|----------|--------------------------------|---|---------------------------------|----------------|-------------|
| Bits | Group | Pin Function 0 | Pin Function 1 | Pin Function 2 | Pin Function 3 | |
| 0–1 | GPIO_0 | $\overline{\text{DMA_DREQ1}}$ | $\overline{\text{GTM1_TOUT1}}$ | — | GPIO_0 | 11 |
| 2–3 | GPIO_1 | $\overline{\text{DMA_DACK1}}$ | GTM1_TIN2/GTM2_TIN1 | — | GPIO_1 | 11 |
| 4–5 | GPIO_2 | DMA_DONE1 | $\overline{\text{GTM1_TGATE2}}/$ $\overline{\text{GTM2_TGATE1}}$ | — | GPIO_2 | 11 |
| 6–7 | GPIO_3 | — | GTM1_TIN3/GTM2_TIN4 | — | GPIO_3 | 11 |
| 8–9 | GPIO_4 | — | $\overline{\text{GTM1_TGATE3}}/$ $\overline{\text{GTM2_TGATE4}}$ | — | GPIO_4 | 11 |
| 10–11 | GPIO_5 | — | $\overline{\text{GTM1_TOUT3}}$ | $\overline{\text{GTM2_TOUT1}}$ | GPIO_5 | 11 |
| 12–13 | GPIO_6 | — | GTM1_TIN4/GTM2_TIN3 | — | GPIO_6 | 11 |
| 14–15 | GPIO_7 | — | $\overline{\text{GTM1_TGATE4}}/$ $\overline{\text{GTM2_TGATE3}}$ | — | GPIO_7 | 11 |
| 16–17 | GPIO_8 | USBDR_DRIVE_VBUS | GTM1_TIN1/GTM2_TIN2 | — | GPIO_8 | 11 |
| 18–19 | GPIO_9 | USBDR_PWRFAULT | $\overline{\text{GTM1_TGATE1}}/$ $\overline{\text{GTM2_TGATE2}}$ | — | GPIO_9 | 11 |
| 20–21 | GPIO_10 | USBDR_PCTL0 | $\overline{\text{GTM1_TOUT2}}$ | $\overline{\text{GTM2_TOUT1}}$ | GPIO_10 | 11 |
| 22–23 | GPIO_11 | USBDR_PCTL1 | $\overline{\text{GTM1_TOUT4}}$ | $\overline{\text{GTM2_TOUT3}}$ | GPIO_11 | 11 |
| 24–25 | ETSEC2_A | TSEC2_COL | — | — | GPIO_26 | 11 |
| | | TSEC2_CRS | — | — | GPIO_27 | |
| 26 | — | | | | | 0 |

- WDT disabling
If the software watchdog timer is not needed, the user must clear SWCRR[SWEN] bit to disable the WDT not later than its timer times out (~32.2 sec. for a 133-MHz system clock).
 - WDT initial servicing
If the software watchdog timer is to be used, the special service sequence, described in [Section 5.5.5.1, “Software Watchdog Timer Unit,”](#) must be executed after system reset and not later than the first WDT time-out (~32.2 sec. for a 133-MHz system clock).
Subsequently, periodical WDT servicing should be performed according to the programming guidelines given in [Section 5.5.5.1, “Software Watchdog Timer Unit.”](#)
- 5.5.5.5, 5-50 In Table 5-51, “RTEVR Bit Settings,” change AIF (bit 30) field description to the following:
“Alarm interrupt flag bit.
Used to indicate the alarm interrupt. The bit is set if the RTC issues an interrupt when the RTC counter value equals RTALR[ALRM].”
- 5.6.5.5, 5-57 In Table 5-60, “PTEVR Bit Settings,” modify PIF (bit) field description to say the following:
“Periodic interrupt flag bit. It is asserted after the SPMPIT counter counts to zero. This status bit should be cleared by software.”
- 5.7.2, 5-59 Change section title to “GTM Features,” and modify, as follows:
- The maximum input clock is the system bus clock
 - Four 16-bit programmable timers
 - Two timers cascaded internally or externally to form a 32-bit timer
 - One timer cascaded internally or externally to form a 64-bit timer
 - Maximum period of ~515 seconds (at 133-MHz bus clock in slow go mode, primary and secondary prescaler = 256) for 16-bit timer
 - Maximum period of ~8246 seconds (at 133-MHz bus clock and prescaler = 256) for 32-bit timer
 - Maximum period of thousands of years (at 133-MHz bus clock and prescaler = 256) for 64-bit timer
 - 7.5-nanosecond timer resolution (at 133-MHz bus clock and no prescaler)
 - Resolution and maximum period can be traded off by selecting prescaler divisor
 - Three programmable input clock sources for the timer prescalers
 - Input capture capability
 - Output compare with programmable mode for the output pin
 - Free run and restart modes
 - Functional and programming compatibility with MPC8260 timers

- 5.7.3, 5-60 Change name of section to “GTM Modes of Operation,” and modify as follows:
The GTM unit can operate in the following modes:
- Cascaded modes
 - Clock source modes
 - Reference modes
 - Capture modes
- 5.7.4.2, 5-62 In Table 5-62, “GTM External Signals—Detailed Signal Descriptions,” in the State Meaning description of signal TGATE_n, change “In a reset gate mode...” to “In a restart gate mode...”.
- In addition, in the State Meaning description of signal TOUT_n, clarify #2 by changing “TOUT_n changes occur on the rising edge of the system clock” to “TOUT_n begins or stops counting, depending on the signal state and the configured mode.”; also clarify Timing by changing “system clock” to “timer input clock.”
- 5.7.6.1, 5-71 In the second paragraph, replace “65,537,” with “65,536.”
- 5.7.6.1, 5-71 Change the last sentence to the following:
“The maximum period (when the reference value is all ones and the prescaler divides by 256) for one 16-bit timer is ~50 ms at 333 MHz.”
- 5.7.6.1, 5-71 Update the sentence in the third paragraph from:
“The best resolution of the timer is one clock cycle (3 ns at a 333-MHz system clock, for example). The maximum period (when the reference value is all ones and the prescaler divides by 256) for one 16-bit timer is ~50 ms at 333 MHz
to:
“The best resolution of the timer is one clock cycle (7.5 ns at 133 MHz system clock, for example). The maximum period (when the reference value is all ones and the prescaler divides by 256) for one 16-bit timer is ~515 s at 133 MHz.”
- 5.7.6.3, 5-72 In the Note, change “the counter begins counting after...” to “the counter begins or stops after...”
- 5.8.2.2, 5-78 In Table 5-76, “PM CER Bit Settings,” add the following to the note in PCI (PME), USB, and eTSEC1 field descriptions: “Also, this bit is not used for wake up from D3 warm.”
- 5.8.2.4, 5-81 In Table 5-78, “PMCCR1 Bit Settings,” add bit field description for ASSERT_PME (bit 25), as follows:
“Normally $\overline{\text{PCI_PME}}$ output is asserted automatically by PMC when a defined wake-up event occurs assuming PMCCR1[PME_EN] = 1 and PCIPMCR1[PME_EN] = 1. A defined wake-up event refers to those events that are registered in bits 23–30 of PM CER. ASSERT_PME allows $\overline{\text{PCI_PME}}$ to be asserted manually, by the e300. This would be done to inform the host of a power state change when PMC does not generate $\overline{\text{PCI_PME}}$ automatically, for example waking from D1 due to CSB bus activity. ASSERT_PME is not qualified by

PMCCR1[PME_EN]. ASSERT_PME is cleared when the $\overline{\text{PCI_PME}}$ signal is asserted.

0 $\overline{\text{PCI_PME}}$ will only be asserted automatically when a defined wake-up event occurs.

1 Assert the $\overline{\text{PCI_PME}}$ signal under software control (manually)."

5.8.3, 5-84

Change bullet "External interrupt" to "IRQ1, IRQ2."

In addition, change bullet "Shut-off the I/Os (for example, DDR, PCI, eLBC, eTSEC2) which are not required by the wake-up logic" to the following: "Shut-off the I/Os (for example, DDR, PCI, eLBC, eTSEC1), which are not required by the wake-up logic."

5.8.3, 5-85

Modify Table 5-80, "MPC8315E Power Rail Nomenclature," as follows:

Table 5-80. MPC8315E Power Rail Nomenclature

| Power Plane | Description |
|-------------|---|
| VDDC | This power plane supplies current in all the power states. The power plane is used to supply the logic needed for wake-up from the new low-power down mode, which include ETSEC2, GPIO, PMC and other related logical blocks. |
| VDD | This supply is switched off in D3Warm state. |

5.8.3, 5-85

In Table 5-81, "Software-Controller Power-Down States—Basic Description," change references to "PMCCR1[FULL_POWER_OFF]" to say "PMCCR1[POWER_OFF]."

5.8.3.7, 5-92

Add the following text to the first paragraph:

"In the D3 Warm state, registers of the reset ([Section 4.5.1, "Reset Configuration Register Descriptions"](#)), clock ([Section 4.5.2, "Clock Configuration Registers"](#)), local memory map ([Section 5.2.3.1, "Local Access Register Memory Map"](#)), system configuration ([Section 5.3.1, "System Configuration Register Memory Map"](#)), GTM ([Section 5.7.5, "GTM Memory Map/Register Definition"](#)), PMC ([Section 5.8.2, "PMC Memory Map/Register Definition"](#)), eTSEC1, eTSEC2 ([Section 19.5.2, "Detailed Memory Map"](#)), and GPIO ([Section 24.3, "Memory Map/Register Definition"](#)) are powered ON. All the registers of these blocks remain intact during D3 Warm state."

5.8.3.7.1, 5-96

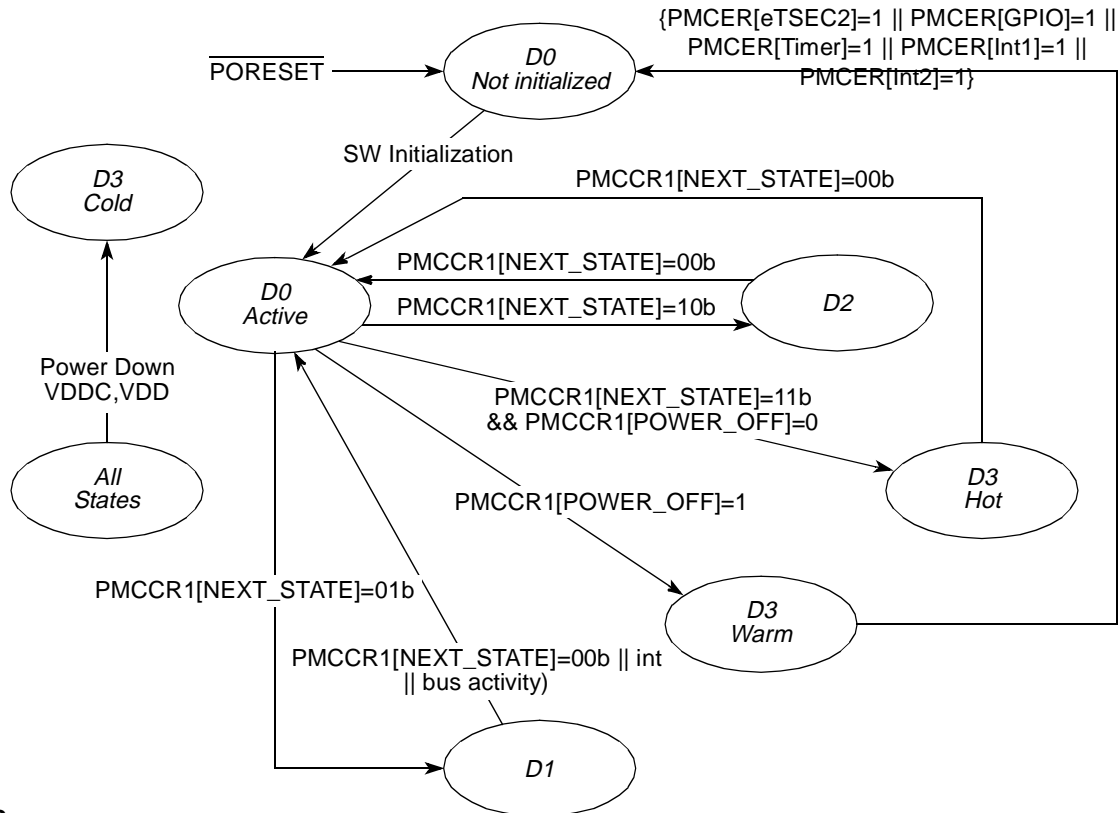
In Figure 5-63, "Power State Transitions Supported," change references to "PMCCR1[FULL_POWER_OFF]" to say "PMCCR1[POWER_OFF]."

In addition, update "{PMCCR[eTSEC1]=1b || PMCCR[GPIO]=1b || PMCR[Timer]=1 || PMCR[Interrupt]=1b}"

to:

"{PMCER[eTSEC2]=1 || PMCER[GPIO]=1 || PMCER[Timer]=1 || PMCER[Int1]=1 || PMCER[Int2]=1}"

as follows:


NOTES:

- 1) Wake-up case #1: Wake-up will occur as a result of one of the following: e300 interrupt or bus activity, $PM\text{CER}[X]$ & $PM\text{CMR}[X] = 1$, where X is one of {USB, TSEC, GPIO, PCI, interrupt, timer}, $PM\text{CCR1}[\text{NEXT_STATE}] = 00$. The device will wake-up directly as host or agent. PCI_PME can be generated manually to host if desired.
- 2) Wake-up case #2: As host ($PM\text{CCR}[\text{PME_EN}] = 0$), Wake-up will occur as a result of one of the following: $PM\text{CER}[X]$ & $PM\text{CMR}[X] = 1$, where X is one of {USB, TSEC, GPIO, PCI, interrupt, timer}. As agent ($PM\text{CCR}[\text{PME_EN}] = 1$) wake-up event will cause PCI_PME to be asserted; the device will wake up when the host sets $PCIPMR1[\text{Power_State}] = 00$, which causes $PM\text{CCR1}[\text{NEXT_STATE}] = 00$.
- 3) Before the transition to D3Warm, VDD power is switched off. VDDC remains constant.
- 4) Before the transition from D3Warm, VDD power is switched on.
- 5) D1, D2 & D3Hot modes are supported through use of the e300 Doze, Nap, and Sleep modes respectively. Transitions from D1–D3Hot are triggered as follows: the e300 internal time base unit invokes a request to exit low-power state or e300 receives an interrupt request, including interrupt from a defined wake-up event.
- 6) Transitions to D3Cold implies power-down of VDDC and VDD supplies.
- 7) Transitions from D3Warm go to the D0-non-initialized state, meaning a portion of the chip (e300, DDR and so on.) will need to be initialized. Portions of the chip which are not powered-off may not need to be initialized.

5.8.3.7.2, 5-100

Under the section titled, “PMC Wake-Up When the MPC8315E is a PCI Host,” modify step 3, as follows:

“3. The PMC then waits for the PMC_PWR_OK signal to be asserted.

PMC_PWR_OK is an indication that the external VDD supply is stable and within spec. The PMC_PWR_OK signal is an input to the PMC that can be supplied from an external source, or can be tied active internally.”

5.8.3.7.2, 5-100

Under the section titled, “PMC Wake-Up When the MPC8315E is a PCI Host,” replace the eighth sentence in step 5 with the following:

“As a guideline for determining the required length of the PMC reset, the core PLL takes 100 ms to lock (as mentioned in Table, “PLL Lock times,” in MPC8315E PowerQUICC™ II Pro Processor Hardware Specifications) after PMC_PWR_OK is asserted.”

5.8.3.7.2, 5-101 Under the section titled, “PMC Wake-Up When the MPC8315E is a PCI Host,” in Figure 5-66, “PMC Wake-Up Sequence from D3Warm,” change “pwr_en signal asserted to switch on the VDD power supply,” to “XT_PWR_CTRL signal asserted to switch on the VDD power supply.”

5.8.3.7.2, 5-102 Under the section titled, “PMC Wake-Up When the MPC8315E is a PCI Host,” in Figure 5-67, “PMC Wake-Up Sequence from D3Warm (continued),” remove “Initialization sequence based on check of PMCCR[WARM_RESET] control bit.”

5.8.3.7.4, 5-104 Modify the section titled, “PMC_PWR_OK Signal,” as follows:
 “PMC_PWR_OK is an external input indication that the VDD that was switched off in the device D3Warm mode has returned to specified levels after a wake-up event occurs. If an external power switch device is used (not a FET), it will typically provide such a signal. When a wake-up event occurs and PMC asserts the EXT_PWR_CTRL signal to turn on power, it will wait until the PMC_PWR_OK signal is asserted before it will proceed to wait for the e300 pll to lock.”

6.2, 6-2 In Table 6-1, “Arbiter Register Map,” add row for Arbiter Event Enable Register as follows:

| | | | | |
|------|--------------------------------------|-----|-------------|---------------------------|
| 0x08 | Arbiter Event Enable Register (AEER) | R/W | 0x0000_007F | 6.2.3/6-5 |
|------|--------------------------------------|-----|-------------|---------------------------|

6.2.1, 6-3 In Table 6-2, “ACR Field Descriptions,” change reserved field descriptions from “Write reserved, read = 0” to “Reserved, write should preserve reset value.”

6.2.3, 6-5 Add Section 6.2.3, “Arbiter Event Enable Register (AEER),” as follows:

6.2.3 Arbiter Event Enable Register (AEER)

Arbiter event enable register (AEER) specifies, which kind of events are considered as error events. If event is defined as non-error event, it also won't be reported neither in event register nor in event attributes and address registers. For transfer types, that are not defined as error events, arbiter also does not end address/data tenures. [Figure 6-3](#) shows the fields of AEER.

Offset 0x08

Access: Read/Write




Figure 6-3. Arbiter Event Enable Register (AEER)

[Table 6-4](#) defines the bit fields of AEER.

Table 6-4. AEER Bit Settings

| Bits | Name | Description |
|------|------|--|
| 0–25 | — | Write reserved, read = 0 |
| 26 | ETEA | External $\overline{\text{TEA}}$. Specifies, whether assertion of $\overline{\text{TEA}}$ signal by one of the slaves is going be reported in arbiter event registers. 0 Assertion of $\overline{\text{TEA}}$ signal by one of the slaves is not reported in arbiter event registers. 1 Assertion of $\overline{\text{TEA}}$ signal by one of the slaves is reported in arbiter event registers. |
| 27 | RES | Reserved transfer type. Specifies, whether transaction with reserved transfer type will be reported in arbiter event registers. 0 Reserved transaction isn't reported in arbiter event registers. 1 Reserved transaction is reported in arbiter event registers. |
| 28 | ECW | External Control Word transfer type. Specifies, whether transaction with external control word transfer type will be reported in arbiter event registers. 0 External control word read/write transaction isn't reported in arbiter event registers. 1 External control word read/write transaction is reported in arbiter event registers. |
| 29 | AO | Address Only transfer type. Specifies, whether transaction with address only transfer type will be reported in arbiter event registers. 0 Address only transaction isn't reported in arbiter event registers. 1 Address only transaction is reported in arbiter event registers. |

Table 6-4. AEER Bit Settings

| Bits | Name | Description |
|------|------|--|
| 30 | DTO | <p>DTO - Data Time Out.</p> <p>Specifies, whether data tenure time out will be reported in arbiter event registers.</p> <p>0 Data time out isn't reported in arbiter event registers.</p> <p>1 Data time out is reported in arbiter event registers.</p> |
| 31 | ATO | <p>ATO - Address Time Out.</p> <p>Specifies, whether address tenure time out will be reported in arbiter event registers.</p> <p>0 Address time out isn't reported in arbiter event registers.</p> <p>1 Address time out is reported in arbiter event registers.</p> |

- 6.2.6, 6-8 Correct the access mentioned in Figure 6-6, “Arbiter Event Attributes Register (AEATR)” from “Read/write” to “Read-only”.
- 6.2.7, 6-10 Correct the access mentioned in Figure 6-7, “Arbiter Event Address Register (AEADR),” from “Read/write” to “Read-only”.
- 6.3.1.3, 6-14 Change “After the completion of snoop copyback, the arbiter grants the bus back to the master that had its transaction ARTRYed,”
to:
“After the completion of snoop copyback, the arbiter grants the bus to the most ahead master among those masters which have an active bus request signal at that time, which may or may not be the same master that had its transaction ARTRYed. Only when a transaction address phase is completed with no ARTRY (and no repeat conditions), the master moves to the end of the line.”
- 7.2, 7-13 In Table 7-2, “Device Revision Level Cross-Reference,” change all instances of “MPC315” and “MPC314” to “MPC8315” and “MPC8314,” respectively.
Table 7-2 now reads as follows:

| MPC8315E Revision | Processor Version Register (PVR) | System Version Register (SVR) |
|-------------------|----------------------------------|---|
| 1.0 | 8085_0020 | <p>MPC8315E (with security): 0x80B4_0010</p> <p>MPC8315 (without security): 0x80B5_0010</p> <p>MPC8314E (with security): 0x80B6_0010</p> <p>MPC8314 (without security): 0x80B7_0010</p> |
| 1.1 | 8085_0020 | <p>MPC8315E (with security): 0x80B4_0011</p> <p>MPC8315 (without security): 0x80B5_0011</p> <p>MPC8314E (with security): 0x80B6_0011</p> <p>MPC8314 (without security): 0x80B7_0011</p> |

- 7.4.1.3.3, 7-21 In Table 7-3, “e300 HID0 Bit Descriptions,” add the following note to EBA and EBD field descriptions: “Do not set this bit; the CSB does not have parity signals.”
- 8.4.2, 8-5 In Table 8-2, “IPIC External Signals—Detailed Signal Descriptions,” modify IRQ[0:7] State Meaning asserted description, as follows:
“Asserted—When an external interrupt request signal is asserted, the priority is checked by the IPIC unit, and the interrupt is conditionally passed to the processor.”

- 8.5.13, 8-24 In Figure 8-16, “System External Interrupt Mask Register (SEMSR),” change the following text appearing under the reset values of 0–15 bits from:
 “The reset values of implemented bits reflect the values of the external IRQ signals. Reserved bits are zeros.”
 to “All zeros”.
- In addition, remove the following second figure footnote:
 “The user should drive all IRQ inputs to an inactive state prior to reset negation”
- 9.3.2.1, 9-6 In Table 9-3, “Memory Interface Signals—Detailed Signal Descriptions,” change signal description of MA[14:0] from:
 “Assertion/Negation—The address is always driven when the memory controller is enabled. It is valid when a transaction is driven to DRAM (when \overline{MCSn} is active).”
 to:
 “Assertion/Negation—The address lines are only driven when the controller has a command scheduled to issue on the address/CMD bus; otherwise they will be at high-Z. It is valid when a transaction is driven to DRAM (when \overline{MCSn} is active).”
- 9.3.2.2, 9-8 In Table 9-4, “Clock Signals—Detailed Signal Descriptions,” modify MCKE description, as follows:

Table 9-4. Clock Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|--------|-----|--|
| MCKE | O | Clock enable. Output signals used as the clock enables to the SDRAM. MCKE can be negated to stop clocking the DDR SDRAM. The MCKE signals should be connected to the same rank of memory as the corresponding MCS and MODT signals. For example, MCKE[0] should be connected to the same rank of memory as MCS[0] and MODT[0]. |
| | | State Meaning Asserted—Clocking to the SDRAM is enabled. Negated—Clocking to the SDRAM is disabled and the SDRAM should ignore signal transitions on MCK or MCK. MCK/MCK are don't cares while MCKE is negated. |
| | | Timing Assertion/Negation—Asserted when DDR_SDRAM_CFG[MEM_EN] is set. Can be negated when entering dynamic power management or self refresh. Are asserted again when exiting dynamic power management or self refresh. High impedance—Always driven. |

- 9.4, 9-9 In Table 9-5, “DDR Memory Controller Memory Map,” add the following registers. Note that only the affected rows are shown:

Table 9-5. DDR Memory Controller Memory Map

| Offset | Register | Access | Reset | Section/Page |
|--------|--|--------|-----------|------------------------------|
| 0xE40 | ERR_DETECT—Memory error detect | w1c | All zeros | 9.4.1.17/-57 |
| 0xE44 | ERR_DISABLE—Memory error disable | R/W | All zeros | 9.4.1.18/-57 |
| 0xE48 | ERR_INT_EN—Memory error interrupt enable | R/W | All zeros | 9.4.1.19/-58 |

Revision History

- 9.4.1.7, 9-18 In Figure 9-8, “DDR SDRAM Control Configuration Register (DDR_SDRAM_CFG),” and Table 9-12, “DDR_SDRAM_CFG Field Descriptions,” change x32_EN (bit 26) to Reserved.
- 9.4.1.7, 9-19 In Table 9-12, “DDR_SDRAM_CFG Field Descriptions,” in 8_BE field description, modify the following note:
 “**Note:** DDR1 (SDRAM_TYPE = 010) must use 8-beat bursts when using 32-bit bus mode (32_BE = 1) and 4-beat bursts when using 64-bit bus mode; DDR2 (SDRAM_TYPE = 011) must use 4-beat bursts, even when using 32-bit bus mode”
- 9.4.1.7, 9-20 In Table 9-12, “DDR_SDRAM_CFG Field Descriptions,” modify HSE field description, as follows:
 “Global half-strength override
 Sets I/O driver impedance to half strength. This impedance is used by the address/command, data, and clock impedance values, but only if automatic hardware calibration is disabled and the corresponding group’s software override is disabled in the DDR control driver register(s) described in [Section 5.3.2.9](#), “[DDR Control Driver Register \(DDRCDR\)](#).” This bit should be cleared if using automatic hardware calibration.
 0 I/O driver impedance is configured to full strength.
 1 I/O driver impedance is configured to half strength.”
- 9.4.1.17, 9-29 Add the sections for registers ERR_DETECT, ERR_DISABLE, and ERR_INT_EN, as follows. Note that these sections are added after Section 9.4.1.17 and prior to Section 9.5.

9.4.1.18 Memory Error Detect (ERR_DETECT)

The memory error detect register stores the detection bits for memory select errors. It is a read/write register. A bit can be cleared by writing a one to the bit. System software can determine the type of memory error by examining the contents of this register. If an error is disabled with ERR_DISABLE, the corresponding error is never detected or captured in ERR_DETECT.

ERR_DETECT is shown in [Figure 9-19](#).




Figure 9-19. Memory Error Detect Register (ERR_DETECT)

Table 9-24 describes the ERR_DETECT fields.

Table 9-24. ERR_DETECT Field Descriptions

| Bits | Name | Description |
|------|------|---|
| 0–30 | — | Reserved |
| 31 | MSE | Memory select error. This bit is cleared by software writing a 1. 0 A memory select error has not been detected. 1 A memory select error has been detected. |

9.4.1.19 Memory Error Disable (ERR_DISABLE)

The memory error disable register, shown in Figure 9-20, allows selective disabling of the DDR controller’s error detection circuitry. Disabled errors are not detected or reported.




Figure 9-20. Memory Error Disable Register (ERR_DISABLE)

Table 9-25 describes the ERR_DISABLE fields.

Table 9-25. ERR_DISABLE Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–30 | — | Reserved |
| 31 | MSED | Memory select error disable 0 Memory select errors are enabled. 1 Memory select errors are disabled. |

9.4.1.20 Memory Error Interrupt Enable (ERR_INT_EN)

The memory error interrupt enable register, shown in Figure 9-21, enables memory select error interrupts. When an enabled interrupt condition occurs, the internal \overline{int} signal is asserted to the programmable interrupt controller (PIC).




Figure 9-21. Memory Error Interrupt Enable Register (ERR_INT_EN)

Table 9-26 describes the ERR_INT_EN fields.

Table 9-26. ERR_INT_EN Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–30 | — | Reserved |
| 31 | MSEE | Memory select error interrupt enable 0 Memory select errors do not cause interrupts. 1 Memory select errors generate interrupts. |

9.5, 9-30

Correct the following sentence in third paragraph:

“Bank sizes up to 2 Gbytes (maximum total physical memory size of 4 Gbytes) are supported, providing up to a maximum of 4 Gbytes of DDR main memory.”
to:

“Bank sizes up to 2 Gbits (maximum total physical memory size of 4 Gbytes) are supported, providing up to a maximum of 4 Gbits of DDR main memory per chip select.”

9.5.2, 9-36

Change title of Table 9-27 to “DDR1 Address Multiplexing for 16-Bit Data Bus,” and change title of Table 9-28 to “DDR2 Address Multiplexing for 16-Bit Data Bus.” Update these tables as follows:

Table 9-27. DDR1 Address Multiplexing for 16-Bit Data Bus

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | | |
|-------------------|--------------------------|--------------------------|---|---|----|----|----|----|----|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|----|--|--|
| | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | | 30 | 31 | | |
| 15 x 11 x 2 | $\overline{\text{MRAS}}$ | | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 15 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 14 x 11 x 2 | $\overline{\text{MRAS}}$ | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 14 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 13 x 11 x 2 | $\overline{\text{MRAS}}$ | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 11 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | |

Table 9-27. DDR1 Address Multiplexing for 16-Bit Data Bus (continued)

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | |
|-------------------|--------------------------|--------------------------|---|---|---|---|---|----|----|----|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|----|--|
| | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | | 30 | 31 | |
| 13 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 13 x 9 x 2 | $\overline{\text{MRAS}}$ | | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 12 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 12 x 9 x 2 | $\overline{\text{MRAS}}$ | | | | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |
| 12 x 8 x 2 | $\overline{\text{MRAS}}$ | | | | | | | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | |

Table 9-28. DDR2 Address Multiplexing for 16-Bit Data Bus

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | |
|-------------------|--------------------------|--------------------------|---|---|----|----|----|----|----|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|----|--|
| | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | | 30 | 31 | |
| 15 x 10 x 3 | $\overline{\text{MRAS}}$ | | | | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 14 x 10 x 3 | $\overline{\text{MRAS}}$ | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 14 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 13 x 10 x 3 | $\overline{\text{MRAS}}$ | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 2 | 1 | 0 | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |

Table 9-28. DDR2 Address Multiplexing for 16-Bit Data Bus

| Row x Col | msb | Address from Core Master | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | lsb | | | |
|-------------------|--------------------------|--------------------------|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|--|--|--|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | | | |
| 13 x 10 x 2 | $\overline{\text{MRAS}}$ | | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |
| 13 x 9 x 2 | $\overline{\text{MRAS}}$ | | | | | | | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | | | | | | | | | | | | | |
| | MBA | | | | | | | | | | | | | | | | | | | | 1 | 0 | | | | | | | | | | | | | |
| | $\overline{\text{MCAS}}$ | | | | | | | | | | | | | | | | | | | | | | | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | | | |

9.5.6, 9-44 Add the following note to the end of the section:

NOTE

Application system board must assert the reset signal on DDR memory devices until software is able to program the DDR memory controller configuration registers, and must deassert the reset signal on DDR memory devices before DDR_SDRAM_CFG[MEM_EN] is set. This ensures that the DDR memory devices are held in reset until a stable clock is provided and, further, that a stable clock is provided before memory devices are released from reset.

9.5.9, 9-50 Modify the last sentence in the first paragraph as follows:

For example, if a write transaction is desired with a size of one word (4 bytes), then the second, third, and fourth beats of data are not written to DRAM, as the width of the data bus is 32 bits.

9.6.1, 9-52 In Table 9-53, “Programming Differences Between Memory Types,” for ODT_PD_EXIT, change it to be set to 0001 for DDR1; for FOUR_ACT, change it to be set for 00001 for DDR1.

9.6.3.1, 9-55 Remove Section 9.6.3.1, “Hardware Based Self-Refresh.”

Chapter 10 Remove references to PLL feature.

Chapter 10 Remove references to atomic operations.

Chapter 10 Update “snoop for CSB transactions” to “No snoop for CSB transactions”

10.1.2, 10-4 Add the following bullet item to the eLBC main features:

Different machines (FCM/GPCM/UPM) share the address, data, and control signals. While the eLBC is servicing a transaction, subsequent transactions are queued until the current transaction has completed

10.1.3, 10-4 Add the following sentence after the first sentence:

“The internal transaction address is limited to 32 bits, so all chip selects must fall within the 4-Gbyte window addressed by the internal transaction address.”

- 10.2, 10-7 In Table 10-2, “Enhanced Local Bus Controller Detailed Signal Descriptions” change the State Meaning description for LA[16:25] to the following:
 “LA is the address bus used to transmit addresses to external RAM devices. Refer to [Section 10.5, “Initialization/Application Information,”](#) for address signal multiplexing.”
 In addition, modify the following sentence for LGPL_n descriptions, as follows: “it is driven with a value programmed into the UPM array.”
- 10.3.1.1, 10-11 In Figure 10-2, “Base Registers (BR_n),” change the reset value of reserved bit 30 from 1 to 0.
- 10.3.1.7, 10-25 In first paragraph, change section beginning with “To avoid race conditions...” to the end of the paragraph to say the following:
 “To avoid race conditions between software and a busy eLBC, registers that affect currently running special operation and LSOR must not be re-written before a pending special operation has been completed. The UPM and FCM have different indications of when such special operations are completed. The behavior of eLBC is unpredictable if special operation modes are altered between LSOR being written and the relevant memory controller completing that access.”
- 10.3.1.12, 10-30 Add the following note prior to Figure 10-16, “Transfer Error Attributes Register (LTEATR)”
 “LTEATR may not capture accurate information for errors that occur when an FCM special operation is in progress.”
- 10.3.1.12, 10-30 In Table 10-20, “LTEAR Field Descriptions,” add the following text to SRCID (bits 11–15) field description:
 “The coding of the source ID debug information is the same as the coding of AEATR[MSTR_ID] (see [Section 6.2.6, ‘Arbiter Event Attributes Register \(AEATR\).’](#))”
- 10.3.1.13, 10-31 Add the following note prior to Figure 10-17, “Transfer Error Address Register (LTEAR).”
 “LTEAR may not capture accurate information for errors that occur when an FCM special operation is in progress.”
- 10.3.1.14, 10-32 In Table 10-21, “LBCR Field Descriptions,” update AHD (bit 10) field description to say the following:
 “Address hold disable. Removes part of the hold time for LAD with respect to LALE in order to lengthen the LALE pulse
- 0 During address phases on the local bus, the LALE signal negates one platform clock period prior to the address being invalidated. For instance, at 133 MHz, this provides 15 ns of additional address hold time at the external address latch. Running the csb at a lower frequency improves the hold time.
 - 1 During address phases on the local bus, the LALE signal negates 0.5 platform clock period prior to the address being invalidated. This halves the address

- hold time, but extends the latch enable duration. This may be necessary for very high frequency designs.”
- 10.3.1.15, 10-33 In Figure 10-19, “Clock Ratio Register (LCRR)” make bit 0 as reserved. In Table 10-22, “LCRR Field Descriptions,” remove bit 0 (PBYP) field description and modify the description for bits 0–13 to say the following: “Reserved. Note bit 0 must remain set for proper operation.”
- 10.3.1.16, 10-34 In Table 10-23, “FMR Field Descriptions,” for CWTO bit field description, change “CW0–CW3” to say “CW0, CW1, RBW and RSW.”
- 10.4, 10-40 Add the following note:
 “Different machines (FCM/GPCM/UPM) share the address, data, and control signals. While the eLBC is servicing a transaction, subsequent transactions are queued until the current transaction has completed.”
- 10.4.1.2, 10-42 Add the following two paragraphs at the end of section:
 “If the RCW is loaded by the eLBC, LALE may remain at an unknown value for up to 8 cycles after PORESET negation. Thus, LALE should be ignored for 8 CLKIN/PCI_SYNC_IN cycles after PORESET negation. In general, it is recommended that a latch be implemented for this adjustment and not a state machine triggered by LALE.
 If the RCW is not loaded by the eLBC (for example, I2C or hard-coded options are used), then LALE is at an unknown value until the PLL is locked and should be ignored until the negation of HRESET.”
- 10.4.1.6, 10-44 Add the following note to end of section describing consequences of long FCM transactions and GPCM/UPM transactions:
 When the FCM is in the middle of a long transaction (such as NAND erase, write, and so on.), another transaction on the GPCM or UPM will trigger the bus monitor to start, even though the GPCM or UPM is waiting for the FCM to finish. If the bus monitor times out, it could corrupt the current NAND flash operation as well as terminate the GPCM or UPM operation. To avoid such cases, it is recommended that the bus monitor timeout be programmed to its maximum setting of LBCR[BMT] = 0 and LBCR[BMTPS] = 0xF.
- 10.4.2, 10-46 Modify the address numbering for A in Figure 10-32, “GPCM Basic Read Timing (XACS = 0, ACS = 1x, TRLX = 0, CLKDIV = 4,8),” as follows:




Figure 10-32. GPCM Basic Read Timing (XACS = 0, ACS = 1x, TRLX = 0, CLKDIV = 4, 8)

10.4.2.1, 10-47

In Table 10-30, “GPCM Read Control Signal Timing,” modify t_{ARCS} and t_{CSRP} columns, as follows:

Table 10-30. GPCM Read Control Signal Timing

| Option Register Attributes | | | | Signal Timing (LCLK clock cycles) ¹ | | | | |
|----------------------------|------|------|-----|--|---|-----------|-----------|----------|
| TRLX | EHTR | XACS | ACS | t_{ARCS} | t_{CSRP} | t_{AOE} | t_{OEN} | t_{RC} |
| 0 | 0 | 0 | 0X | 0 | 2+SCY | 1 | 0 | 2+SCY |
| 0 | 0 | 0 | 10 | $\frac{1}{4}$ ($\frac{1}{2}$) | $1\frac{3}{4}$ +SCY (2+SCY) | 1 | 0 | 2+SCY |
| 0 | 0 | 0 | 11 | $\frac{1}{2}$ | $1\frac{1}{2}$ +SCY | 1 | 0 | 2+SCY |
| 0 | 0 | 1 | 0X | 0 | 2+SCY | 1 | 0 | 2+SCY |
| 0 | 0 | 1 | 10 | 1 | 1+SCY | 1 | 0 | 2+SCY |
| 0 | 0 | 1 | 11 | 2 | 1+SCY | 2 | 0 | 3+SCY |
| 0 | 1 | 0 | 0X | 0 | 2+SCY | 1 | 1 | 3+SCY |
| 0 | 1 | 0 | 10 | $\frac{1}{4}$ ($\frac{1}{2}$) | $1\frac{3}{4}$ +SCY ($1\frac{1}{2}$ +SCY) | 1 | 1 | 3+SCY |
| 0 | 1 | 0 | 11 | $\frac{1}{2}$ | $1\frac{1}{2}$ +SCY | 1 | 1 | 3+SCY |
| 0 | 1 | 1 | 0X | 0 | 2+SCY | 1 | 1 | 3+SCY |
| 0 | 1 | 1 | 10 | 1 | 1+SCY | 1 | 1 | 3+SCY |
| 0 | 1 | 1 | 11 | 2 | 1+SCY | 2 | 1 | 4+SCY |
| 1 | 0 | 0 | 0X | 0 | 2+2xSCY | 1 | 4 | 6+2xSCY |
| 1 | 0 | 0 | 10 | $1\frac{1}{4}$ ($1\frac{1}{2}$) | $1\frac{3}{4}$ +2xSCY ($1\frac{1}{2}$ +2xSCY) | 2 | 4 | 7+2xSCY |
| 1 | 0 | 0 | 11 | $1\frac{1}{2}$ | $1\frac{1}{2}$ +2xSCY | 2 | 4 | 7+2xSCY |
| 1 | 0 | 1 | 0X | 0 | 2+2xSCY | 1 | 4 | 6+2xSCY |

Table 10-30. GPCM Read Control Signal Timing (continued)

| Option Register Attributes | | | | Signal Timing (LCLK clock cycles) ¹ | | | | |
|----------------------------|------|------|-----|--|------------------------|------------------|------------------|-----------------|
| TRLX | EHTR | XACS | ACS | t _{ARCS} | t _{CSRP} | t _{AOE} | t _{OEN} | t _{RC} |
| 1 | 0 | 1 | 10 | 2 | 1+2xSCY | 2 | 4 | 7+2xSCY |
| 1 | 0 | 1 | 11 | 3 | 1+2xSCY | 3 | 4 | 8+2xSCY |
| 1 | 1 | 0 | 0X | 0 | 2+2xSCY | 1 | 8 | 10+2xSCY |
| 1 | 1 | 0 | 10 | 1¼ (½) | 1¾+2xSCY (1½+2xSCY) | 2 | 8 | 11+2xSCY |
| 1 | 1 | 0 | 11 | 1½ | 1½+2xSCY | 2 | 8 | 11+2xSCY |
| 1 | 1 | 1 | 0X | 0 | 2+2xSCY | 1 | 8 | 10+2xSCY |
| 1 | 1 | 1 | 10 | 2 | 1+2xSCY | 2 | 8 | 11+2xSCY |
| 1 | 1 | 1 | 11 | 3 | 1+2xSCY | 3 | 8 | 12+2xSCY |

¹ Times in parentheses are specific for the case LCRR[CLKDIV] = 2; other times apply to all CLKDIV values.

10.4.2.1, 10-47

In Table 10-31, “GPCM Write Control Signal Timing,” modify t_{AWCS} and t_{WC} columns, as follows:

Table 10-31. GPCM Write Control Signal Timing

| Option Register Attributes | | | | Signal Timing (LCLK clock cycles) ¹ | | | | |
|----------------------------|------|-----|------|--|-------------------|------------------|------------------|--------------------|
| TRLX | XACS | ACS | CSNT | t _{AWCS} | t _{CSWP} | t _{AWE} | t _{WEN} | t _{WC} |
| 0 | 0 | 00 | 0 | 0 | 2+SCY | 1 | 0 | 2+SCY |
| 0 | 0 | 10 | 0 | ¼ (½) | 1¾+SCY (2+SCY) | 1 | 0 | 2+SCY |
| 0 | 0 | 11 | 0 | ½ | 1½+SCY | 1 | 0 | 2+SCY |
| 0 | 1 | 00 | 0 | 0 | 2+SCY | 1 | 0 | 2+SCY |
| 0 | 1 | 10 | 0 | 1 | 1+SCY | 1 | 0 | 2+SCY |
| 0 | 1 | 11 | 0 | 2 | 1+SCY | 2 | 0 | 3+SCY |
| 0 | 0 | 00 | 1 | 0 | 2+SCY | 1 | ¼ (0) | 2+SCY |
| 0 | 0 | 10 | 1 | ¼ (½) | 1½+SCY | 1 | 0 | 1¾+SCY (1½+SCY) |
| 0 | 0 | 11 | 1 | ½ | 1¼+SCY (1+SCY) | 1 | 0 | 1¾+SCY (1½+SCY) |
| 0 | 1 | 00 | 1 | 0 | 2+SCY | 1 | ¼ (0) | 2+SCY |
| 0 | 1 | 10 | 1 | 1 | ¾+SCY (½+SCY) | 1 | 0 | 1¾+SCY (1½+SCY) |
| 0 | 1 | 11 | 1 | 2 | ¾+SCY (½+SCY) | 2 | 0 | 2¾+SCY (2½+SCY) |

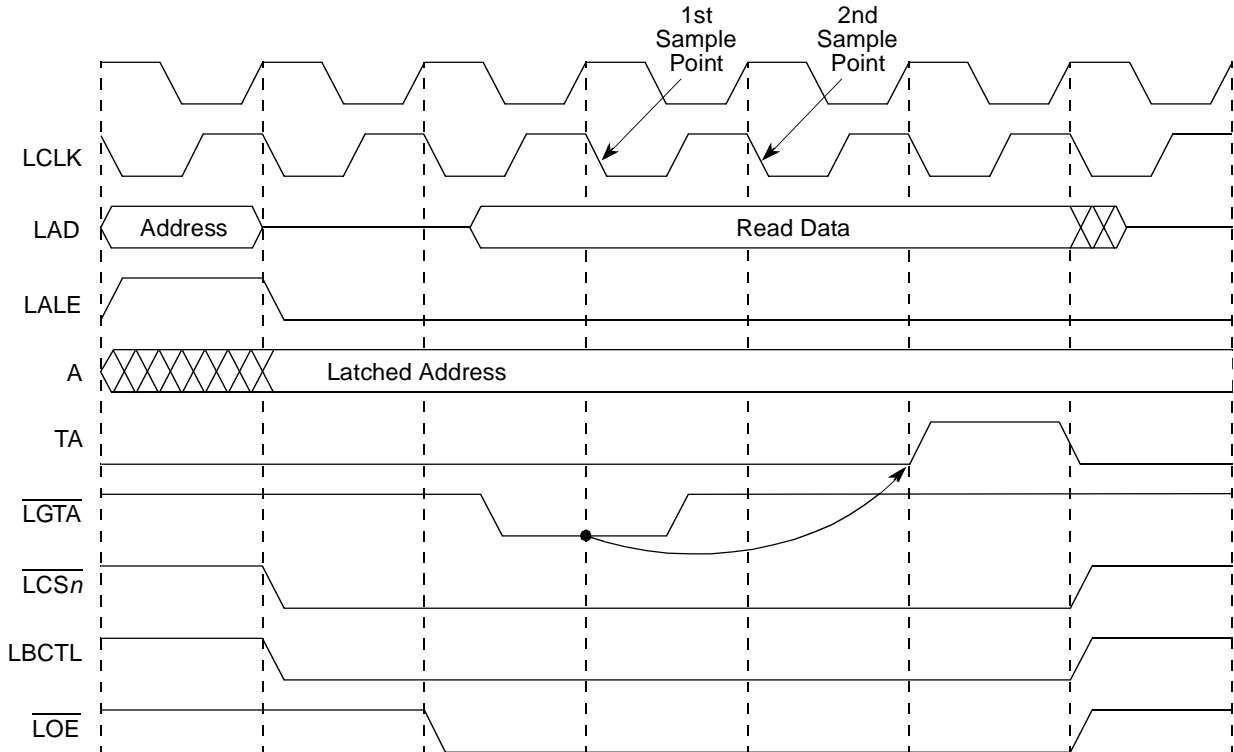
Table 10-31. GPCM Write Control Signal Timing (continued)

| Option Register Attributes | | | | Signal Timing (LCLK clock cycles) ¹ | | | | |
|----------------------------|------|-----|------|--|-----------------------|------------------|------------------|------------------------|
| TRLX | XACS | ACS | CSNT | t _{AWCS} | t _{CSWP} | t _{AWE} | t _{WEN} | t _{wc} |
| 1 | 0 | 00 | 0 | 0 | 2+2xSCY | 1 | 0 | 2+2xSCY |
| 1 | 0 | 10 | 0 | 1¼ (1½) | 1¾+2xSCY (2+2xSCY) | 2 | 0 | 3+2xSCY |
| 1 | 0 | 11 | 0 | 1½ | 1½+2xSCY | 2 | 0 | 3+2xSCY |
| 1 | 1 | 00 | 0 | 0 | 2+2xSCY | 1 | 0 | 2+2xSCY |
| 1 | 1 | 10 | 0 | 2 | 1+2xSCY | 2 | 0 | 3+2xSCY |
| 1 | 1 | 11 | 0 | 3 | 1+2xSCY | 3 | 0 | 4+2xSCY |
| 1 | 0 | 00 | 1 | 0 | 3+2xSCY | 1 | 1¼ (1) | 3+2xSCY |
| 1 | 0 | 10 | 1 | 1¼ (1½) | 1½+2xSCY | 2 | 0 | 2¾+2xSCY (2½+2xSCY) |
| 1 | 0 | 11 | 1 | 1½ | 1¼+2xSCY (1+2xSCY) | 2 | 0 | 2¾+2xSCY (2½+2xSCY) |
| 1 | 1 | 00 | 1 | 0 | 3+2xSCY | 1 | 1¼ (1) | 3+2xSCY |
| 1 | 1 | 10 | 1 | 2 | ¾+2xSCY (½+2xSCY) | 2 | 0 | 2¾+2xSCY (2½+2xSCY) |
| 1 | 1 | 11 | 1 | 3 | ¾+2xSCY (½+2xSCY) | 3 | 0 | 3¾+2xSCY (3½+2xSCY) |

¹ Times in parentheses are specific for the case LCRR[CLKDIV] = 2; other times apply to all CLKDIV values.

10.4.2.3.2, 10-51 In paragraph directly following Figure 10-35, “GPCM Basic Write Timing (XACS = 0, ACS = 00, CSNT = 1, SCY = 1, TRLX = 0, CLKDIV = 4, 8),” change “OR_n[CSNT] controls the timing for the appropriate strobe negation in write cycles” to say “OR_n[CSNT], along with OR_n[TRLX], control the timing for the appropriate strobe negation in write cycles.”

10.4.2.4, 10-56 Replace Figure 10-42, “External Termination of GPCM Access” as follows:



10.4.3.4.1, 10-70 In Table 10-36, “Boot Bank Field Values after Reset for FCM as Boot Controller,” change 010 to “From por_cfg_scy[1:3]” for SCY.

10.4.4.4.1, 10-81 In Table 10-39, “RAM Word Field Descriptions,” add the following note to fields LAST and UTA:

“In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.”

10.4.4.4.5, 10-83 Add second list item, as follows:

“Loop start word should not have an AMX change with regard to the previous word.”

10.4.4.4.7, 10-84 Modify last two sentences of first paragraph to say:

“The next address (NA) bit of the RAM word does not affect LA signals, unless AMX = 00 and chooses the column address for NA = 1.”

10.4.4.4.7, 10-84 Add the following note regarding bank addresses on multiple-bank DRAM and SDRAM devices.

Multiple-bank DRAM and SDRAM devices require that the bank address be driven during both RAS and CAS cycles. The UPM does not support a persistent bank address on both RAS and CAS cycles. Therefore, external logic must be used to supply a bank address to these devices.

| | |
|-------------------|---|
| 10.4.4.4.9, 10-86 | Add the following paragraph: “In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.” |
| 10.4.4.5, 10-87 | Remove Section 10.4.4.5, “Synchronous Sampling of LUPWAIT for Early Transfer Acknowledge.” |
| 10.5.1.3, 10-89 | Modify last sentence of second paragraph to say: “Typical values for the two propagation delays are in the order of 3–6 ns.” |
| 10.5.4.4, 10-94 | Change FMR[OP]=01 to FMR[OP]=11. |
| 10.5.4.5, 10-95 | Change FMR[OP]=01 to FMR[OP]=11. |
| 10.5.4.6, 10-95 | Change FMR[OP]=01 to FMR[OP]=11. |
| 10.5.7, 10-103 | Remove Section 10.5.7, “Interfacing to DSP Host Ports.” |
| 11.4.1, 11-3 | In Table 11-2, “POTAR _n Field Descriptions,” add the following to TA (bits 12–31) field description: “The translation address must be aligned based on the window’s size.” |
| 12.2.1, 12-2 | In Table 12-1, “DMA Interface Signals—Detailed Signal Descriptions,” change DREQ[0:3], DACK[0:3], DDONE[0:3] to DREQ[0:1], DACK[0:1], DDONE[0:1]. |
| 12.4.8.1, 12-12 | In Table 12-11, “DMAMR _n Field Descriptions,” modify bit field description for PRC, bits 11–10, as follows: “PCI read command. This field indicates the type of PCI read command to use. 00 Reserved 01 PCI read line 10 PCI read multiple 11 Reserved” |
| 12.4.8.1, 12-13 | In Table 12-11, “DMAMR _n Field Descriptions,” modify TEM (bit 3) field description, as follows: “0 The DMA will halt when a transfer error occurs. 1 The DMA will complete the transfer regardless of whether a transfer error occurs. Note: Regardless of the setting of TEM, if an error condition was detected during the DMA transfer, it will cause DMASR _n [TE] to be set.” |
| 12.4.8.2, 12-14 | In Table 12-12, “DMASR _n Field Descriptions,” modify TE (bit 7) field description, as follows: “Transfer error. Set when there is an error condition during the DMA transfer.” In addition, modify beginning of CB (bit 2) field description, as follows: “Channel busy. This bit indicates whether the channel is busy. It is cleared as a result of any of the following conditions: an error or completion of the DMA transfer.” |

- 12.5.3, 12-20 Update paragraph beginning with “Accesses to CSB memory depend...” to read as follows:
 “Accesses to CSB memory depend on the alignment of the source and destination addresses and the size of the transfer. The DMA controller transfers a full cache line whenever possible. On misaligned addresses, full cache line transfers (32 byte bursting) occur if the source and destination offsets end in the same byte offset. For example, if the destination address is 0x4000_1050 and the source address is 0x9000_2050, the transfer bursts because both end in 0x50. However, if the destination address is 0x4000_1050 and the source is 0x9000_2000, the transfer does not burst because the last byte offset is not the same. On misaligned destination addresses, subtransfers of less than a cache line occur on the initial and final beats of the transfer while full cache lines occur on intermediate beats.”
- 13.3.2.11, 13-22 Add the following text to the introductory paragraph:
 “Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an outbound window, or where an outbound translation window points back into an inbound window, are not allowed.”
- 13.3.2.11, 13-23 In Table 13-18, “PITAR_n Field Descriptions,” revise TA bit field description, as follows:
 “Translation address. Contains the starting address of the inbound translated address. TA corresponds to the 20 highest-order bits of a 32-bit local address. Note that window 0 does not support 64-bit address. Therefore, only bits 12–31 will be used to define the address bits 31–12. The specified address must be aligned to the window size, as defined by PIWAR_n[IWS].”
- 13.3.2.12, 13-23 In Table 13-19, “PIBAR_n Field Descriptions,” revise BA bit field description, as follows:
 Base address. Contains the starting address in the PCI memory space of the inbound window. This field corresponds to bits 43–12 of a 64-bit address. In PIBAR₀, the upper 12 bits are reserved because only a 32-bit address is supported. The specified address must be aligned to the window size, as defined by PIWAR_n[IWS].
- 13.4.8, 13-60 Add the following section after the Section 13.4.7, “CompactPCI Hot Swap Specification Support”:

13.4.8 Byte Ordering

Whenever data must cross a bridge between two busses, the byte ordering of data on the source and destination buses must be considered. The internal platform bus of this device is inherently big endian and the PCI bus interface is inherently little endian.

There are two methods to handle ordering of data as it crosses a bridge—address invariance and data invariance. Address invariance preserves the addressing of bytes within a scalar data element, but not the

relative significance of the bytes within that scalar. Conversely, data invariance preserves the relative significance of bytes within a scalar, but not the addressing of the individual bytes that make up a scalar.

This device uses address invariance as its byte ordering policy

As stated above, address invariance preserves the byte address of each byte on an I/O interface as it is placed in memory or moved into a register. This policy can have the effect of reversing the significance order of bytes (most significant to least significant and vice versa), but it has the benefit of preserving the format of general data structures. Provided that software is aware of the endianness and format of the data structure, it can correctly interpret the data on either side of the bridge.

Figure 13-55 shows the transfer of a 4-byte scalar, 0x4142_4344, from a big endian source across an address invariant bridge to a little endian destination.




Figure 13-55. Address Invariant Byte Ordering—4 bytes Outbound

Note that although the significance of the bytes within the scalar have changed, the address of the individual bytes that make up the scalar have not changed. As long as software is aware that the source of the data used a big endian format, the data can be interpreted correctly.

Figure 13-56 shows data flowing the other way, from a little endian source to a big endian destination.




Figure 13-56. Address Invariant Byte Ordering—4 bytes Inbound

Figure 13-57 shows an outbound transfer of an 8-byte scalar, 0x5455_1617_CDCE_2728, using address invariance.




Figure 13-57. Address Invariant Byte Ordering—8 bytes Outbound

Figure 13-58 shows an inbound transfer of a 2-byte scalar, 0x5837, using address invariance.




Figure 13-58. Address Invariant Byte Ordering—2 bytes Inbound

Note that in all of these examples, the original addresses of the individual bytes within the scalars (as created by the source) have been preserved.

13.4.1.21 Byte Order for Configuration Transactions

All internal memory-mapped registers in the CCSR space use big endian byte ordering. However, the PCI specification defines PCI configuration registers as little endian. All accesses to the PCI configuration port, CFG_DATA, including the those targeting the internal PCI configuration registers, use the address invariance policy as shown in Figure 13-59. Therefore, software must access CFG_DATA with little-endian formatted data—either using the `lwbrx/stwbrx` instructions or by manipulating the data before writing to and after reading from CFG_DATA.




Figure 13-59. CFG_DATA Byte Ordering

Chapter 14, 14-1

Add the following note:

“The PCI Express engine does not support mis-aligned byte transfers. It must be DWORD aligned to the CSB bus.”

14.1, 14-3

Update last paragraph to say the following:

“As an initiator, the device supports memory read and write operations with a maximum payload of 128 bytes and I/O transactions. In addition, outbound configuration are supported if the device is in RC mode. As a target interface, the device accepts read and write operations to local memory space. Furthermore, as an EP device, the device accepts configuration transactions to the internal PCI Express configuration registers. Inbound I/O transactions are not supported.”

14.3.1, 14-6

Update Table 14-3, “PCI Express Memory Map,” as follows. Note that only the affected rows are shown:

Table 14-3. PCI Express Memory Map

| Offset | Register | Access | Reset | Section/Page |
|--|--|---------|-------------|----------------------------------|
| PCI Express 1—Block Base Address 0x0_9000 PCI Express 2—Block Base Address 0x0_A000 | | | | |
| PCI Express1 Core Configuration Header Registers | | | | |
| 0x00F | PCI Express BIST Register | R | All zeros | 14.4.1.8/14-20 |
| 0x010– 0x014 | Base Address Registers 0 and 1 (BAR0/BAR1) (EP mode only) | Mixed | 0x0008 | 14.4.2.1.1/14-22 |
| 0x018– 0x020 | Base Address Registers 2 and 4 (BAR2/BAR4) (EP mode only) | Mixed | 0x0000_000C | 14.4.2.1.2/14-23 |
| 0x01C– 0x024 | Base Address Registers 3 and 5 (BAR3/BAR5) (EP mode only) | R/W | All zeros | 14.4.2.1.3/14-23 |
| 0x02C | PCI Express Subsystem Vendor ID Register (EP mode only) | Special | All zeros | 14.4.2.2/14-24 |
| 0x02E | PCI Express Subsystem ID Register (EP mode only) | Special | All zeros | 14.4.2.3/14-24 |
| 0x134 | PCI Express Error Source Identification Register | R | All zeros | 14.4.5.11/14-60 |
| PCI Express Core Control and Status Registers (CSRs) | | | | |
| 0x41C | PCI Express N_FTS Control Register (PEX_NFTS_CTRL) | R/W | 0x0000_4040 | 14.4.6.2/14-62 |
| 0x438 | PCI Express ACK Replay Timeout Register (PEX_ACKRPLY_TO) | R/W | 0x005B_2090 | 14.4.6.3/14-63 |
| 0x45C | PCI Express ASPM Request Timer Register (PEX_ASPM_REQTMR) (RC mode only) | R/W | 0x0000_0629 | 14.4.6.7/14-64 |
| 0x478 | PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE) | R/W | All zeros | 14.4.6.8/14-65 |
| 0x47C | PCI Express Device Capabilities Update Register (PEX_DEVCAP_UPDATE) | R/W | All zeros | 14.4.6.9/14-66 |
| 0x480 | PCI Express Link Capabilities Update Register (PEX_LINKCAP_UPDATE) | R/W | 0x0000_3D41 | 14.4.6.10/14-67 |
| 0x490 | PCI Express Slot Capabilities Update Register (PEX_SLCAP_UPDATE) | R/W | 0x0000_07C0 | 14.4.6.11/14-68 |
| PCI Express BAR Configuration Registers (EP Mode) | | | | |
| 0x4D4 | PCI Express BAR Enable Register (PEX_BAR_ENABLE) | R/W | 0x0000_000F | 14.4.7.1/14-71 |

Table 14-3. PCI Express Memory Map

| | | | | |
|---|--|-----|-------------|--------------------------------|
| 0x4E0 | PCI Express BAR Select Configuration Register (PEX_BAR_SEL) | R/W | 0x0000_0000 | 14.4.7.3/14-72 |
| 0x504 | PCI Express BAR Prefetch Configuration Register (PEX_BAR_PF) | R/W | 0x0000_0000 | 14.4.7.4/14-73 |
| PCI Express Extended Status and Control Register | | | | |
| 0x590 | PCI Express PME_To_Ack Timeout Register (PEX_PME_TO_ACK_TOR) | R/W | 0x00FD_4BC0 | 14.4.8.1/14-74 |
| 0x594 | PME_To_Ack Status Register (PEX_PME_TO_ACK_SR) | w1c | All zeros | 14.4.8.2/14-75 |
| 0x5A0 | PCI Express PCI Interrupt Mask Register (PEX_SS_INTR_MASK) | R/W | 0x0000_003F | 14.4.8.3/14-76 |
| Power Management Registers | | | | |
| 0xC80 | PCI Express PM Control Register (PEX_PM_CTRL) | R/W | All zeros | 14.5.9.1/14-93 |
| 0xC88 | PCI Express slot control misc register | R/W | All zeros | 14.5.9.2/14-99 |

14.4.1, 14-15 Add the following note:

NOTE

The registers described in this section use little-endian byte ordering. Software running on the local processor in big-endian mode must byte-swap the data. No byte swapping occurs when the registers are accessed from the PCI Express bus.

14.4.1.3, 14-17 In Table 14-6, “PCI Express Command Register Field Descriptions,” change the Bus master (bit 2) field description to say the following:
 “Enables/disables this PCI Express device to behave as a PCI Express bus master.
 0 Disables the ability to generate PCI Express accesses.
 1 Enables this PCI Express controller to behave as a bus master.”

14.4.1.3, 14-17 In Table 14-6, “PCI Express Command Register Field Descriptions,” change the Memory space (bit 1) field description to say the following:
 “Controls whether this PCI Express device (as a target) responds to memory accesses.
 0 Device does not respond to PCI Express memory space accesses.
 1 Device responds to PCI Express memory space accesses”

14.4.2.2, 14-24 In Figure 14-16, “PCI Express Subsystem Vendor ID Register,” change the access from “Read-only” to “Special,” and in Table 14-16, “PCI Express Subsystem Vendor ID Register Field Descriptions,” modify the field description to say the following:
 “Subsystem Vendor ID. The value of this register can be set by programming the SSVID field of the PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE). See [Section 14.4.6.8, “PCI Express Subsystem Vendor ID Update Register \(PEX_SSVID_UPDATE\).”](#)”

This value has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration.”

14.4.2.3, 14-24

In Figure 14-17, “PCI Express Subsystem ID Register,” change the access from “Read-only” to “Special,” and in Table 14-17, “PCI Express Subsystem ID Register Field Descriptions,” modify the field description as follows:

“Subsystem ID. The value of this register can be set by programming the SSID field of the PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE). See [Section 14.4.6.8, “PCI Express Subsystem Vendor ID Update Register \(PEX_SSVID_UPDATE\).”](#)”

This value has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration.”

14.4.2.6, 14-26

In Table 14-20, “PCI Express Interrupt Pin Register Field Descriptions” change the field description to say the following:

| Bits | Name | Description |
|------|---------------|---|
| 7–0 | Interrupt pin | Legacy INTx message used by this device. 0x01 INTA only is supported by this device. |

14.4.3.18, 14-35

In Table 14-39, “PCI Express Interrupt Pin Register Field Descriptions” modify the description of bits 7–0 as follows:

| Bits | Name | Description |
|------|---------------|---|
| 7–0 | Interrupt Pin | Legacy INTx message used by this device. 0x01 INTA only is supported by this device. |

14.4.4.3, 14-39

Modify Table 14-43, “PCI Express Power Management Capabilities Register Field Descriptions,” as follows:

Table 14-43. PCI Express Power Management Capabilities Register Fields Description

| Bits | Name | Description |
|-------|-------------|--|
| 15–11 | PME Support | For a device, this 5-bit field indicates the power states in which the device may generate a PME. PME can be issued from D0, D1, D2 and D3hot. |
| 10 | D2 | D2 power state is supported. |
| 9 | D1 | D1 power state is supported. |
| 8–6 | AUX Curr | AUX Current. Vaux and D3cold is not supported by this device. |
| 5 | DSI | A Device Specific Initialization is not required. |
| 4 | — | Reserved |
| 3 | PME CLK | Does not apply to PCI Express |
| 2–0 | Version | 0x02 indicates compatibility to the PCI EXPRESS BASE SPECIFICATION, REV. 1.0a |

| Bits | Name | Description |
|-------|-------------|---|
| 14–12 | L0s_EX_LAT | L0s exit latency. 0b011 indicates 256 ns to less than 512 ns |
| 11–10 | ASPM | Active state power management (ASPM) Support, L0s Entry Supported |
| 9–4 | MAX_LINK_W | Maximum link width 0b000001 x1 |
| 3–0 | MAX_LINK_SP | Maximum link speed, 0b0001 indicates 2.5 Gb/s |

14.4.4.15, 14-45 Update Section 14.4.4.15, “PCI Express Slot Capabilities Register,” as follows:

14.4.4.15 PCI Express Slot Capabilities Register

The PCI Express slot capabilities register is shown in [Figure 14-21](#). For End Point applications implementing a slot, the content of this register can be modified by using the PCI Express Slot Capabilities Update Register as described in [Section 14.4.6.11, “PCI Express Slot Capabilities Update Register \(PEX_SLCAP_UPDATE\).”](#)




Figure 14-60. PCI Express Slot Capabilities Register

Table 14-49. PCI Express Slot Capabilities Register Field Descriptions

| Bits | Name | Description |
|-------|----------------------|--|
| 31–19 | Physical Slot Number | This field indicates the physical slot number attached to this Port. |
| 18–17 | — | Reserved |
| 16–15 | SPLS | Slot power limit scale. |
| 14–17 | SPLV | Slot power limit value |
| 6 | HPD | Hot plug capable |
| 5 | HPS | Hot plug surprise |
| 4 | PIP | Power indicator present |
| 3 | AIP | Attention indicator present |
| 2 | MRLSP | MRL sensor present |

14.4.6, 14-61 Modify entire section and its subsections, as follows:

14.4.6 PCI Express Controller Internal Control and Status Registers (CSRs)

14.4.6.1 PCI Express LTSSM State Status Register (PEX_LTSSM_STAT)

PEX_LTSSM_STAT, shown in [Figure 14-80](#), provides details on link training status. This register is useful for debugging link training failures.



Figure 14-80. PCI Express LTSSM State Status Register (PEX_LTSSM_STAT)

The fields of the PEX_LTSSM_STAT are described in [Table 14-7](#).

Table 14-7. PEX_LTSSM_STAT Field Descriptions

| Bits | Name | Description |
|------|-------------|---|
| 31–7 | — | Reserved |
| 6–0 | Status code | Status code. See Table 14-78 for encodings. |

[Table 14-78](#) provides the encodings for the status code field of the PEX_LTSSM_STAT register.

Table 14-78. PEX_LTSSM_STAT Status Codes

| Status Code (Hex) | LTSSM State Description | Status Code (Hex) | LTSSM State Description |
|-------------------|-------------------------------------|-------------------|-------------------------|
| 00 | Detect quiet | 27 | TX L0s FTS; RX L0s FTS |
| 01 | Detect active (0) | 28 | L0 to L1 (0) |
| 02 | Detect active (1) | 29 | L0 to L1 (1) |
| 03 | Detect active (2) | 2A | L1 entry |
| 04 | Polling active (0) | 2B | L1 idle (0) |
| 05 | Polling active (1) | 2C | L1 idle (1) |
| 06 | Polling config (0) | 2D | L0 to L2 (0) |
| 07 | Polling config (1) | 2E | L0 to L2 (1) |
| 08 | Polling compliance | 2F | L2 entry |
| 09 | Configuration link width start (0) | 30 | L2 idle (0) |
| 0A | Configuration link width start (1) | 31 | L2 idle (1) |
| 0B | Configuration link width accept (0) | 32 | Recovery lock (0) |
| 0C | Configuration link width accept (1) | 33 | Recovery lock (1) |
| 0D | Configuration lane number wait (0) | 34 | Recovery lock (2) |
| 0E | Configuration lane number wait (1) | 35 | Recovery cfg (0) |

Table 14-78. PEX_LTSSM_STAT Status Codes (continued)

| Status Code (Hex) | LTSSM State Description | Status Code (Hex) | LTSSM State Description |
|-------------------|--|-------------------|-------------------------------------|
| 0F | Configuration lane number wait (2) | 36 | Recovery cfg (1) |
| 10 | Configuration lane number wait (3) | 37 | Recovery idle (0) |
| 11 | Configuration lane number accept | 38 | Recovery idle (1) |
| 12 | Configuration complete (0) | 39 | Recovery to configuration |
| 13 | Configuration complete (1) | 3A | Recovery cfg to configuration |
| 14 | Configuration idle (0) | 3F | L0 no training |
| 15 | Configuration idle (1) | 7F | Detect quiet EI |
| 16 | L0 | 49 | Configuration link width start—RC |
| 17 | TX L0; RX L0s entry | 4A | Configuration link width accept—RC |
| 18 | TX L0; RX L0s idle | 4B | Configuration lane number wait—RC |
| 19 | TX L0; RX L0s fast training sequence (FTS) | 4C | Configuration lane number accept—RC |
| 1A | TX L0s entry (0); RX L0 | 60 | Loopback slave active (0) |
| 1B | TX L0s entry (0); RX L0s idle | 61 | Loopback slave active (1) |
| 1C | TX L0s entry (0); RX L0s FTS | 62 | Loopback slave exit |
| 1D | TX L0s entry (1); RX L0 | 68 | Hot reset (0) |
| 1E | TX L0s entry (1); RX L0s idle | 69 | Hot reset (1) |
| 1F | TX L0s entry (1); RX L0s FTS | 6A | Hot reset (0)—RC |
| 20 | TX L0s idle; RX L0 | 6B | Hot reset (1)—RC |
| 21 | TX L0s idle; RX L0s entry | 75 | Disabled (0) |
| 22 | TX L0s idle; RX L0s idle | 71 | Disabled (1) |
| 23 | TX L0s idle; RX L0s FTS | 72 | Disabled (2) |
| 24 | TX L0s FTS; RX L0 | 73 | Disabled (3) |
| 25 | TX L0s FTS; RX L0s entry | 74 | Disabled (4) |
| 26 | TX L0s FTS; RX L0s idle | 78 | L0 to L1/L2—RC |

14.4.6.2 PCI Express N_FTS Control Register (PEX_NFTS_CTRL)

The PCI Express N_FTS Control Register, shown in [Figure 14-81](#), is used to set the N_FTS value that is advertised by the PCI Express controller during link training. It is preferable to set it before the PCI Express core internal reset is negated. If this value is changed after the link is up, the new value will take effect during the next link training. The N_FTS value is should be set according to the L0s exit latency of the Rx link of PHY. At a given time, either N_FTS or N_FTS_COM value is used based on the setting of common clock configuration bit in the configuration space.

received for some more time due to the link and processing latency involved before downstream device processes the NAK msg and stops the request. Meanwhile, these request DLLPs should not be considered as a new L1 entry request and responded with Ack DLLP. For more details refer to the PCI Express Base Specification Rev 1.0a errata, C7. ASPM & PCI-PM L1.

The fields of the PCI Express ASPM Request Timer Register are described in [Table 14-84](#).




Figure 14-86. PCI Express ASPM Request Timer Register

Table 14-84. PCI Express ASPM Request Timer Register Field Descriptions

| Bits | Name | Description |
|-------|-----------|--|
| 31–13 | — | Reserved |
| 12–0 | ASPML1TMR | ASPM L1 request timer value. This is the time-out interval after sending NAK message, before a new ASPM L1 request from a downstream device is treated as a new ASPM L1 entry request. For example: If the upstream device rejects an ASPM L1 entry request from a downstream device with ASPM NAK message, the next ASPM L1 entry request from downstream device will be entertained only after this timeout interval or only after the Rx link of the downstream port enters L0s state. This value is specified in terms of system clock cycles (CSB clock). The default value is equivalent to 9.5 μs in a 166 MHz system clock. This value can be calculated as [Time in microseconds * SYSTEM_CLK in MHz]. For example, 9.5[μs]*166[MHz] = 1577 (0x629). This register is used only in RC mode. |

14.4.6.8 PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE)

The PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE) shown in [Figure 14-87](#) is used to configure Subsystem Vendor ID and Subsystem ID fields of configuration header (offset 0x2C) for End Point devices. This register has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration.




Figure 14-87. PCI Express Subsystem Vendor ID Update Register (PEX_SSVID_UPDATE)

Table B-85. PEX_SSVID_UPDATE Field Descriptions

| Bits | Name | Description |
|-------|-------|---------------------|
| 31–16 | SSID | Subsystem ID |
| 15–0 | SSVID | Subsystem Vendor ID |

Table 14-87. PCI Express Link Capabilities Update Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 7–6 | ASPM | ASPM Support. Indicates the level of ASPM supported on the given PCI Express Link. Defined encodings are: 00b Reserved 01b L0s Entry Supported 10b Reserved 11b Reserved (L0s and L1 not supported by this device) |
| 5–0 | MLW | Maximum Link Width of the given PCI Express Link. Defined encodings are: 000001b x1 Other: Reserved |

14.4.6.11 PCI Express Slot Capabilities Update Register (PEX_SLCAP_UPDATE)

The PCI Express Slot Capabilities Update Register shown in [Figure 14-90](#) is used to set the values to the PCI Express Slot Capabilities Register in the PCI Express configuration header (offset 0x60). It can be used when the device is configured as an End Point to make the correct slot information available to the upstream device. This register has to be programmed before setting the config-ready bit in the PCI Express Configuration Ready Register so that the host reads the correct information during enumeration.




Figure 14-90. PCI Express Slot Capabilities Update Register

Table 14-88. PCI Express Slot Capabilities Update Register Field Descriptions


| Bits | Name | Description |
|-------|------|---|
| 31–30 | — | Reserved |
| 29–17 | PSN | Physical Slot Number. Physical Slot number attached to the port. |
| 16–15 | SPLS | Slot Power Limit Scale. Specifies the scale used for the Slot Power Limit Value. Range of Values: 00 1.0x 01 0.1x 10 0.01x 11 0.001x |
| 14–7 | SPLV | Slot power limit value. In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot. Power limit (in Watts) calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. |
| 6 | HPC | Hot plug capable |
| 5 | HPS | Hot plug surprise |
| 4 | PIP | Power indicator present |
| 3 | AIP | Attention indicator present |
| 2 | MRL | MRL sensor present |

Table 14-88. PCI Express Slot Capabilities Update Register Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--------------------------|
| 1 | PCP | Power controller present |
| 0 | ABP | Attention button present |

14.4.6.12 PCI Express Configuration Ready Register

The PCI Express configuration ready register, shown in [Figure 14-91](#), indicates configuration complete status to the transaction layer. The transaction layer handles configuration requests from external hosts only after the CFG_READY bit is set. All the configuration requests received from external hosts before the CFG_READY bit is set are completed with configuration request retry status (CRS). To ensure that the external host reads the correct capability advertisements during enumeration, the CFG_READY bit in this register should be set only after all relevant configuration registers are programmed.


Figure 14-91. PCI Express Configuration Ready Register (PEX_CFG_READY)

The fields of the PCI Express configuration ready register are described in [Table 14-89](#).

Table 14-89. PEX_CFG_READY Field Descriptions

| Bits | Name | Description |
|------|-----------|---|
| 31-1 | — | Reserved |
| 0 | CFG_READY | Configuration ready 1 The transaction layer accepts inbound configuration requests. 0 The transaction layer responds to all inbound configuration requests with retry (CRS) |

14.4.7, 14-65 Add the section for PCI Express BAR Enable Register (PEX_BAR_ENABLE) prior to PEX_BAR_SIZEL register as follows. Note that the subsequent sections are re-numbered.

14.4.7.1 PCI Express BAR Enable Register (PEX_BAR_ENABLE)

PEX_BAR_ENABLE, shown in [Figure 14-85](#), is used to enable BARs. It supports a maximum of four BARs for the controller. To enable a given BAR, the bit corresponding to the BAR has to be set.

Revision History




Figure 14-85. PCI Express BAR Enable Register (PEX_BAR_ENABLE)

The fields of the PEX_BAR_ENABLE register are described in [Table 14-83](#).

Table 14-83. PEX_BAR_ENABLE Field Descriptions

| Bits | Name | Description |
|------|------------|---|
| 31-4 | — | Reserved. Must be zeros |
| 3-0 | BAR_ENABLE | BAR enable signal for all the BARs in PCI Express. A given BAR will be enabled only if the corresponding bit in this register is enabled. Bit 0 corresponds to BAR0, bit1 to BAR1 and so on. For example, if this field is 0000001111, all four BARs are enabled. |

- 14.4.7.1, 14-66 In Table 14-83, “PEX_BAR_SIZEL Field Descriptions,” add the note “for 64 bit BAR only” against 4 Gigabytes window to read as follows:
“0000...0000 - 4 Gigabytes window (for 64 bit BAR only).”
- 14.4.7.2, 14-66 In Figure 14-86, “PCI Express BAR Select Configuration Register (PEX_BAR_SEL),” update reset value of PEX_BAR_SEL to all zeros.
- 14.4.7.3, 14-67 In Figure 14-87, “PCI Express BAR Prefetch Configuration Register (PEX_BAR_PF)” update the reset value of PEX_BAR_PF to all zeros
- 14.4.8.2, 14-68 Add PME_To_Ack Status Register (PEX_PME_TO_ACK_SR) at offset address 0x594 as follows:

14.4.8.2 PME_To_Ack Status Register (RC Mode Only)

The PME_To_Ack Status Register (PEX_PME_TO_ACK_SR) shown in [Figure 14-89](#), can be used by the Power Manager software to decide whether it is safe to switch off power to downstream devices, after PME_Turn_Off message has been broadcast by the Root Port.




Figure 14-89. PME_To_Ack Status Register (PEX_PME_TO_ACK_SR)

The fields of the PEX_PME_TO_ACK_SR are described in [Table 14-87](#).

Offset 0xBFC

Access: w1c




Figure 14-121. CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR)

Table 14-119. PEX_CSMISR Register Field Descriptions

| Bits | Name | Description |
|-------|---------|---|
| 31–30 | — | Reserved |
| 29 | ERRD | Error detected. Indicates that a PCI Express event occurred. The PCI Express event is reported by the secondary status register (PCI Express secondary status register) at address 0x901E. Note that there is a secondary mask, the PCI Express interrupt mask register (PEX_SS_INTR_MASK), at address 0x95A0. Valid only for RC. This bit must be cleared after the interrupt is serviced and after the associated status registers in the PCI Express controller causing the interrupt are cleared. |
| 28–21 | — | Reserved |
| 20 | PSCH | Power state change. Indicates that a device power state change occurred. Change in Device power state (D state) for function-0. D-state can transition between the supported values of D0, D1, D2 and D3-hot. PM software changes D-state with a configuration write to PMCSR register in PM capability of the PCI Express. The new D-state is available in the corresponding field of the PMCSR register. This bit must be cleared after the interrupt is serviced. |
| 19 | EXL2L3 | Exited L2/L3 state. Indication that the L2/L3 ready state has been exited and the current link state is L0. Traffic can be re-started on the link. This bit is set when the link switches from the L2/L3 ready to L0 state in response to an Exit L2 command. After issuing this command, the user must wait until the exited EXL2L3 bit is set before initiating traffic. This bit is valid only in RC mode. Setting this bit causes an interrupt to be sent to CSB side. The bit must be cleared after the interrupt is serviced. |
| 18 | ENL2L3 | Entered L2/L3 state. Indication to the Power manager that it is safe to switch off power to the downstream device 100nsec after this bit is set. It is set when the PCI Express controller enters L2/L3 ready state. This bit is valid only in RC mode. Setting this bit causes an interrupt to be sent to the CSB side. The bit must be cleared after the interrupt is serviced. |
| 17 | PMETMO | PME Turn Off Ack timeout event. Indication to power manager software that it is safe to switch off power to the downstream device. It is set when the PCI Express controller detects that the timeout interval for receiving a PME_To_Ack message from the downstream device has expired. This bit is valid only in RC mode. This bit must be cleared after the interrupt is serviced. |
| 16 | RPMETMO | Received PME Turn off message. Notifies that main power to the device is to be removed. After this notification is received, Uplink must not try to transmit TLPs or initiate PME requests because the power may be switched off. After this message is received, the user should indicate the readiness to lose power by setting the Initiate L2/L3 entry bit. This bit is valid only in EP mode. Setting this bit causes an interrupt to be sent to the CSB side. The bit must be cleared after the interrupt is serviced. |
| 15 | ROF | Receive overflow error. Indicates that the PCI Express reported a receive overflow error. |
| 14 | ECRC | ECRC error. Indicates that a TLP received by the PCI Express failed the ECRC check. |

Table 14-119. PEX_CSMISR Register Field Descriptions (continued)

| Bits | Name | Description |
|------|-------|---|
| 13 | PTLP | Poisoned TLP. Indicates that a poisoned TLP was received by the PCI Express. |
| 12 | MFTLP | Malformed TLP. Indicates that a malformed TLP was received by the PCI Express |
| 11 | CTO | Completion timeout. Indicates that a PCI Express completion timeout occurred. |
| 10 | CA | Completer abort. Indicates that a PCI Express completion Abort was received. |
| 9 | UR | Unsupported request. Indicates that an unsupported request was received. |
| 8 | INTD | PCI Express INTD. Indicates that an INTD interrupt was received on the PCI Express link. Valid for RC applications only. |
| 7 | INTC | PCI Express INTC. Indicates that an INTC interrupt was received on the PCI Express link. Valid for RC applications only. |
| 6 | INTB | PCI Express INTB. Indicates that an INTB interrupt was received on the PCI Express link. Valid for RC applications only. |
| 5 | INTA | PCI Express INTA. Indicates that an INTA interrupt was received on the PCI Express link. Valid for RC applications only. |
| 4–2 | — | Reserved |
| 1 | RST | PCI Express reset. Indicates that a PCI Express reset or link down occurred. |
| 0 | IMB | Inbound mailbox ready. Indicates that the inbound mailbox control register ready bit (PEX_IMBCR[READY]) was set and the CSB host can read the mailbox data. |

14.5.9, 14-93 Add Section 14.5.9, “PCI Express Power Management Registers,” as follows. Note that the subsequent sections, figures, and tables are re-numbered

14.5.9 PCI Express Power Management Registers

14.5.9.1 PCI Express Power Management Control Register (PEX_PM_CTRL)

This PCI Express PM Control Register shown in [Figure 14-122](#), is used to control the link power management by the PCI Express controller.




Figure 14-122. PCI Express PM Control Register (PEX_PM_CTRL)

Table 14-121 describes the PCI Express slot control misc register fields.

Table 14-121. PCI Express Slot Control Misc Register Field Descriptions

| Bits | Name | Description |
|------|-------------------|---|
| 0 | Slot implemented | When set, this bit indicates that the PCI Express link associated with this port is connected to a slot. |
| 1 | Slot clock config | This bit indicates that the component uses the same physical reference clock that the platform provides on the connector. If the device uses an independent clock regardless of the presence of a reference on the connector, this bit must be cleared. |
| 2–31 | — | Reserved |

14.5.9.1, 14-94 In Table 14-120, “PEX_OWAR0–PEX_OWAR3 Register Field Descriptions,” modify NSNP (bit 4) field description, as follows:

| | | |
|---|------|---|
| 4 | NSNP | No snoop enable. When this bit and the PCI Express device control register [Enable No Snoop] bit are set, the No Snoop bit for the packet is enabled. This attribute is not applicable and must be cleared for configuration requests, I/O requests, and memory requests that are Message Signaled Interrupts. 0 PCI Express TLP snoop enabled 1 PCI Express TLP snoop disabled |
|---|------|---|

14.5.9.2, 14-95 Add the following to the register description:
“Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an outbound window, or where an outbound translation window points back into an inbound window, are not allowed.”

14.5.10.1, 14-97 Add the following to the register description:
“Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an outbound window, or where an outbound translation window points back into an inbound window, are not allowed.”

14.5.11.2, 14-99 Add the following to the register description:
“Inbound and outbound windows for the same bus should not overlap. Therefore, situations where an inbound window translation points back into an outbound window, or where an outbound translation window points back into an inbound window, are not allowed.”

14.5.11.2, 14-99 Modify Figure 14-128, “PCI Express RC Inbound Window Translation Address Register *n* (PEX_RCIWTAR0–PEX_RCIWTAR3),” as follows:



Figure 14-128. PCI Express RC Inbound Window Translation Address Register *n* (PEX_RCIWTAR0–PEX_RCIWTAR3)

14.5.11.3, 14-99 Modify Figure 14-129, “PCI Express RC Inbound Window Base Address Register Low n (PEX_RCIWBARL0–PEX_RCIWBARL3),” as follows:




Figure 14-129. PCI Express RC Inbound Window Base Address Register Low n (PEX_RCIWBARL0–PEX_RCIWBARL3)

14.6.1, 14-103 Replace Section 14.6.1.2, “Byte Swapping,” Section 14.6.1.3 Outbound Byte Swapping, and Section 14.6.1.4 “Inbound Byte Swapping” with the following:

14.6.1.2 Byte Ordering

Whenever data must cross a bridge between two busses, the byte ordering of data on the source and destination busses must be considered. The internal platform bus of this device is inherently big endian and the PCI Express bus interface is inherently little endian.

There are two methods to handle ordering of data as it crosses a bridge—address invariance and data invariance. Address invariance preserves the addressing of bytes within a scalar data element, but not the relative significance of the bytes within that scalar. Conversely, data invariance preserves the relative significance of bytes within a scalar, but not the addressing of the individual bytes that make up a scalar.

This device uses address invariance as its byte ordering policy.

As stated above, address invariance preserves the byte address of each byte on an I/O interface as it is placed in memory or moved into a register. This policy can have the effect of reversing the significance order of bytes (most significant to least significant and vice versa), but it has the benefit of preserving the format of general data structures. Provided that software is aware of the endianness and format of the data structure, it can correctly interpret the data on either side of the bridge.

Figure 14-134 shows the transfer of a 4-byte scalar, 0x4142_4344, from a big endian source across an address invariant bridge to a little endian destination.




Figure 14-134. Address Invariant Byte Ordering—4 bytes Outbound

Note that although the significance of the bytes within the scalar have changed, the address of the individual bytes that make up the scalar have not changed. As long as software is aware that the source of the data used a big endian format, the data can be interpreted correctly.

Figure 14-135 shows data flowing the other way, from a little endian source to a big endian destination.




Figure 14-135. Address Invariant Byte Ordering—4 bytes Inbound

Figure 14-136 shows an outbound transfer of an 8-byte scalar, 0x5455_1617_CDCE_2728, using address invariance.




Figure 14-136. Address Invariant Byte Ordering—8 bytes Outbound

Figure 14-137 shows an inbound transfer of a 2-byte scalar, 0x5837, using address invariance.




Figure 14-137. Address Invariant Byte Ordering—2 bytes Inbound

Note that in all of these examples, the original addresses of the individual bytes within the scalars (as created by the source) have been preserved.

14.6.1.2.1 Byte Order for Configuration Transactions

All internal memory-mapped registers in the CCSR space use big endian byte ordering. However, the PCI Express specification defines PCI Express configuration registers as little endian. All accesses to the PCI Express configuration port, PEX_CONFIG_DATA, including the those targeting the internal PCI Express configuration registers, use the address invariance policy as shown in Figure 14-138. Therefore, software must access PEX_CONFIG_DATA with little-endian formatted data—either using the **lwbrx/stwbrx** instructions or by manipulating the data before writing to and after reading from PEX_CONFIG_DATA.

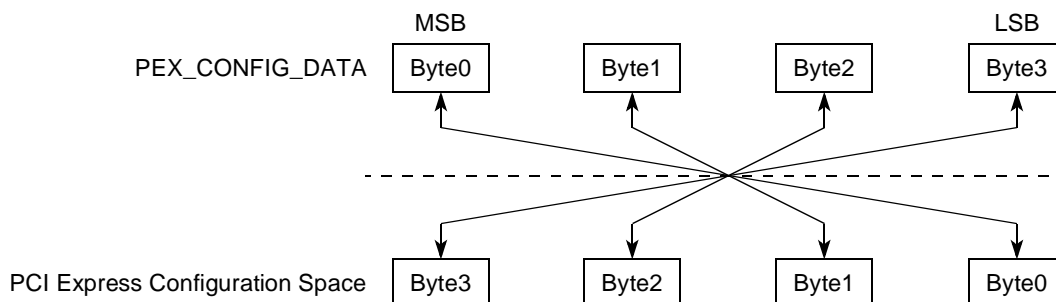


Figure 14-138. PEX_CONFIG_DATA Byte Ordering

14.6.1.8, 14-104 Revise entire section and its subsections, as follows:

14.6.1.8 Configuration Space Access

To access the PCI Express controller's internal configuration header by MPC8315E itself, the only mechanism supported is the direct access via the CSB, since the whole internal configuration space is memory-mapped. This is true regardless the PCI Express controller is configured as RC or EP.

If the PCI Express controller is configured as RC,

- Inbound configuration transaction is not supported.
- Outbound configuration transaction to access downstream PCI Express devices is supported. The only mechanism can be used to initiate either Type 0 or Type 1 configuration cycle is via outbound ATMU windows.

If the PCI Express controller is configured as EP,

- Outbound configuration transaction is not supported. In other words, the PCI Express EP controller does not generate configuration transactions in EP mode.
- Inbound configuration transaction to access the PCI Express EP controller's configuration space is supported.

14.6.1.8.1 Outbound ATMU Configuration Transaction Generation (RC)

In RC mode, the PCI Express controller can generate both Type 0 and Type 1 configuration cycles to access downstream PCI Express devices via the outbound ATMU windows mechanism.

As RC the PCI Express controller configuration access mechanism utilizes a flat memory-mapped address space to access device configuration registers. To achieve this, software can program the TYPE field of the PEX_OWAR_n register to 0x0 in one of the outbound ATMU windows to perform a configuration access. The other bit fields of the PEX_OWAR_n register should be programmed as below:

- TC = 0x0;
- NSNP = 0;
- RLXO = 0;
- EN = 1;

The SIZE field of the PEX_OWARN register should be set to a value to correspond with the BA (base address) field of the base address register (PEX_OWBARn), normally based on the Bus Number(s) of the downstream PCI Express device(s) to be accessed. The base address registers, PEX_OWBARn, set the CSB address window for the configuration transactions. The translation address registers, PEX_OWTLn, can be used to define the translated PCI Express address of the CSB-based configuration transaction.

Once the PCI Express outbound window attributes register (PEX_OWARN), base address register (PEX_OWBARn) and translation address register (PEX_OWTLn) are fully defined, a CSB-based memory transaction hitting the defined base address register will be converted to an external PCI Express configuration transaction cycle appeared on the downstream link of the PCI Express RC controller. In this case, the CSB memory address determines the configuration register accessed and the memory data returns the contents of the addressed register. Software must only issue 4-byte or less access to the ATMU configuration window and the access cannot cross a 4-byte boundary.

The mapping from the CSB local address to PCI Express configuration space address is defined in [Table 14-131](#). Note that there is no byte swapping for the address itself, although the programming of the registers content do require byte swapping. The table also translates between the bit ordering commonly used by the Power PC terminology (0:31) and the bit ordering used by the PCI Express terminology (31:0).

Table 14-131. Configuration Address Mapping

| CSB Address Bits | PCI Express Address Bits | PCI Express Configuration Space |
|------------------|--------------------------|---------------------------------|
| 0:7 | 31:24 | Bus number |
| 8:12 | 23:19 | Device number |
| 13:15 | 18:16 | Function number |
| 16:19 | 15:12 | Reserved |
| 20:23 | 11:8 | Extended register number |
| 24:29 | 7:2 | Register number |
| 30:31 | 1:0 | Reserved |

Note: It is the user's responsibility to set the reserved bit fields to zero (especially bits 15–12).

Note: The configuration cycle generation mechanism does not differentiate from internal or external configuration cycle. This means that any transaction which hits a configuration window will be passed to the PCI Express link with a relevant transaction type.

The PCI Express RC controller initiates the Type 0 or Type 1 configuration cycle on its downstream link based on the following rules:

- If the bus number of the CSB-initiated transaction equals the secondary bus number from Type 1 header of RC's configuration space and the device number is 0, a Type 0 configuration cycle will be sent to the link.
- If the bus number of the CSB-initiated transaction does not equal the RC's primary bus number, and does not equal the secondary bus number (from RC's Type 1 header), and is less than or equal to the subordinate bus number (from RC's Type 1 header), a Type 1 configuration cycle will be sent to the link. Note that according to PCI and PCI Express base specifications, the relationship where

the Secondary Bus Number \leq Subordinate Bus Number must be ensured when configuring the two bus numbers within RC's Type 1 header.

- For all other cases, the PCI Express RC controller will issue a configuration cycle on the link whenever an outbound configuration window is hit on the CSB side, no matter what the parameters are. It is the software driver's responsibility to block transactions with unsupported bus, device, and function numbers and return "1"s for such reads. If the bus number in the CSB-initiated transaction equals the primary bus number of RC hitting the outbound configuration window, software error will occur which must be handled by the driver.

The following is an example showing how to configure the related registers of one of the MPC8315E outbound windows for configuration transaction generation purpose. To simplify the illustration, this example has the following assumption:

- The Boot ROM location is configured to locate within 0x0000_0000 to 0x007F_FFFF.
- The overall PCI Express system has a total of 16 buses to be configured.

Sixteen buses mean that the bus number can range from 0x00 to 0x0F. In general, if there are 2^n bus numbers to be configured, n number of address bits are needed to represent the variation of bus number ranging from 0 to $2^n - 1$. For this example, four address bits from CSB[4:7] are required to represent the bus number variation and therefore become the most significant four bits within the total 28 bit offset (CSB[4:31], including reserved bits) of the outbound window to be configured. The CSB[0:3] in this case will not participate in the bus number mapping process as shown in [Table 14-131](#), where all eight bits of CSB[0:7] are mapped to PCI Express address bits [31:24] to represent the possible of 256 bus numbers within a very large system. In other words, in the example, CSB[0:3] become the four most significant bits of the base address of the outbound window to be configured and will be translated to a new "address" in the PCI Express space as defined by the corresponding PEX_OWTARLn.

Based on the above information, the most significant four bits of the base address of the outbound window came from CSB[0:3] must be unique within the total 4 Gbyte local CSB memory space. This essentially defines a window with size of 256 Mbytes, since the lower 28 bits are offset. For this example, assume a 256 Mbyte outbound window is therefore defined between 0x5000_0000 and 0x5FFF_FFFF in local CSB space. This yields the PEX_OWBARn's BA[31:12] to be 0x5000_0, with the actual base address of the outbound window as 0x5000_0000. The SIZE[31:12] field of the PEX_OWARn register is 0x1000_0 to reflect the actual size of this outbound window as 0x1000_0000 or 256 Mbytes.

Note that the final configuration transaction to be generated is based on the information gathered from two areas: the most significant four bits from the defined TAL (translation address low) bit field of PEX_OWTARLn and the lower order bits from the CSB[4:31] offsets directly mapped to PCI Express address bits [27:0]. As mention in the note section of [Table 14-131](#), software should ensure all the reserved bits within CSB[4:31] are filled with zeros.

Since TAL[31:24] (total of eight bits) of PEX_OWTARLn could be used for mapping the bus number bit field for a general configuration transaction, while in this example the bus number to be configured ranges from 0x00 to 0x0F, the most significant four bits (TA[31:28]) are not needed for the bus number mapping and therefore must be configured as 0x0. The TA[27:12] bit field of PEX_OWTARLn is left as all zeros. This yields the PEX_OWTARLn's TAL[31:12] to be 0x0000_0, with the actual translation address of the outbound window as 0x0000_0000. Since the size of this window is 256 Mbytes, the upper limit of the translation address is then locate at 0x0FFF_FFFF.

With the outbound window configured as above, if software intends to scan the PCI Express bus and attempts to probe the first bus immediately underneath the PCI Express RC, the software will need to issue a Configuration Transaction to read Register Number 0 (Vendor ID Register) at Bus Number/Device Number/Function Number of 1/0/0, assuming the RC's secondary bus number has been configured as 0x1 along with subordinate bus number being configured with a big number like 0xFF initially. As long as the LAW is configured to ensure that the local address space between 0x5000_0000 and 0x5FFF_FFFF is configured for PCI Express, a CSB-based memory read transaction to local address 0x5100_0000 initiated by software will be translated to a transaction hitting PCI Express RC controller with PCI Express address of 0x0100_0000. Upon receiving this CSB-based transaction, the RC controller checks the Type attribute of this outbound window and realizes the transaction is of a configuration type, instead of directly using the translated address as in usual memory transaction, the RC controller starts compose a configuration transaction with header information from this translation address based on the mapping defined in [Table 14-131](#). The decode of the mapping process yields a type 0 configuration transaction to be generated on the PCI Express link with Bus Number=1, Device Number=0, Function Number=0, and Register Number=0.

Similarly, if the software intends to read the above EP's configuration space offset 0x440, it can generate a CSB-based memory transaction to address 0x5100_0440. Once the transaction hitting the above configured outbound window, the ATMU translates it into a transaction with PCI Express address 0x0100_0440. This RC controller, once receiving the transaction, will issue a Type 0 configuration transaction on the PCI Express link with Bus Number=1, Device Number=0, Function Number=0, Extended Register Number=0x4, and Register Number=0x40.

14.6.1.8.2 EP Configuration Register Access

When the PCI Express controller is configured as an EP device it responds to remote host generated configuration cycles. This is indicated by decoding the configuration command along with type 0 access in the packet. A remote host can access up to 4096 bytes of the PCI Express configuration area. While in EP mode, the PCI Express controller does not support generating configuration accesses as a master. There is no configuration mechanism supported in EP mode using the ATMU window. If the outbound ATMU window is configured to issue a configuration transaction, all posted transactions hitting this window are ignored and all non-posted transactions will get a response with an error and can lead to unexpected results.

14.6.1.8, 14-105 In [Table 14-131](#), "Configuration Address Mapping," add the following table footnotes:

Note: It is the user's responsibility to set the reserved bit fields to zero (especially bits 15–12).

Note: The configuration cycle generation mechanism does not differentiate from internal or external configuration cycle. This means that any transaction which hits a configuration window will be passed to the PCI Express link with a relevant transaction type.

14.6.1.9.1, 14-105 Add the following text to the first paragraph:

The actual events are logged in the PCI Express Root Error Status Register and in the CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR). See [Section 14.4.5.10](#), "PCI Express Root Error Status Register, [Section 14.5.10.4](#), "CSB System Miscellaneous Interrupt Enable Register (PEX_CSMIER) and

Section 14.5.10.8, “CSB System Miscellaneous Interrupt Status Register (PEX_CSMISR) for further details.

14.6.2.1, 14-107 Revise entire section and its subsections, as follows:

14.6.2.1 EP Interrupt Generation

Both hardware INTx messages generation and Hardware MSI generation are supported, for the interrupt events described in Section 14.5.7, “PCI Express Host Interrupt Registers.” The INTx message and MSI mechanism are mutually exclusive. Only one of these two mechanisms should be enabled and used at a given time.

14.6.2.1.1 Hardware INTx Message Generation

Hardware INTx message is generated when any interrupt event occurs and the corresponding interrupt is enabled in PEX_HIER register, if the interrupt disable bit is cleared in the PCI Express EP’s command register (see Section 14.4.1.3, “PCI Express Command Register”) and the MSI interrupt mechanism is not enabled.

Only INTA message is supported by this device.

Note that MSI interrupt mechanism will be used as long as it is enabled, regardless the setting of the interrupt disable bit in the PCI Express EP’s command register.

14.6.2.1.2 Software INTx Message Generation

Software INTx message generation is not supported.

14.6.2.1.3 Hardware MSI Generation

Host software must set up the MSI capability registers to enable MSI mode and put the correct MSI address and data values into the MSI capability registers prior to setting up various interrupt event enable bits in the PEX_HIER register to enable the generation of the correct MSI cycle to RC.

Note that the value being programmed by the host software into the MSI Data register of EP’s Type 0 configuration space is a 16-bit base message data pattern, which is referred to as “base vector [15:0]” in the description below:

- If only one MSI message is desired, the EP software can directly use this “base vector” as the interrupt vector for MSI interrupt generation. In such case, there is no need to program any of the “vector registers” among PEX_HOPIVR, PEX_HIPIVR, PEX_HWDIVR, PEX_HRDIVR and PEX_HMIVR. The PEX_HIER register setting determines which event can trigger this single MSI interrupt message to RC. Note that multiple interrupt events are allowed to share the same MSI interrupt vector.
- If multiple MSI messages are desired, the Multiple Message Capable bit field of EP’s MSI Message Control Register can be used to indicate how many MSI messages (in the power of two, up to 32 messages allowed per EP) the EP wants to use. During configuration stage, after examining the above desired value, the Host software will allocate the actual number of MSI messages for an EP to use by configuring the Multiple Message Enable bit field in the same register, in addition to programming the “base vector” in EP’s MSI Data Register. Once this is accomplished, the EP

software can program the IVEC bit field of each individual “vector register” based on the number of MSI messages allocated by host. The IVEC value must be unique for the same EP and start from 0x00. At last, the EP software can set the corresponding interrupt event enable bits in the PEX_HIER to enable the Hardware MSI generation. The actual MSI data value for a given interrupt event used by the EP in its MSI message to RC is formed by the concatenation of the “base vector [15:5]” and the IVEC [4:0] value of the corresponding interrupt event.

As an example, if the value of the Multiple Message Enable bit field allocated by host software is 010b (4 MSI messages allocated) and the “base vector” being programmed by host software in EP’s MSI Data Register is 0x55A0 (lower-16bits little-endian), the actual MSI data values of all possible MSI messages can be used by the EP are 0x0000_55A0, 0x0000_55A1, 0x0000_55A2, and 0x0000_55A3. The EP software only needs to program the four possible lower-order bits (0x00, 0x01, 0x02, and 0x03) as unique IVEC values in its “vector registers” among PEX_HOPIVR, PEX_HIPIVR, PEX_HWDIVR, PEX_HRDIVR and PEX_HMIVR. If both OPAIE and OPCIE bit fields are enabled in the PEX_HIER register, assuming PEX_HOPIVR register’s IVEC bit field is programmed as 0x03h, when any one of these two interrupt events occurs, the EP will use 0x0000_55A3 as the actual MSI data value in its MSI message sent to RC. Once receiving such MSI message, the device driver running at RC is responsible to issue a read to EP’s PCI Express Interrupt Status Register (PEX_HISR) to find out exactly which of the two interrupt events caused the interrupt.

14.6.2.1.4 Software MSI Generation

Host software needs to set up the MSI capability registers to enable MSI mode and put the correct values for the MSI address and data register. Next, local software must read the MSI address in the MSI capability register and configure the outbound ATMU window to map the translated address to the MSI address. Software determines the number of allocated messages in the MSI capability register and allocates the appropriate data values to use. A write to the MSI ATMU window with the appropriate data value generates the MSI transaction to the RC.

14.6.1.9.1, 14-105 Modify Table 14-132, “PCI Express RC Inbound Message Handling,” as follows:

Table 14-132. PCI Express RC Inbound Message Handling

| Name | Code[7:0] | Routing[2:0] | Action |
|---------------------|-----------|--------------|------------------|
| Assert_INTA | 0010 0000 | 100 | Set INTA event |
| Assert_INTB | 0010 0001 | 100 | Set INTB event |
| Assert_INTC | 0010 0010 | 100 | Set INTC event |
| Assert_INTD | 0010 0011 | 100 | Set INTD event |
| Deassert_INTA | 0010 0100 | 100 | Clear INTA event |
| Deassert_INTB | 0010 0101 | 100 | Clear INTB event |
| Deassert_INTC | 0010 0110 | 100 | Clear INTC event |
| Deassert_INTD | 0010 0111 | 100 | Clear INTD event |
| PM_Active_State_Nak | 0001 0100 | 100 | No action taken |

Table 14-132. PCI Express RC Inbound Message Handling (continued)

| Name | Code[7:0] | Routing[2:0] | Action |
|---------------------------|-----------|--------------|--|
| PM_PME | 0001 1000 | 000 | Set PM_PME event |
| PME_Turn_Off | 0001 1001 | 011 | No action taken |
| PME_TO_Ack | 0001 1011 | 101 | Log entered_l23_state in PME and message detect register and generate interrupt to IPIC if enabled |
| ERR_COR | 0011 0000 | 000 | Set correctable error event |
| ERR_NONFATAL | 0011 0001 | 000 | Set non-fatal error event |
| ERR_FATAL | 0011 0011 | 000 | Set fatal error event |
| Unlock | 0000 0000 | 000 | No action taken |
| Set_Slot_Power_Limit | 0101 0000 | 100 | No action taken |
| Vendor_Defined Type 0 | 0111 1110 | — | No action taken |
| Vendor_Defined Type 1 | 0111 1111 | — | No action taken |
| Attention_Indicator_On | 0100 0001 | 100 | No action taken |
| Attention_Indicator_Blink | 0100 0011 | 100 | No action taken |
| Attention_Indicator_Off | 0100 0000 | 100 | No action taken |
| Power_Indicator_On | 0100 0101 | 100 | No action taken |
| Power_Indicator_Blink | 0100 0111 | 100 | No action taken |
| Power_Indicator_Off | 0100 0100 | 100 | No action taken |
| Attention_Button_Pressed | 0100 1000 | 100 | Log in PME and message detect register if enabled. Send interrupt if enabled. |

14.6.2.2.2, 14-108 Change references to “MSIR” to say “MSIIR.”

14.6.5, 14-111 Revise section, as follows:

“When a hot reset condition occurs, the controller (in both RC and EP mode) initiates a cleanup of all outstanding transactions and goes into suspend mode. The RC controller then needs to be reset (issuing CBRST in PECR1/PECR2) to bring it back to the idle state followed by a reprogramming of all the CSB bridge registers (offset 0x800–0xFFC, such as the ATMU windows). The EP controller also needs to be reset, followed by a reprogramming of the entire configuration space and CSB bridge registers (offset 0x800–0xFFC, such as the ATMU windows). All configuration register bits that are non-sticky are reset. Link training takes place subsequently. The device is permitted to generate a hot reset condition on the bus when it is configured as an RC device by setting the secondary bus reset bit in the bridge control register in the configuration space. In EP mode, the device is not permitted to generate a hot reset condition; it can only detect a hot reset condition and initiate the cleanup procedure appropriately.”

14.7.1, 14-113 In Figure 14-136, “DMA Descriptor Format,” and Table 14-136, “DMA Descriptor Bit Fields Description,” change bit 3 name to “No snoop for CSB,” and modify field description as follows:
 “No snoop for CSB transactions.
 0 The memory transaction is broadcast on the CSB as non-global (that is, not snooped)
 1 The memory transaction is broadcast on the CSB as global (that is, snooped)”
 In addition, modify field description for bits 127–96 and 95–64, as follows:

| | | | |
|--------|----|---------|--|
| 127–96 | 32 | Control | Destination address. Software programs this field to indicate the destination address. For a write DMA, the destination address is a CSB address that is mapped to a PCI Express memory address using the outbound window address translation. For a read DMA, the destination address is the CSB address. |
| 95–64 | 32 | Control | Source address. Software programs this field to indicate the source address. For a write DMA, the source address is the CSB address. For a read DMA, the source address is a CSB address that is mapped to a PCI Express memory address using the outbound window address translation. |

14.7.2, 14-114 Update Section 14.7.2, “Write DMA,” as follows:

14.7.2 Write DMA

The PCI Express or CSB software can program the write DMA engines to send data from the CSB system to the PCI Express. After the control registers are programmed, the write DMA engine issues a CSB read request through the DMA read master. To improve system performance, the DMA request is segmented according to a natural aligned CSB address boundary of maximum transfer size (32 bytes).

The entire address space accessed by the DMA controller is prefetchable, so the CSB read request for DMA operation always reads all the bytes. All segments use the same ID and are posted without waiting for the previous read response.

When a response to a CSB read request is received, the segments are packed into the PCI Express memory write request according to the PCI Express MPS. All data requested by the DMA controller is processed as a single data stream as described in this section. For example, when a 256-byte request starting from address 0 is segmented to eight 32-byte CSB read requests, then it is segmented to two 128-byte PCI Express write requests (assuming a 128-byte MPS).

If all data can be packed into one PCI Express write request, no segmentation is performed. Otherwise, the first segment starts from the start address and ends at the MPS address boundary. Byte enables are set according to DMA control register settings. All other segments except the last one start from the MPS address boundary and end at the next boundary. That is, the size of a write request is equal to MPS. All bytes are enabled. Remaining data, if applicable from the last MPS address boundary to the PCI Express end address, is packed into the last segment. Byte enables are set according to DMA control register settings. When all segments get an CSB response with a status of OKAY, the DMA data transfer has completed successfully.

If any data within an CSB read response is aborted by the CSB slave (SLVERR response), all the data received before that point is packed and sent. Any remaining data returned from the CSB slave is discarded, and an error is logged.

If any segment gets a response of DECERR, all data received before that point is packed and sent. Any remaining data returned from the CSB slave is discarded, and an error is logged.

The CSB read master can issue multiple pending CSB read requests if the requests belong to a single DMA request. When all segments are successfully received (that is, they get an CSB response with a status of OKAY), the CSB read master starts processing the next DMA request.

14.7.3, 14-115 Update Section 14.7.3, “Read DMA,” as follows:

14.7.3 Read DMA

The read DMA engines can be programmed by the PCI Express or CSB software to send data from the PCI Express system to the CSB system. After the control registers are programmed, the read DMA engine issues a PCI Express read request. The DMA request is segmented according to the PCI Express MRRS natural aligned address boundary. If all data can be requested in one read request, no segmentation is performed. Otherwise, the first segment starts from the start address and ends at the first MRRS boundary. All other segments except the last one start from the MRRS address boundary and end at the next boundary. That is, the length of the read request is equal to MRRS. Any remaining data, if applicable from the last MRRS address boundary to the end address, is requested in the last segment.

The entire address space accessed by the DMA controller is considered as prefetchable, so a PCI Express read request for DMA operation always enables all byte enables. All segments have a unique tag, so multiple read requests can be pending at any given time.

To improve system performance, when a PCI Express completion comes back, it is segmented according to the CSB maximum size natural aligned address boundary. If all data can be packed into one CSB write request, no segmentation is performed, and the write is performed according to the DMA control register settings. Otherwise, the first segment starts from the start address and ends at the first CSB address boundary. All other segments except the last one start from the CSB address boundary and end at the next boundary. That is, the size of the write request is equal to the CSB maximum size packet. All bytes within each data phase are enabled.

Any remaining data, if applicable from the last CSB address boundary to the PCI Express end address, is packed into the last segment.

Each completion is segmented as a single data stream according to these rules. For example, when a 256-byte request starting from address 0 is converted to two 128-byte PCI Express read requests (assuming a 128-byte MPS/MRRS), the bridge issues two read requests if a tag and a completion buffer are available. When any completion comes back, it is segmented as described in this section. No scatter gather is performed. Because PCI Express can return completions in any order, the bridge may not issue an CSB write request in address order.

When all segments get a response with a PCI Express completion that has a status of successful, the DMA data transfer has completed successfully. If any segment gets a response with a status other than successful, or a completion timeout occurs, the DMA data transfer completes with an error status after all requests

complete either normally or abnormally. The data returned for prior requests is still processed and sent to the CSB accordingly.

14.7.4.1, 14-116 Update Section 14.7.4.1, “Chain Descriptor,” as follows:

14.7.4.1 Chain Descriptor

Chain descriptors form an n -way chain in which each descriptor implicitly or explicitly contains the address of the next descriptor. This enables the host to use memory efficiently to store the descriptors even when contiguous memory locations are not available. When the host needs to initiate another transfer, it adds another descriptor in its memory and modifies the pointer of the last descriptor in the chain to the location of the new descriptor. Figure 14-123 illustrates the chain descriptor organization in host memory.




Figure 14-139. n -Way Chain Descriptor Organization in Host Memory

The number of ways, n , is software configurable. In this mode, n descriptors are written in contiguous memory locations. The implicit address of the next descriptor is the next memory location. The last descriptor in the contiguous block contains the explicit address pointer of the next set of n descriptors. Address 0x0 will never be part of the chain since it should close the chain. Non-contiguous valid descriptors are not supported. If the valid bit of a descriptor in the chain is not set, all of the succeeding descriptors should also have the valid bit as zero.

The software need to follow the following sequence on receiving a chain transfer done interrupt:

1. Software receives an interrupt for chain transfer done.
2. Software polls the main memory and waits for descriptor done bit to be set (bit 32 in the descriptor). This guarantees that the final data has been written into memory for read DMA. For write DMA, it guarantees that the final data has been sent to PCI Express controller. If a read is sent to the same location, PCI Express controller guarantees that the outbound read will not bypass the outbound write.

The host can reuse the descriptor memory after the DMA/bridge logic processes it. The exact handshake between the hardware and software is described later in this chapter. The DMA registers are described in Section 12.4.8, “DMA Registers.”

When $n > 1$, the hardware uses this knowledge to prefetch descriptors in advance, thereby reducing possible holes in transmission. This might be useful for applications that request several small DMA transfer requests, such as an Ethernet traffic. For applications that request large data transfers, the effect on performance due to descriptor fetching is not significant.

14.7.4.1, 14-117 Add the following to the end of the second paragraph:
 “Non-contiguous valid descriptors are not supported. If the valid bit of a descriptor in the chain is not set, all of the succeeding descriptors should also have the valid bit as zero.”

14.7.4.4, 14-118 Change the first bullet to read as follows:
 “The DMA controller automatically fetches the same set of descriptors again after the time indicated in PEX_DMA_DSTMR[DSRT] (see [Section 14.5.2.2, “PCI Express DMA Descriptor Timer Register \(PEX_DMA_DSTMR\)”](#)). If the descriptor is ready now, it continues execution or else checks again later.”
 In addition, clarify sentence by adding “prefetched” to: “... it first executes any remaining prefetched valid descriptors...”

15.3.3.2, 15-16 In Figure 15-14, “SATA Interface Error Register (SError)” make bits 10 and 24 reserved as follows:

Offset 0x1_8104 Access: w1c

| DIAG Decode | | | | | | | | | | | | | ERR Decode | | | | | | | | | | | | | | |
|-------------|--|-----------|----|-----|-----|-----|----|-----|-----|-----|-----|-----|------------|-----|-----|----|-----|----|-----|-----|---|---|-----|-----|---|---|---|
| | | 31 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 12 | 11 | 10 | 9 | 8 | 7 | 2 | 1 | 0 | | |
| R | | — | | A | X | F | — | S | H | C | D | B | W | I | N | N | — | E | — | C | T | — | — | M | I | T | G |
| W | | — | | w1c | w1c | w1c | — | w1c | w1c | w1c | w1c | w1c | w1c | w1c | w1c | — | w1c | — | w1c | w1c | — | — | w1c | w1c | — | — | |
| Reset | | All zeros | | | | | | | | | | | | | | | | | | | | | | | | | |

15.3.3.2, 15-17 In Table 15-14, “SError Field Descriptions,” make bits 10 and 24 reserved.

15.3.4.8, 15-26 In Figure 15-24, “PHY Control Configuration Register1 (PhyCtrlCfg1),” and Table 15-24, “PhyCtrlCfg1 Field Descriptions,” change bit 13 (HOST_MODE) to be reserved with the stipulation that the reset value needs to be preserved.

15.3.4.10, 15-29 In Table 15-26, “PhyCtrlCfg2 Field Descriptions,” add the following note to the TX_AMP_CNTRL (bits 24–25) field description: “This bit is available only for SATA1.”

15.5.3.5.2, 15-68 In Table 19-38, “MACCFG2 Field Descriptions,” append the following sentence to the PreAM RxEN] = 1 (bit 24) field description: “If the preamble is less than 7 bytes, 0’s are prepended to pad it to 7 bytes.”

16.1.2, 16-1 In list of features, change first bullet to read, “Support for either two x1 PCI Express or two x1 SGMII”

16.1.2, 16-2 Add “eSATA” to list of features.

16.3, 16-4 In Table 16-2, “SerDes PHY Block Memory Map,” change reset value for SRDSCR1 to 0x0000_0040, and change reset value for SRDSCR2 to 0x0080_0000.

16.3.1, 16-5 In Table 16-3, “SRDSCR0 Field Descriptions,” remove from RXEQA and RXEQE bit field bit descriptions: references/explanation about RXEQA[2] and RXEQA[3], RXEQE[2] and RXEQE[3], respectively.

RXEQA bit field description now reads:

“Receive equalization selection bus for lane A—when asserted in PCI Express mode:

- 00 No equalization
- 01 2 dB of equalization

10 4 dB of equalization

11 Reserved

Recommended setting per protocol:

PCI Express: 01

SGMII: 01”

RXEQE bit field description now reads:

“Receive equalization selection bus for lane E—when asserted in PCI Express mode:

00 No equalization

01 2 dB of equalization

10 4 dB of equalization

11 Reserved

Recommended setting per protocol:

PCI Express: 01

SGMII: 01”

16.3.2, 16-8

In Figure 16-4, “SerDes Control Register 1 (SRDSCR1),” change reset value, as follows:




Figure 16-4. SerDes Control Register 1 (SRDSCR1)

| Signal | I/O | Description |
|----------------|-----|--|
| USBDR_PWRFAULT | I | Power fault. USBDR_PWRFAULT indicates whether a power fault occurred on the USB port Vbus. |
| | | State Meaning Asserted—Indicates that a Vbus fault occurred. Applications that support power switching must shut down Vbus power. Negated—Indicates normal operation. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_PCTL0 | O | Port control 0. USBDR_PCTL0 controls the port status indicator LED 0 when in host mode. |
| | | State Meaning Asserted—LED on. Negated—LED off. |
| | | Timing Synchronous to PHY_CLK. |
| USBDR_PCTL1 | O | Port control 1. USBDR_PCTL1 controls the port status indicator LED 1 when in host mode. |
| | | State Meaning Asserted—LED on. Negated—LED off. |
| | | Timing Synchronous to PHY_CLK. |

17.3, 17-5 In Table 17-3, “USB Interface Memory Map,” modify reset value of DCCPARAMS, as follows:

Table 17-3. USB Interface Memory Map

| Offset | Register | Access | Reset | Section/Page |
|----------|--|--------|-------------|--------------------------------|
| 0x2_3124 | DCCPARAMS—Device controller parameters | R | 0x0000_0183 | 17.3.1.6/17-11 |

17.3, 17-6 In Table 17-3, “USB Interface Memory Map,” replace the following USBDR register reset values:

- Address offset=0x2_3104: Value=0x0001_0011
- Address offset=0x2_3140: Value=0x0008_0000
- Address offset=0x2_3184: Value=0x1000_0000
- Address offset=0x2_31A4: Value=0x0000_0C20

17.3.1.3, 17-8 In Figure 17-4, “Host Controller Structural Parameters (HCSPARAMS),” change reset value to 0x0001_0011.

17.3.2.1, 17-11 In Figure 17-8, “USB Command Register (USBCMD),” change reset value to 0x0008_0000.

17.3.1.6, 17-11 Modify reset value in Figure 17-7, “Device Control Capability Parameters (DCCPARAMS),” as follows:

1 32-byte fetch”

17.7.1, 17-124

In Figure 17-61, “Endpoint Queue Head Layout,” change “Total Bytes” field to be 15-bits-wide instead of 14-bits-wide as follows:

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|--------------------------------------|--------------------------|-----|----|----|----|-----------------------|----|----|----|-----------------------------|------------------|----|-----|----|--------------------|--------|----|--------------------|---------------------|-------------------|----|---|-------------------|---|---|---|---|------|---|---|---|--------|
| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | offset |
| Mult | | zlt | | 00 | | Maximum Packet Length | | | | | | | | | | ios | | 000_0000_0000_0000 | | | | | | | | | | 0x00 | | | | |
| Current dTD Pointer ¹ | | | | | | | | | | | | | | | | 0_0000 | | 0x04 | | | | | | | | | | | | | | |
| Next dTD Pointer ¹ | | | | | | | | | | | | | | | | 0000 | | T ¹ | 0x08 ² | | | | | | | | | | | | | |
| 0 | Total Bytes ¹ | | | | | | | | | | ioc ¹ | | 000 | | MultO ¹ | | 00 | | Status ¹ | | | | 0x0C ² | | | | | | | | | |
| Buffer Pointer (Page 0) ¹ | | | | | | | | | | Current Offset ¹ | | | | | | | | | | 0x10 ² | | | | | | | | | | | | |
| Buffer Pointer (Page 1) ¹ | | | | | | | | | | Reserved | | | | | | | | | | 0x14 ² | | | | | | | | | | | | |
| Buffer Pointer (Page 2) ¹ | | | | | | | | | | Reserved | | | | | | | | | | 0x18 ² | | | | | | | | | | | | |
| Buffer Pointer (Page 3) ¹ | | | | | | | | | | Reserved | | | | | | | | | | 0x1C ² | | | | | | | | | | | | |
| Buffer Pointer (Page 4) ¹ | | | | | | | | | | Reserved | | | | | | | | | | 0x20 ² | | | | | | | | | | | | |
| Reserved | | | | | | | | | | | | | | | | 0x24 | | | | | | | | | | | | | | | | |
| Setup Buffer Bytes 3–0 ¹ | | | | | | | | | | | | | | | | 0x28 | | | | | | | | | | | | | | | | |
| Setup Buffer Bytes 7–4 ¹ | | | | | | | | | | | | | | | | 0x2C | | | | | | | | | | | | | | | | |

¹ Device controller read/write; all others read-only.

² Offsets 0x08 through 0x20 contain the transfer overlay.

17.7.2, 17-126

In Figure 17-62, “Endpoint Transfer Descriptor (dTD),” change “Total Bytes” field to be 15-bits-wide instead of 14-bits-wide as follows:

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------------------------|--------------------------|----|----|----|----|----|----|----|----|-----------------------------|-----|---------------------------|-----|----|-------|------|----|----|---------------------|------|----|---|------|---|---|---|---|---|---|---|---|--------|
| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | offset |
| Next Link Pointer | | | | | | | | | | | | | | | | 0000 | | T | 0x00 | | | | | | | | | | | | | |
| 0 | Total Bytes ¹ | | | | | | | | | | ioc | | 000 | | MultO | | 00 | | Status ¹ | | | | 0x04 | | | | | | | | | |
| Buffer Pointer (Page 0) | | | | | | | | | | Current Offset ¹ | | | | | | | | | | 0x08 | | | | | | | | | | | | |
| Buffer Pointer (Page 1) | | | | | | | | | | 0 | | Frame Number ¹ | | | | | | | | 0x0C | | | | | | | | | | | | |
| Buffer Pointer (Page 2) | | | | | | | | | | 0000_0000_0000 | | | | | | | | | | 0x10 | | | | | | | | | | | | |
| Buffer Pointer (Page 3) | | | | | | | | | | 0000_0000_0000 | | | | | | | | | | 0x14 | | | | | | | | | | | | |
| Buffer Pointer (Page 4) | | | | | | | | | | 0000_0000_0000 | | | | | | | | | | 0x18 | | | | | | | | | | | | |

¹ Device controller read/write; all others read-only.

18.1, 18-3

Remove the figure below the first paragraph.

18.1.2, 18-5 Modify Section 18.1.2 and rename it “Polychannel Overview,” as follows:

18.1.2 Polychannel Overview

The polychannel block implements four channels for processing descriptors. Each channel contains the following addressable structures:

- A fetch FIFO, which holds a queue of pointers to descriptors waiting to be processed
- A configuration register, which allows the user a number of options for SEC event signalling
- Monitor registers containing information about the transaction in process
- A status register containing an indication of the last unfulfilled bus request
- A descriptor buffer memory used to store the active descriptor and other temporary

Whenever a channel is idle and its fetch FIFO is non-empty, the channel reads the next descriptor pointer from the fetch FIFO. Using this pointer, the channel fetches the descriptor and places it in its descriptor buffer. The channel’s processing of descriptors is described in more detail in #####”

A channel can signal to the host that it is done with a descriptor via interrupt and/or by a writeback of the descriptor header into host memory. In the case of writeback, the value written back is identical to the header that was read, with the exception that a DONE byte is set to 0xFF. The channels’ done signaling is described in more detail in #####”.

An EU operation can include generating an ICV and then comparing it against a received ICV. The result of the ICV checking can be signalled to the host either via interrupt or by a writeback of the descriptor header (but not by both methods). In the case of writeback, the user can opt to do it at end of every descriptor, or only at the end of descriptors that call for ICV checking.

In case of an error condition in a channel or its reserved EUs, the channel issues an interrupt to the host. The channel can be configured to either abort the current descriptor and proceed to the next one, or halt and wait for host intervention.

For more about configuring signalling see [Section 18.5.4.1, “Channel Configuration Register \(CCR\),”](#) and for detail on the writeback fields see [Section 18.3.4, “Link Table Format.”](#)

Many security protocols involve both encryption and hashing of packet payloads. To accomplish this without requiring two passes through the data, channels can configure data flows through two EUs. In such cases, one EU is designated the “primary EU”, and the other as the “secondary EU”. The primary EU receives its data from memory via the controller, and the secondary EU receives its data by “snooping” the SEC buses.

There are two types of snooping:

- Input data can be fed to the primary EU and the same input data snooped by the secondary EU. This is called “in-snooping”.
- Output data from the primary EU can be snooped by the secondary EU. This is called “out-snooping”.

In the SEC, only MDEU and CRCU are used as secondary EUs.

For more information, refer to [Section 18.5, “Polychannel.”](#)

- 18.1.3, 18-7 Modify Section 18.1.3 and rename it “Controller Overview,” as follows:
- “The controller manages the master and slave interfaces to the system bus and the internal buses that connect all the various modules. It receives service requests from the host (via the slave interface) and from the channels, and schedules the required data transfers. The system bus interface and access to system memory are critical factors in performance, and the 64-bit master and slave interfaces of the SEC controller enable it to achieve performance unattainable on secondary buses. The controller provides for two ways of operating the execution units, channel-controlled access and host-controlled access.
- In channel-controlled access (the SEC’s normal operating mode), all interactions with EUs are directed by a channel executing a descriptor. The host is involved only in initially supplying the descriptor pointer and in handling results once descriptor processing is complete.
 - In host-controlled access (intended primarily for debug purposes), the host moves data in and out of execution units directly through memory-mapped EU registers. No descriptor is involved.

For more information about the controller (including more details about channel-controlled and host-controlled access), refer to #####.”

- 18.2, 18-12 Remove reference to IMMRBAR.

- 18.3, 18-18 Modify section, as follows:
- “The host processor maintains a record of current secure sessions and the corresponding keys and contexts of those sessions. Once the host has determined that a security operation is required, it creates a “descriptor” containing all the information the SEC needs to perform the security operation. The host creates the descriptor in main memory, then writes a pointer to the descriptor into the Fetch FIFO of one of the SEC channels. The channel uses this pointer to read the descriptor into its descriptor buffer. Once it obtains the descriptor, the SEC uses its bus mastering capability to obtain inputs and write results, thus off-loading data movement and encryption operations from the host processor.
- Descriptors are only used in channel-controlled accesses to SEC, and not in host-controlled accesses. For more information about host-controlled access, see #####.”

- 18.3.1, 18-19 Revise Section 18.3.1, “Descriptor Structure,” as follows:

18.3.1 Descriptor Structure

SEC descriptors are conceptually similar to descriptors used by most devices with DMA capability. SEC descriptors are designed to support the cryptographic computation of a single packet using a single descriptor. SEC descriptors have a fixed length of 64 bytes, such as, eight 64-bit words (referred to as dwords). A descriptor consists of one “header dword” and seven “pointer dwords,” as seen in [Figure 18-3](#).

| | 0 | 15 | 16 | 17 | 23 | 24 | 27 | 28 | 31 | 32 | 63 | |
|---------------------|---------------|----|---------|----|-------|----------|----|----|----|-----------------|----|--|
| Header Dword | Header | | | | | | | | | Reserved | | |
| Pointer Dword 0 | Length0 | J0 | Extent0 | — | Eptr0 | Pointer0 | | | | | | |
| Pointer Dword 1 | Length1 | J1 | Extent1 | — | Eptr1 | Pointer1 | | | | | | |
| Pointer Dword 2 | Length2 | J2 | Extent2 | — | Eptr2 | Pointer2 | | | | | | |
| Pointer Dword 3 | Length3 | J3 | Extent3 | — | Eptr3 | Pointer3 | | | | | | |
| Pointer Dword 4 | Length4 | J4 | Extent4 | — | Eptr4 | Pointer4 | | | | | | |
| Pointer Dword 5 | Length5 | J5 | Extent5 | — | Eptr5 | Pointer5 | | | | | | |
| Pointer Dword 6 | Length6 | J6 | Extent6 | — | Eptr6 | Pointer6 | | | | | | |

Figure 18-3. Descriptor Format

As shown in [Figure 18-3](#), the first and second halves of the header dword are denoted as descriptor control and descriptor feedback fields, respectively. The descriptor control field of the header dword specifies the security operation to be performed, the execution unit(s) needed, and the modes for each execution unit. The descriptor feedback field is written to by the security engine upon completion of descriptor processing, when the “channel done writeback” feature is enabled. Further details about the header dword may be found in [Section 18.3.2, “Descriptor Format: Header Dword.”](#)

The pointer dwords, all of which have the same format, contain pointer and length information for locating input or output parcels (such as keys, context, or text data). The large number of pointers provided in the descriptor allows for multi-algorithm operations that require fetching of multiple keys, as well as fetch and return of contexts. Any pointer dword that is not needed may be assigned a length of zero. Further details about the pointer dwords may be found in [Section 18.3.3, “Descriptor Format: Pointer Dwords.”](#)

SEC descriptors include scatter/gather capability, which means that each pointer in a descriptor can be either a direct pointer to a contiguous parcel of data, or a pointer to a “link table” which is a list of pointers and lengths used to assemble the parcel. When a link table is used to read input data, this is referred to as a “gather” operation; when used to write output data, it is referred to as a “scatter” operation. Further details about scatter/gather capability may be found in [Section 18.3.4, “Link Table Format.”](#)

Revision History

18.5.4.1, 18-34 Change the title of Figure 18-10 to “Channel Status Register (CSR),” and change the fields to be Read/Write, as follows:

Offset Channel_1 0x01110 Access: Read/Write
 Channel_2 0x01210
 Channel_3 0x01310
 Channel_4 0x01410




Figure 18-10. Channel Status Register (CSR)

18.5.4.2, 18-38 In Table 18-13, “CSR Error Field Definitions,” modify field descriptions for WDT, RSI, and RSG, as follows. Note that the following rows are only part of Table 18-13:

Table 18-13. CSR Error Field Definitions

| Value | Error | Description |
|-------|-------|--|
| 56 | WDT | Watchdog timeout. The main state machine stayed asleep too long. This timer runs only after EUs have been reserved, and does not run if the primary EU is the RNGU or PKEU. The timeout interval is controlled by the FCC field of the Channel Configuration Register. This error halts the channel. Restarting the channel clears this bit. |
| 58 | RSI | RAID Size Incorrect. The channel was provided with a descriptor of type RAID_XOR with data sizes not permitted. To clear this error, the host must write a ‘1’ to this bit. |
| 59 | RSG | RAID scatter gather error. The channel was provided with a descriptor of type RAID_XOR with a j bit set. Use of scatter/gather is not permitted with RAID_XOR type descriptors. To clear this error, the host must write a ‘1’ to this bit. |

18.7.1.10, 18-63 Modify Table 18-27, “AESU Context Registers,” as follows:

Table 18-7. AESU Context Registers

| Context Register # (byte address) | Cipher Mode | | | | | | |
|-----------------------------------|-------------|--------------------------|-----|-----------------------|--------------|--------------|----------------------|
| | ECB | CBC/CBC-R BP/OFB/CF B128 | CTR | CCM | GCM | GCM - GHASH | SRT |
| 1 (0x34100) | — | IV ¹ | -- | IV ¹ / MAC | Computed MAC | Computed MAC | Counter ¹ |
| 2 (0x34108) | | | | | | | |

Table 18-7. AESU Context Registers (continued)

| Context Register # (byte address) | Cipher Mode | | | | | | |
|--------------------------------------|-------------|--------------------------------|---|---|------------------------------|-----------------------|---------------------------------|
| | ECB | CBC/CBC-R BP/OFB/CF B128 | CTR | CCM | GCM | GCM - GHASH | SRT |
| 3 (0x34110) | — | — | — | Encrypted MAC ² / Decrypted MAC / Encrypted Counter | Received MAC ² | — | Counter Modulus ¹ |
| 4 (0x34118) | | | | — | | | |
| 5 (0x34120) | — | — | Counter ¹ | Counter ¹ | Counter | — | — |
| 6 (0x34128) | | | | | | | |
| 7 (0x34130) | — | — | Counter Modulus Exponent ¹ | Counter Modulus Exponent ¹ (header size/ MAC size) ³ | len(AAD) ⁴ | len(AAD) ⁴ | — |
| 8 (0x34138) | | | | — | | | |
| 9 (0x34140) | — | — | — | — | Y ₀ | H | — |
| 10 (0x34148) | | | | | | | |
| 11 (0x34150) | — | — | — | — | len(AAD) ⁵ | len(AAD) ⁵ | — |
| 12 (0x34158) | | | | | len(IV) ⁵ | — | |

¹ Must be written at start of new message, except if zero

² Needed only in ICV mode—must be written at start of new CCM decryption

³ The header and MAC sizes are internally constructed by the AES engine; then, that information is included inside context register 7 for context switching purposes

⁴ Length of total data (in bits)

⁵ Length of data processed with current descriptor (in bits)

— Don't care

18.7.1.10.1, 18-64 Modify the bullets as follows:

- Context Register 1 holds the *most* significant bytes of the initialization vector (bytes 1–8).
- Context Register 2 holds the *least* significant bytes of the initialization vector (bytes 9–16).

18.7.2.6, 18-81 In Figure 18-35, “CRCU Control Register,” change access for all fields to Read/Write, as follows:




Figure 18-35. CRCU Control Register

- 18.7.6.9, 18-123 Add a note, as follows:
 “Host reads of the RNGB FIFO should be performed on an 8-byte basis, regardless of how many bits of random number is actually required. Partial host reads can leave the RNGB FIFO in a state that will result in a channel error.”
- Chapter 19 Remove all references to extraction of data to allocate in the L2 cache, throughout.
- 19.1, 19-1 Add the following note after the second paragraph:
 “The eTSECs do not support TBI, GMII, and FIFO operating modes, so all references to these interfaces and features should be ignored for this device.”
- 19.2, 19.4 Change the last bullet to say the following: “Hardware assist for 1588-compliant timestamping (1588 not supported in conjunction with SGMII 10/100)”
- 19.4, 19-6 In Table 19-1, “eTSEC_n Network Interface Signal Properties,” for RGMII and RTBI protocols, changed description of TSEC_RX_ER from “Unused, output driven low” to “Unused”
- 19.4, 19-7 In Table 19-1, “eTSEC_n Network Interface Signal Properties,” add the following row for TSEC_TMR_PP3:

Table 19-1. eTSEC_n Network Interface Signal Properties

| Signal Name | Function | Reset State |
|--------------|------------------|-------------|
| TSEC_TMR_PP3 | 1588—Pulse out 3 | 0 |

- 19.4.1, 19-10 In Table 19-2, “eTSEC Signals—Detailed Signal Descriptions,” add the following row for TSEC_TMR_PP3:

Table 19-2. eTSEC Signals—Detailed Signal Descriptions

| Signal | I/O | Description |
|--------------|-----|--|
| TSEC_TMR_PP3 | O | 1588 pulse out 3. Timer pulse per period 3. It is phase aligned with 1588 timer clock (chip external output pin) |

- 19.5, 19-11 Add the following paragraph to the end of the section:
 “The ten-bit interface (TBI) and reduced ten-bit interface (RTBI) module MII registers are also described in this section. The TBI/RTBI registers are defined like PHY registers and, as such, are accessed through the MII management interface in the same way the PHYs are accessed. For detailed descriptions of the TBI/RTBI registers (the MII register set for the ten-bit interface) refer to [Section 19.5.4, “Ten-Bit Interface \(TBI\).”](#)”
- 19.5.1, 19-11 In Table 19-3, “Module Memory Map Summary” made bits A00–AFF reserved as FIFO mode is not supported for 8315. The table now shows:

| Address Offset | Function |
|----------------|---|
| 000–0FF | eTSEC general control/status registers |
| 100–2FF | eTSEC transmit control/status registers |
| 300–4FF | eTSEC receive control/status registers |
| 500–5FF | MAC registers |
| 600–7FF | RMON MIB registers |
| 800–8FF | Hash table registers |
| 900–AFF | — |
| B00–BFF | DMA system registers |
| C00–C3F | Lossless Flow Control registers |
| C40–DFF | — |
| E00–EFF | 1588 Hardware Assist |

- 19.5.2, 19-13 Modify Table 19-4, “Module Memory Map” for the following registers. Note that only affected rows are shown.

Table 19-4. Module Memory Map

| eTSEC1 Offset | Name | Access | Reset | Section/Page |
|-------------------|---|--------|-------------|------------------------------------|
| 0x2_4030 | TBIPA—TBI PHY address register | R/W | All zeros | 19.5.3.1.9/19-36 |
| 0x2_4034–0x2_40FC | Reserved | — | — | — |
| 0x2_4338 | RQFCR*—Receive queue filing table control register | R/W | 0xnnnn_nnnn | 19.5.3.3.7/19-57 |
| 0x2_433C | RQFPR*—Receive queue filing table property register | R/W | 0xnnnn_nnnn | 19.5.3.3.8/19-58 |
| 0x2_4730 | CAR1—Carry register one register | w1c | All zeros | 19.5.3.6.44/14-100 |
| 0x2_4734 | CAR2—Carry register two register | w1c | All zeros | 19.5.3.6.45/14-101 |
| 0x2_4E04 | TMR_TEVENT* - time stamp event register | w1c | All zeros | 19.5.3.10.2/14-112 |

- 19.5.2, 19-15 In Table 19-4, “Module Memory Map,” remove the row containing the ATRELLI register.

- 19.5.4, 19-22 Modify the first sentence to say the following:
 “This section describes the ten-bit interface (TBI), reduced ten-bit interface (RTBI), and the TBI/RTBI MII set of registers. TBI and RTBI operate in the same manner (the only difference is that RTBI has reduced I/O signalling).”
- 19.5.3.1.2, 19-23 In Figure 19-3, “TSEC_ID2 Register,” and in Table 19-6, “TSEC_ID2 Field Descriptions,” extend the size of TSEC_INT field to be 8–15 bits. Update the field description to say: “Interface mode support. See [Table 19-7](#) for settings.”
- 19.5.3.1.2, 19-24 Add Table 19-7, “TSEC_ID2[TSEC_INT] Field Settings,” as follows:

Table 19-7. TSEC_ID2[TSEC_INT] Field Settings

| Bit | Mode |
|-------|--|
| 8 | 0 1588 protocol not supported 1 1588 protocol supported |
| 9 | 0 SGMII not supported 1 SGMII supported |
| 10 | 0 Ethernet mode not supported 1 Ethernet mode supported |
| 11–13 | Reserved |
| 14 | 0 Can be configured to run in Ethernet normal/full mode 1 Ethernet normal/full mode off |
| 15 | 0 Can be configured to run in Ethernet reduced mode 1 Ethernet reduced mode off |

- 19.5.3.1.3, 19-25 Add the following sub bullet to the end of the bulleted list #3:
 — Special function interrupts are: FGPI, MSRO, MMRD, and MMRW
- 19.5.3.1.3, 19-26 In Table 19-7, “IEVENT Field Descriptions,” replace second sentence of the CRL (bit 14) field description to say the following “The frame is discarded without being transmitted and the queue halts (TSTAT[THLTn] set to 1).”
- 19.5.3.1.6, 19-32 In Figure 19-7, “ECNTRL Register Definition,” and Table 19-10, “ECNTRL Field Descriptions,” make bit 16 reserved.
- 19.5.3.1.6, 19-32 In Table 19-10, “ECNTRL Field Descriptions,” update CLRCNT (bit 17) field description to read as follows:
 “Clear all statistics counters and carry registers.
 0 Allow MIB counters to continue to increment and keep any overflow indicators.
 1 Reset all MIB counters and CAR1 and CAR2.
 This bit is self-resetting.”
- 19.5.3.1.6, 19-32 In Table 19-10, “ECNTRL Field Descriptions,” update AUTOZ (bit 18) field description to read as follows:
 “Automatically zero MIB counter values and carry registers.
 0 The user must write the addressed counter zero after a host read.

1 The addressed counter value is automatically cleared to zero after a host read. This is a steady state signal and must be set prior to enabling the Ethernet controller and must not be changed without proper care.”

19.5.3.1.6, 19-32

Modify Figure 19-7, “ECNTRL Register Definition,” and Table 19-10, “ECNTRL Field Descriptions,” as follows. Note that only the affected rows are shown:




Figure 19-7. ECNTRL Register Definition

Table 19-10. ECNTRL Field Descriptions

| Bits | Name | Description |
|------|-------|--|
| 25 | GMIIM | GMII interface mode. Not supported. |
| 26 | TBIM | (Reduced) ten-bit interface mode. If this bit is set, reduced ten-bit interface (RTBI) mode is enabled. This bit can be pin-configured at reset to set or clear. See Section 4.4.2, “Power-On Reset Flow.” 0 MII or RMIIM mode interface 1 RTBI mode interface |
| 29 | RMM | Reduced-pin mode for 10/100 interfaces. If this bit is set, an RMIIM pin interface is expected. RMM must be 0 if RPM = 1. This register can be pin-configured at reset to 0 or 1 0 Non-RMIIM interface mode 1 RMIIM interface mode |

19.5.3.1.6, 19-33

In Table 19-10, “ECNTRL Field Descriptions,” add the following note to ECNTRL[R100M] bit description:
“This bit must be cleared for 1-Gbps SGMII operation.”

19.5.3.1.6, 19-34

Modify Table 19-11, “eTSEC Interface Configurations,” as follows:

| Interface Mode | ECNTRL Field | | | | | | | MACCFG2 Field |
|-----------------|-------------------|--------------------|------|-----|-------|-----|--------|---------------|
| | FIFM ¹ | GMIIM ² | TBIM | RPM | R100M | RMM | SGMIIM | I/F Mode |
| MII 10/100 Mbps | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 01 |
| RMIIM 100 Mbps | 0 | 0 | 0 | 0 | 1 | 1 | 0 | 01 |
| RMIIM 10 Mbps | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 01 |
| RGMII 1Gbps | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 10 |

| Interface Mode | ECNTRL Field | | | | | | | MACCFG2 Field |
|-----------------|-------------------|--------------------|------|-----|-------|-----|--------|---------------|
| | FIFM ¹ | GMIIM ² | TBIM | RPM | R100M | RMM | SGMIIM | I/F Mode |
| RGMIIM 100 Mbps | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 01 |
| RGMIIM 10 Mbps | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 01 |
| RTBI 1Gbps | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 10 |
| SGMIIM 1 Gbps | 0 | 0 | 1 | 0 | 0 | 0 | 1 | 10 |
| SGMIIM 100 Mbps | 0 | 0 | 1 | 0 | 1 | 0 | 1 | 01 |
| SGMIIM 10 Mbps | 0 | 0 | 1 | 0 | 0 | 0 | 1 | 01 |

¹ FIFM bit is not supported.

² GMIIM bit is not supported.

19.5.3.1.8, 19-36 In Table 19-13, “DMACTRL Field Descriptions,” change TOD field definition as follows:

“1 eTSEC immediately fetches a new TxBD from ring 0.”

19.5.3.1.9, 19-36 Add Section 19.5.3.1.9, “TBI Physical Address Register (TBIPA),” as follows:

19.5.3.1.9 TBI Physical Address Register (TBIPA)

The TBIPA, shown in Figure 19-10, is writable by the user to assign a physical address to the TBI (or RTBI) for MII management configuration. The TBI registers are accessed at the offset of TBIPA. For detailed descriptions of the TBI registers (the MII register set for the ten-bit interface) refer to Section 19.5.4, “Ten-Bit Interface (TBI).”

Offset eTSEC1:0x2_4030; eTSEC2:0x2_5030;
eTSEC3:0x2_6030; eTSEC4:0x2_7030

Access: Read/Write




Figure 19-10. TBIPA Register Definition

Table 19-14 describes the fields of the TBIPA register.

Table 19-14. TBIPA Field Descriptions

| Bits | Name | Description |
|-------|-------|---|
| 0–26 | — | Reserved |
| 27–31 | TBIPA | This field is used to program the PHY address of the ten-bit interface’s MII management bus. To access the TBI register the user must write the TBIPA value to the MIIMADD [PHY Address] register located in the MAC register section. PHY Address 0 is reserved. Refer to Section 19.5.3.5.8, “MII Management Address Register (MIIMADD).” |

19.5.3.2.1, 19-38 In Table 19-14, “TCTRL Field Description,” change TXSCHED field description for 01 state to read as follows:

- 19.5.3.2.2, 19-38 “01 Priority scheduling mode. Frames from enabled TxBD rings are serviced in ascending ring index order.”

In Table 19-15, “TSTAT Field Descriptions,” add the following to the field descriptions of THLT n :

“Repeatable error conditions which cause halt include:

Bus error:

 - Invalid BD or data address
 - Uncorrectable error on BD or data read”
- 19.5.3.2.4, 19-43 In Table 19-17, “TXIC Field Descriptions,” remove references to CCB clock.
- 19.5.3.3.1, 19-49 Add description of bit 16 in Table 19-25, “RCTRL Field Descriptions,” as follows:

| Bits | Name | Description |
|------|------|-------------|
| 16 | — | Reserved |

- 19.5.3.3.3, 19-53 In Table 19-27, “RXIC Field Descriptions,” remove references to CCB clock.
- 19.5.3.3.4, 19-54 Modify Figure 19-24, “RQUEUE Register Definition,” and Table 19-28, “RQUEUE Field Descriptions,” as follows:




Figure 19-24. RQUEUE Register Definition

Table 19-28. RQUEUE Field Descriptions

| Bits | Name | Description |
|------|------|--|
| 0–23 | — | Reserved |
| 24 | EN0 | Receive queue 0 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 25 | EN1 | Receive queue 1 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 26 | EN2 | Receive queue 2 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 27 | EN3 | Receive queue 3 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |

Table 19-28. RQUEUE Field Descriptions (continued)

| Bits | Name | Description |
|------|------|--|
| 28 | EN4 | Receive queue 4 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 29 | EN5 | Receive queue 5 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 30 | EN6 | Receive queue 6 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |
| 31 | EN7 | Receive queue 7 enable. 0 RxBd ring is not queried for reception. In effect the receive queue is disabled. 1 RxBd ring is queried for reception. |

- 19.5.3.3.5, 19-55 In Table 19-29, “RBIFX Field Descriptions,” change the definition of the 10 encoding of *BnCTL* to: “Byte 0 is located in the received frame at offset B0OFFSET bytes from the byte after the last byte of the layer 2 header.”
In addition, change the definition of the 11 encoding of *BnCTL* to: “Byte 0 is located in the received frame at offset B0OFFSET bytes from the byte after the last byte of the layer 3 header.”
- 19.5.3.3.7, 19-57 In Figure 19-27, “Receive Queue Filer Table Control Register Definition,” change the default reset value to be undefined.
- 19.5.3.3.8, 19-58 In Figure 19-29, “Receive Queue Filer Table Property ID1 Register Definition,” change the default reset value to be undefined.
- 19.5.3.5.1, 19-67 In Table 19-37, “MACCFG1 Field Descriptions,” add the following text to RxFlow (bit 26) and TxFlow (bit 27) field descriptions: “Must be 0 if MACCFG2[Full Duplex] = 0.”
- 19.5.3.5.1, 19-67 In Table 19-37, “MACCFG1 Field Descriptions,” add the following note to fields Tx_Flow and Rx_Flow (bits 26 and 27): “Should not be set when operating in Half-Duplex mode.”
- 19.5.3.5.2, 15-68 In Table 19-38, “MACCFG2 Field Descriptions,” update the following bit field descriptions. Note that only affected rows are shown:

Table 19-38. MACCFG2 Field Descriptions

| Bits | Name | Description |
|-------|-----------------|--|
| 16–19 | Preamble Length | This field determines the length in bytes of the preamble field preceding each Ethernet start-of-frame delimiter byte. Values from 0x3 to 0xF are supported by the controller. The default value of 0x7 should not be altered in order to guarantee reliable operation with IEEE 802.3 compliant hardware. |
| 24 | PreAM RxEN | User defined preamble enable for received frames. This bit is cleared by default. 0 The MAC skips the Ethernet preamble without returning it. 1 The MAC recovers the received Ethernet preamble and passes it to the driver at the start of each received frame. If the preamble is less than 7 bytes, 0’s are prepended to pad it to 7 bytes. Not applicable to or RMII 10/100 modes. |

Table 19-38. MACCFG2 Field Descriptions

| Bits | Name | Description | | | | | | | | | | | | | | | | | | | | |
|---------------------|------------------------|--|---------------------------|--------------|-------------------|---------------------------|---------------------|------------------------|-----|-----|---------|------------------------|----|----|----------|------------------------|----|-----|---------------------|------------------------|----|----|
| 26 | Huge Frame | <p>Huge frame enable. This bit is cleared by default.</p> <p>0 Limit the length of frames received to less than or equal to the maximum frame length value (MAXFRM[Maximum Frame]) and limit the length of frames transmitted to less than the maximum frame length. See Section 19.6.7, “Buffer Descriptors,” for further details of buffer descriptor bit updating.</p> <table border="1"> <thead> <tr> <th>Frame type</th> <th>Frame length</th> <th>Packet truncation</th> <th>Buffer descriptor updated</th> </tr> </thead> <tbody> <tr> <td>Receive or transmit</td> <td>> maximum frame length</td> <td>yes</td> <td>yes</td> </tr> <tr> <td>Receive</td> <td>= maximum frame length</td> <td>no</td> <td>no</td> </tr> <tr> <td>Transmit</td> <td>= maximum frame length</td> <td>no</td> <td>yes</td> </tr> <tr> <td>Receive or transmit</td> <td>< maximum frame length</td> <td>no</td> <td>no</td> </tr> </tbody> </table> <p>1 Frames are transmitted and received regardless of their relationship to the maximum frame length. Note that if Huge Frame is cleared, the user must ensure that adequate buffer space is allocated for received frames. See Section 19.5.3.5.5, “Maximum Frame Length Register (MAXFRM),” for further information.</p> | Frame type | Frame length | Packet truncation | Buffer descriptor updated | Receive or transmit | > maximum frame length | yes | yes | Receive | = maximum frame length | no | no | Transmit | = maximum frame length | no | yes | Receive or transmit | < maximum frame length | no | no |
| Frame type | Frame length | Packet truncation | Buffer descriptor updated | | | | | | | | | | | | | | | | | | | |
| Receive or transmit | > maximum frame length | yes | yes | | | | | | | | | | | | | | | | | | | |
| Receive | = maximum frame length | no | no | | | | | | | | | | | | | | | | | | | |
| Transmit | = maximum frame length | no | yes | | | | | | | | | | | | | | | | | | | |
| Receive or transmit | < maximum frame length | no | no | | | | | | | | | | | | | | | | | | | |
| 29 | PAD/CRC | <p>Pad and append CRC. This bit is cleared by default. This bit must be set when in half-duplex mode (MACCFG2[Full Duplex] is cleared).</p> <p>0 Frames presented to the MAC have a valid length and contain a CRC.</p> <p>1 The MAC pads all transmitted short frames and appends a CRC to every frame regardless of padding requirement.</p> | | | | | | | | | | | | | | | | | | | | |

- 19.5.3.5.4, 19-70 In Figure 19-37, “Half-Duplex Register Definition,” and in Table 19-40, “HAFDUP Field Descriptions,” change field size of “collision window” to be 26–31 bits wide, and make bits 20–25 Reserved.
- 19.5.3.5.5, 19-71 In Table 19-41, “MAXFRM Descriptions” remove reference to FIFO from Maximum Frame (bits 16–31) field description.
- 19.5.3.5.5, 19-71 In Table 19-41, “MAXFRM Descriptions,” modify the first paragraph of “Maximum Frame” (bits 16–31) field description, as follows:
 “This field is set to 0x0600 (1536 bytes) by default and always must be set to a value greater than or equal to 0x0040 (64 bytes), but not greater than 0x2580 (9600 bytes). It sets the maximum Ethernet frame size in both the transmit and receive directions. (Refer to MACCFG2[Huge Frame].) It does not affect the size of packets sent or received via the FIFO packet interface.
- 19.5.3.5.6, 19-72 In Table 19-42, “MIIMCFG Field Descriptions,” remove references to CCB clock.
- 19.5.3.5.6, 19-72 In Table 19-42, “MIIMCFG Field Descriptions,” update the note in the MIIMCFG[MgmtClk] field description to read as follows:
 The eTSEC system clock is selected by the SCCR register. (See Chapter 4, “Reset, Clocking, and Initialization.”)

Revision History

- 19.5.3.5.9, 19-74 Change Figure 19-42, “MII Mgmt Control Register Definition,” to be a write-only register.
- 19.5.3.6, 19-79 Add the following note to the end of this section:

NOTE

The transmit and receive frame counters (TR64, TR127, TR 255, TR511, TR1K, TRMAX, and TRMGV) do not increment for aborted frames (collision retry limit exceeded, late collision, underrun, EBERR, TxFIFO data error, frame truncated due to exceeding MAXFRM, or excessive deferral).

- 19.5.3.6.17, 19-87 Table 19-69, “RFLR Field Descriptions,” add the following content to the RFLR (bits 16–31) field description: “Frames tagged with a single VLAN tag are checked for valid length based on bytes 17–18 (rather than 13–14). Frames tagged (stacked) with multiple VLAN tags are not checked for valid length.”
- 19.5.3.6.25, 19-91 Modify the second sentence of the TBYT field description in Table 19-77, “TBYT Field Descriptions,” from:
 “This count does not include preamble/SFD or jam bytes.”
 to read as:
 “This count does not include preamble/SFD or jam bytes, except for half-duplex flow control (back-pressure triggered by TCTRL[THDF]=1). For THDF, the sum total of ‘phantom’ preamble bytes transmitted for flow control purposes is included in the TBYT increment value of the next frame to be transmitted, up to 65,535 bytes of frame and phantom preamble.”
- 19.5.3.6.41, 19-99 In Table 19-93, “TOVR Field Descriptions,” update TOVR[TOVR] field description to read as follows:
 “Transmit oversize frame counter. Increments each time a frame is transmitted which exceeds 1518 bytes (non VLAN) or 11522 bytes (VLAN) with a correct FCS value.”
- 19.5.3.6.44, 19-100 In Figure 19-93, “Carry Register 1 (CAR1) Register Definition,” change access from “Read only” to “w1c.”
- 19.5.3.6.45, 19-101 In Figure 19-94, “Carry Register 2 (CAR2) Register Definition,” change access from “Read only” to “w1c.”
- 19.5.3.8.1, 19-107 Modify Section 19.5.3.8.1, as follows:

19.5.3.8.1 Attribute Register (ATTR)

The attribute register defines memory access attributes and transaction types used to access buffer descriptors, to write receive data, and to read transmit data. Snoop enable attributes may be set for reading buffer descriptors and for reading transmit data.

Figure 19-100 describes the definition for the ATTR register.




Table 19-100. ATTR Register Definition

Table 19-103 describes the fields of the ATTR register.

Table 19-103. ATTR Field Descriptions


| Bits | Name | Description |
|-------|---------|---|
| 0–23 | — | Reserved |
| 24 | RDSSEN | Rx data snoop enable. 0 Disables snooping of all receive frames data to memory 1 Enables snooping of all receive frames data to memory. |
| 25 | RBDSSEN | RxBD snoop enable 0 Disables snooping of all receive BD memory accesses. 1 Enables snooping of all receive BD memory accesses. |
| 26–31 | — | Reserved |

- 19.5.3.8.2, 19-108 Remove Section 19.5.3.8.2, “Attribute Extract Length and Extract Index Register (ATTRELI).”
- 19.5.3.9.2, 19-110 In Figure 19-103, “RFBPTR0–RFBPTR7 Register Definition,” modify offset to say: “eTSEC1:0x2_4C44+8×n; eTSEC2:0x2_5C44+8×n; eTSEC3:0x2_6C44+8×n; eTSEC4:0x2_7C44+8×n”
- 19.5.3.10, 19-110 Add the following note after the third paragraph:
“IEEE 1588 timestamping is not supported in conjunction with the SGMII 10/100 interface mode.”
- 19.5.3.10.1, 19-111 In Table 19-107, “TMR_CTRL Register Field Descriptions,” modify first sentence of ETEP2 (bit 22) field description, as follows: “External trigger 2 edge polarity,” and modify first sentence of ETEP1 (bit 23) field description, as follows: “External trigger 1 edge polarity.”
- 19.5.3.10.1, 19-111 In Table 19-107, “TMR_CTRL Register Field Descriptions,” add the following text to the end of TCLK_PERIOD (bits 6–15) field description:
“For nanosecond granularity on 1588 timer counter rate, the TCLK_PERIOD should be calculated using the following equation:
 $TCLK_PERIOD = 10^9 / \text{Nominal_Frequency}$ ”
- 19.5.3.10.1, 19-111 In Table 19-107, “TMR_CTRL Register Field Descriptions,” replace CIPH (bit 25) field description with the following:
“Oscillator input clock phase.

| | | |
|----------------------|---|---|
| | 0 | non-inverted timer input clock |
| | 1 | inverted timer input clock (NOTE: this setting is reserved if CKSEL=01.)” |
| 19.5.3.10.2, 19-112 | | In Figure 19-105, “TMR_TEVENT Register Definition,” change access to w1c. |
| 19.5.3.10.9, 19-118 | | In Figure 19-110, “TMR_ACC Register Definition,” change access from “Read only” to “Read/Write.” |
| 19.5.3.10.12, 19-119 | | In Figure 19-113, “TMR_ALARM1-2_H/L Register Definition,” change access from “Mixed” to “Read/Write.” |
| 19.5.3.10.13, 19-121 | | In Figure 19-114, “TMR_FIPERn Register Definition,” change access from “Mixed” to “Read/Write.” |
| 19.5.4.3.10, 19-134 | | In Table 19-134, “TBICON Field Descriptions,” update Clock Select (bit 10) field description, as follows: “Clock select. This bit selects how the on-chip TBI PHY is clocked. This bit is cleared by default. |
| | 0 | The TBI PHY is clocked by dual split-phase 62.5 MHz receive clocks. These external signals must be provided via TBI receive clock 0 (TSEC _n _RX_CLK) and TBI receive clock 1 (TSEC _n _TX_CLK). If operating in SGMII mode, clearing this bit effectively disables the TBI PHY clock. |
| | 1 | The TBI PHY is clocked by a single 125 MHz receive clock (required for SGMII operation). This single clock, if operating in a non-SGMII (parallel) mode, must be provided via the TBI receive clock 0 (TSEC _n _RX_CLK) external signal. If operating in SGMII mode, this clock is provided on-chip by the SerDes block.” |
| 19.6.1.4, 19-138 | | Add the following section describing the connection to an SGMII physical interface. |

19.6.1.4 Serial Gigabit Media-Independent Interface (SGMII)

This section describes the serial gigabit media-independent interface (SGMII) intended to be used between a SerDes PHY and the eTSEC to implement a serial gigabit version of a media-independent interface. [Figure 19-130](#) depicts the basic components of the SGMII including the signals required to establish eTSEC module connection with a PHY. Note that in SGMII the eTSEC utilizes the on-board TBI PHY in addition to the SerDes interface.



¹ The management signals (MDC and MDIO) may be common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

Figure 19-130. eTSEC-SGMII Connection

19.6.1.5, 19-141

Add the following Table 19-138, “SGMII Signalling” after Table 19-137, “RGMII Signals Multiplexing.” Note that the subsequent tables are re-numbered.

Table 19-138. SGMII Signalling

| Signals | I/O | No. of Signals | Function |
|--------------------------|-----|----------------|------------------------------------|
| SD _n _RX[n] | I | 2 | SGMII receive data (differential) |
| SD _n _TX[n] | O | 2 | SGMII transmit data (differential) |
| SD _n _REF_CLK | I | 2 | Reference clock (differential) |
| Sum | | 6 | — |

19.6.2.9, 19-153

Update second sentence of third paragraph to say, “The controller completes any frame in progress before stopping transmission and does not commence counting the pause time until transmit is idle.”

19.6.2.10, 19-153

Modify the sub-bullets underneath the first bullet, as follows:

- Receive data frame interrupts, when bits RXB or RXF in IEVENT are set
- Transmit data frame interrupts, when bits TXB or TXF in IEVENT are set
- Error, diagnostic, and special interrupts (all bits in IEVENT other than RXB, RXF, TXB, or TXF)

Revision History

- | | |
|-------------------|--|
| 19.6.2.11, 19-156 | <p>Add the following to bulleted list after the first paragraph:</p> <ul style="list-style-type: none"> • All BDs for any multiple-BD frame reside in the same cache line. • TCP/UDP and IP Checksum generation are disabled in each frame's TxFCB, or in TCTRL, or frames are limited to 1200 bytes in length. |
| 19.6.2.13, 19-157 | <p>In Table 19-147, “Reception Errors,” add the following note to “Parser error” description: “Any values in the length/type field between 1500 and 1536 is treated as a length, however, only illegal packets exist with this length/type since these are not valid lengths and not valid types. These are treated by the MAC logic as out of range.”</p> |
| 19.6.4, 19-163 | <p>Add Section 19.6.4.1, “Receive Parser,” as follows:</p> |

19.6.4.1 Receive Parser

The receive parser parses the incoming frame data and generates filer properties and frame control block (FCB). The receive parser composes of the Ethernet header parser and L3/L4 parser.

The Ethernet header parser parses only L2 (ethertype) headers. It is enabled by RCTRL[PRSDEP] != 0. It has the following key features:

- Extraction of 48-bit MAC destination and source addresses
- Extraction and recognition of the first 2-byte ethertype field
- Extraction and recognition of the final 2-byte ethertype field
- Extraction of 2-byte VLAN control field
- Walk through MPLS stack and find layer 3 protocol
- Walk through VLAN stack and find layer 3 protocol
- Recognition of the following ethertypes for inner layer parsing
 - LLC and SNAP header
 - JUMBO and SNAP header
 - IPV4
 - IPV6
 - VLAN
 - MPLSU/MPLSM
 - PPPOES

For stack L2 (that is, more than one ethertypes) header, the Ethernet parser traverses through the header until it finds the last valid ethertype or the ethertype is unsupported. [Table 19-150](#) describes what the Ethernet header parser recognizes for stack L2 header.

Table 19-150. Supported Stack L2 Ethernet Headers

| Column—Current L2 Ethertype Row—Next Supported L2 Ethertype | LLC/ SNAP | JUMBO/ SNAP | IPV4 | IPV6 | VLAN | MPLSU | MPLSM | PPOES |
|--|--------------|----------------|------|------|------|-------|-------|-------|
| LLC/SNAP | N | N | Y | Y | Y | Y | Y | Y |
| JUMBO/SNAP | N | N | Y | Y | Y | Y | Y | Y |
| IPV4 | N | N | N | N | N | N | N | N |
| IPV6 | N | N | N | N | N | N | N | N |
| VLAN | Y | Y | Y | Y | Y | Y | Y | Y |
| MPLSU | N | N | Y* | Y* | N | y | Y | N |
| MPLSM | N | N | Y* | Y* | N | Y | Y | N |
| PPOES | N | N | Y | Y | N | Y | Y | N |

Note: * means that it is the next protocol

The L3 parser is enabled by `RCTRL[PRSDEP] = 10` or `11`. It begins when the Ethernet parser ends and a valid IPv4/v6 ethertype is found. The L4 header is enabled by `RCTRL[PRSDEP] = 11`. It begins when the L3 parser ends and a valid TCP/UDP next protocol is found and no fragment frame is found. The primary functionalities of L3(IPv4/6) and L4(TCP/UDP) parsers are as follows:

- IP recognition (v4/v6, encapsulated protocol)
- IP header checksum verification
- IPv4/6 over IPv4/6 (tunneling)—parse headers and find layer 4 protocol
- IP layer 4 protocol/next header extraction
- Stop parsing on unrecognized next header/protocol
- IPv4 support
 - IPv4 source and destination addresses
 - 8-bit IPv4 type of service
 - IP layer 4 protocol / next header support
 - IPV4
 - IPV4 Fragment. Parser stops after a fragment is found
 - TCP/UDP
- IPv6 support
 - The first 4 bytes of the IPv6 source address extraction
 - The first 4 bytes of the IPv6 destination address extraction
 - IPv6 source address hash for pseudo header calculation

Revision History

- IPv6 destination address hash for pseudo header calculation
- 8-bit IPv6 traffic class field extraction
- Payload length field extraction
- IP layer 4 protocol/next header support
 - IPV6
 - IPV6 fragment. Parser stops after a fragment is found
 - IPV6 route
 - IPV6 hop/destination
 - TCP/UDP
- L4 (TCP/UDP) support
 - Extraction of 16-bit source port number extraction
 - Extraction of 16-bit destination port number extraction
 - TCP checksum calculation (including pseudo header)
 - UDP checksum calculation if the checksum field is not zero (including pseudo header)

19.6.4.1.3, 19-166 In Table 19-150, “Special Filer Rules,” add RQCTRL[GPI] column as follows:

| Rule Description | RQCTRL Fields | | | | | | | RQPROP Word | RQCTRL Word ¹ |
|--|---------------|------------------|-----|-----|---------|-----|------|-------------|--------------------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | |
| Default file—Always file frame to ring Q | 0 | 0 | 0 | 0 | Q | 01 | 0000 | All zeros | 0x0000_ <i>qq</i> 20 |
| Default reject—Always discard frame | 0 | 0 | 1 | 0 | 000_000 | 01 | 0000 | All zeros | 0x0000_0120 |
| Empty rule in AND—Always matches | 0 | 0/1 ² | 0 | 1 | 000_000 | 01 | 0000 | 0xFFFF_FFFF | 0x0000_00A0 |
| Empty rule in rule set—Always fails | 0 | 0/1 ³ | 0 | 0 | 000_000 | 11 | 0000 | 0xFFFF_FFFF | 0x0000_0060 |

¹ Hexadecimal digits *qq* denotes field Q shifted left 2 bits.

² Set CLE = 1 if the empty rule guards a cluster.

³ Set CLE = 1 if the empty rule occurs at the end of a cluster.

19.6.4.1.6, 19-167 In Table 19-152, “Filer Table Example—802.1p Priority Filing,” add RQCTRL[GPI] column as follows.

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---------------------------|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 0 | 0 | 0 | 0 | 0 | 000_000 | 00 | 1001 | 0x0000_0007 | File priority 7 to ring 0 | 0x0000_0009 |
| 1 | 0 | 0 | 0 | 0 | 000_001 | 00 | 1001 | 0x0000_0006 | File priority 6 to ring 1 | 0x0000_0409 |
| 2 | 0 | 0 | 0 | 0 | 000_010 | 00 | 1001 | 0x0000_0005 | File priority 5 to ring 2 | 0x0000_0809 |
| 3 | 0 | 0 | 0 | 0 | 000_011 | 00 | 1001 | 0x0000_0004 | File priority 4 to ring 3 | 0x0000_0C09 |
| 4 | 0 | 0 | 0 | 0 | 000_100 | 00 | 1001 | 0x0000_0003 | File priority 3 to ring 4 | 0x0000_1009 |
| 5 | 0 | 0 | 0 | 0 | 000_101 | 00 | 1001 | 0x0000_0002 | File priority 2 to ring 5 | 0x0000_1409 |

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|--|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 6 | 0 | 0 | 0 | 0 | 000_110 | 00 | 1001 | 0x0000_0001 | File priority 1 to ring 6 | 0x0000_1809 |
| 7 | 0 | 0 | 0 | 0 | 000_111 | 00 | 1001 | All zeros | File undefined 802.1p or priority 0 to ring 7—Default always matches | 0x0000_1C09 |

19.6.4.1.7, 19-168 In Table 19-153, “Filer Table Example—IP Diff-Serv Code Points Filing,” add RQCTRL[GPI] column as follows.

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 0 | 0 | 0 | 0 | 0 | 001_000 | 01 | 1010 | 0x0000_00E0 | File class 7 to queue 8 (TOS >= 0xE0) | 0x0000_202A |
| 1 | 0 | 0 | 0 | 0 | 001_001 | 01 | 1010 | 0x0000_00C0 | File class 6 to queue 9 (TOS >= 0xC0) | 0x0000_242A |
| 2 | 0 | 0 | 0 | 0 | 001_010 | 01 | 1010 | 0x0000_00A0 | File class 5 to queue 10 (TOS >= 0xA0) | 0x0000_282A |
| 3 | 0 | 0 | 0 | 0 | 001_011 | 01 | 1010 | 0x0000_0080 | File class 4 to queue 11 (TOS >= 0x80) | 0x0000_2C2A |
| 4 | 0 | 0 | 0 | 0 | 000_100 | 01 | 1010 | 0x0000_0060 | File class 3 to queue 4 (TOS >= 0x60) | 0x0000_102A |
| 5 | 0 | 0 | 0 | 0 | 001_100 | 01 | 1010 | 0x0000_0040 | File class 2 to queue 12 (TOS >= 0x40) | 0x0000_302A |
| 6 | 0 | 0 | 0 | 0 | 010_100 | 01 | 1010 | 0x0000_0020 | File class 1 to queue 20 (TOS >= 0x20) | 0x0000_502A |
| 7 | 0 | 0 | 0 | 0 | 011_100 | 01 | 1010 | All zeros | File class 0 to queue 28 (TOS >= 0x00) or file to ring 4 by default | 0x0000_702A |

19.6.4.1.8, 19-169 In Table 19-154, “Filer Table Example—TCP and UDP Port Filing,” add RQCTRL[GPI] column as follows:

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 0 | 0 | 1 | 0 | 1 | 000_000 | 00 | 1011 | 0x0000_0006 | Enter cluster if layer 4 is TCP | 0x0000_028B |
| 1 | 0 | 0 | 0 | 1 | 000_000 | 01 | 1111 | 0x0000_0014 | AND rule—FTP from TCP ports 20 and 21: file to ring 2 | 0x0000_00AF |
| 2 | 0 | 0 | 0 | 0 | 000_010 | 11 | 1111 | 0x0000_0016 | | 0x0000_086F |
| 3 | 0 | 0 | 0 | 0 | 000_011 | 00 | 1111 | 0x0000_0017 | telnet from TCP port 23: file to ring 3 | 0x0000_0C0F |
| 4 | 0 | 0 | 0 | 0 | 000_000 | 00 | 1111 | All zeros | <i>empty entry reserved for future use</i> | 0x0000_000F |
| 5 | 0 | 0 | 0 | 0 | 000_000 | 00 | 1111 | All zeros | <i>empty entry reserved for future use</i> | 0x0000_000F |
| 6 | 0 | 1 | 0 | 0 | 000_001 | 01 | 0000 | All zeros | end cluster; default TCP: file to ring 1 | 0x0000_0620 |

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment | RQCTRL Word |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|--|-------------|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | | |
| 7 | 0 | 1 | 0 | 1 | 000_000 | 00 | 1011 | 0x0000_0011 | Enter cluster if layer 4 is UDP | 0x0000_028B |
| 8 | 0 | 0 | 0 | 0 | 000_101 | 00 | 1111 | 0x0000_0801 | NFS from UDP port 2049 | 0x0000_140F |
| 9 | 0 | 0 | 0 | 0 | 000_111 | 00 | 1111 | 0x0000_0208 | Route from UDP port 520 | 0x0000_000F |
| 10 | 0 | 0 | 0 | 0 | 000_110 | 00 | 1111 | 0x0000_0045 | TFTP from UDP port 69 | 0x0000_180F |
| 11 | 0 | 1 | 0 | 0 | 000_100 | 01 | 0000 | All zeros | End cluster; default UDP: file to ring 4 | 0x0000_1220 |
| 12 | 0 | 0 | 0 | 0 | 000_000 | 01 | 0000 | All zeros | By default, file to ring 0 | 0x0000_0020 |

19.6.4.1.9, 19-169 Add the following filer example section describing interrupt from deep sleep. The subsequent sections and tables are re-numbered.

19.6.4.1.9 Filer Example—Interrupt from Deep Sleep

The example in [Table 19-155](#) shows how the filer can facilitate exit from deep sleep if any of the following packets arrive:

- ARP packet with Target IP address matching either of two IP addresses (either static or link local address)
- IPv4/UDP multicast DNS query
- IPv4/UDP SNMP broadcast query

These packets are also be stored in memory. All other packets are dropped.

Table 19-155. Filer Example—Interrupt from Deep Sleep

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|--|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | |
| 0 | 0 | 0 | 0 | 0 | 000_000 | 11 | 0000 | 0x0001_0000 | Check for ARP request; set mask register to mask off everything but the ARP request flag. |
| 1 | 0 | 1 | 0 | 1 | 000_000 | 00 | 0001 | 0x0001_0000 | Check to see if ARP request flag is set by doing a =1 comparison. Enter the “ARP Request Cluster” if true. |
| 2 | 0 | 0 | 0 | 0 | 000_000 | 11 | 0000 | 0xFFFF_FFFF | ARP Cluster: Set Mask to unmask everything (Reset mask to all F's) |
| 3 | 1 | 0 | 0 | 0 | 000_001 | 00 | 1100 | 0XXXXX_XXXX | Compare the ARP Target IP address against “MY_IP_1”, which is indicated by the user-defined value of 0XXXXX_XXXX; if they match, accept the frame and generate an interrupt. |
| 4 | 1 | 0 | 0 | 0 | 000_001 | 00 | 1100 | 0YYYYY_YYYY | Compare the ARP Target IP address against “MY_IP_2”, which is indicated by the user-defined value of 0YYYYY_YYYY; if they match, accept the frame. |

Table 19-155. Filer Example—Interrupt from Deep Sleep (continued)

| Table Entry | RQCTRL Fields | | | | | | | RQPROP | Comment |
|-------------|---------------|-----|-----|-----|---------|-----|------|-------------|---|
| | GPI | CLE | REJ | AND | Q | CMP | PID | | |
| 5 | 0 | 1 | 1 | 0 | 000_000 | 01 | 0000 | 0x0000_0000 | Default rule that will always discard the packet; inserted here because an ARP request was received, but the Target IP address did not match either local IP addresses; therefore drop the packet and exit the cluster. |
| 6 | 0 | 0 | 0 | 0 | 000_000 | 11 | 0000 | 0x0000_02D0 | Set Mask for IP4 Packet (2D0), with IPv4 checksum checked and verified, and UDP header located. |
| 7 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0001 | 0x0000_02D0 | Check to see if IP4 Packet, with IPv4 checksum checked and verified, and UDP header. |
| 8 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0000 | 0xFFFF_FFFF | Set Mask to unmask everything (Reset mask to all F's). |
| 9 | 0 | 1 | 0 | 1 | 000_000 | 00 | 1011 | 0x0000_0011 | Check against UDP protocol; if this passes, enter the cluster - all packets in the cluster are IPv4 packets with UDP protocol identified as the L4 protocol type. |
| 10 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0011 | 0x00XX_XXXX | Compare upper L2 DA bits to XX_XXXX (for multicast DNS query) |
| 11 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0100 | 0x00YY_YYYY | Compare lower L2 DA bits to YY_YYYY |
| 12 | 0 | 0 | 0 | 1 | 000_000 | 00 | 1100 | 0xZZZZ_ZZZZ | Compare L3 Destination IP address to ZZZZ_ZZZZ |
| 13 | 0 | 0 | 0 | 1 | 000_000 | 00 | 1110 | 0x0000_XXXX | Compare L4 destination port to XXXX |
| 14 | 1 | 0 | 0 | 0 | 000_001 | 00 | 1111 | 0x0000_YYYY | If all of the previously consecutive ANDed conditions pass, multicast DNS Query has matched. |
| 15 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0011 | 0x00XX_XXXX | Compare upper L2 DA bits to XX_XXXX (for SNMP broadcast query). |
| 16 | 0 | 0 | 0 | 1 | 000_000 | 00 | 0100 | 0x00YY_YYYY | Compare lower L2 DA bits to YY_YYYY. |
| 17 | 1 | 0 | 0 | 0 | 000_001 | 00 | 1110 | 0x0000_ZZZZ | If all of the previously consecutive ANDed conditions pass, SNMP broadcast Query has matched. |
| 18 | 0 | 1 | 1 | 0 | 000_000 | 01 | 0000 | 0x0000_0000 | Cluster End: IPv4, UDP Comparison Default rule that will always discard the packet; inserted here because an IPv4 packet with L4=UDP request was received, but the profiles didn't match anything "interesting"; therefore drop the packet and exit the cluster. |
| 19 | 0 | 0 | 1 | 0 | 000_000 | 01 | 0000 | 0x0000_0000 | Default rule that will always discard the packet; inserted here no matches for any "interesting" packets were received that are used to wake up the CPU. All packets that reach this rule are discarded. |

Revision History

19.6.4.2.1, 19-170 Replace Section 19.6.4.2.1, “Priority-Based Queuing (PBQ),” with the following:
 “PBQ is the simplest scheduler decision policy. The enabled TxBD rings are assigned a priority value based on their index. Rings with a lower index have precedence over rings with higher indices, with priority assessed on a frame-by-frame basis. For example, frames in TxBD ring 0 have higher priority than frames in TxBD ring 1, and frames in TxBD ring 1 have higher priority than frames in TxBD ring 2, and so on.

The scheduling decision is then achieved as follows:

```

loop
  # start or S/W clear of TSATn
  ring = 0;
  while ring <= 7 loop
    if enabled(ring) and not ring_empty(ring) then
      transmit_frame(ring);
      ring = 0;
    else
      ring = ring + 1;
    endif
  endloop
endloop

```

19.6.5, 19-171 Remove sentence, “The current hardware initiated back-pressure mechanism (the RX_PAUSE register) is designed to allow for momentary backlogging of the Rx DMA (due to host memory contention), not a lack of Rx BDs.”

19.6.6, 19-174 Add the following note after the third paragraph of this section:
 “NOTE:
 IEEE 1588 timestamping is not supported in conjunction with the SGMII 10/100 interface mode.”

19.6.7, 19-177 Revise section, as follows:
 “The eTSEC buffer descriptor (BD) is modeled after the MPC8260 Fast Ethernet controller BD for ease of reuse across the PowerQUICC network processor family. Drawing from the MPC8260 FEC BD programming model, the eTSEC descriptor base registers point to the beginning of BD rings. The eTSEC BD also expands upon the MPC8260 BD model to accommodate the eTSEC’s unique features. However, the 8-byte data BD format is designed to be compatible with the existing MPC8260 BD model.”

19.6.6.5, 19-177 Add Section 19.6.6.5, “Time-Stamp Insertion on Transmit Packets,” as follows:

19.6.6.5 Time-Stamp Insertion on Transmit Packets

Software has the option to write the time stamp of the transmitted frame to memory in the padding alignment bytes (PAL) located between the TxFCB and the frame data. It is required that a minimum of two TxBDs are used. The first points to the start of the 8 byte TxFCB. The second points to the start of frame data. In memory, the TxFCB, and at least the first 16 bytes of the TxPAL must be adjacent, such as, located in contiguous memory locations, as depicted in [Figure 19-56](#).

The first TxBD[TOE] bit is set. When the TMR_CTRL[Record Time-stamp In PAL Enable] and TxFCB[PTP] bits are set, the timestamp is written to memory location TxBD[Data Buffer Pointer]+16.

The second TxBD's Data Length must either contain the full frame length, or a value greater than the TxThreshold setting. Refer to [Table 19-141](#). When time-stamps are inserted into the TxPAL, the TMR_TXTSn_H/L and TMR_TXTSn_ID registers still function normally.

19.6.6.5.1 Interrupts

The TxPAL is updated with a time-stamp before closing the second TxBD. The TxBD[I] bit can be set for the second TxBD frame to cause an interrupt (via IEVENT[TXF]) after the time-stamp has been written to the TxPAL.

When time-stamps are inserted into the TxPAL, the TMR_TXTSn_H/L and TMR_TXTSn_ID registers still function normally. Therefore, the 1588 interrupt can be triggered by using the TMR_PEVENT register bits TXP1, and TXP2.

Table 19-56. Time-Stamp Insertion Programming Requirements

| Requirement | Behavior if requirement is not met |
|--|--|
| TMR_CTRL[RTPE]=1 | If TMR_CTRL[RTPE]=0, then no time-stamp is written to a TxPAL. |
| TxBD[TOE]=1 | If TxBD[TOE]=0, then no time-stamp is written to a TxPAL. |
| First TxBD[Data Buffer Pointer] is 8-byte aligned | The time-stamp will be written to address First TxBD[Data Buffer Pointer] + 0x10 rounded down to the nearest 8-byte aligned address, except at 4K page boundaries, in which case the time-stamp may be invalid, and the Second TxBD close status will be lost. |
| First TxBD[Data Length]=8, 8 bytes for TxFCB | If L2 or frame data is included in the Length, the buffer immediately following the FCB is transmitted on the line and the frame data stored in memory will be overwritten with a time-stamp value after the frame is transmitted. |
| TxFCB[PTP]=1 | If TxBD[PTP]=0, then no time-stamp is written to a TxPAL. |
| The TxFCB is followed immediately by a minimum of 16 bytes for the TxPAL | The time-stamp will be written to address First TxBD[Data Buffer Pointer] + 0x10. |
| Second TxBD[Data Buffer Pointers] points to start of L2 or frame data | If there is only one TxBD used to transfer a PTP frame, then no time-stamp is written to a TxPAL. |
| Second TxBD[Data Length] >= FIFO_TX_THR or includes the entire frame | If this condition is not true, the time-stamp in TxPAL is invalid. |

Figure 19-141 depicts the buffer format requirements for time-stamp insertion on transmit packets.

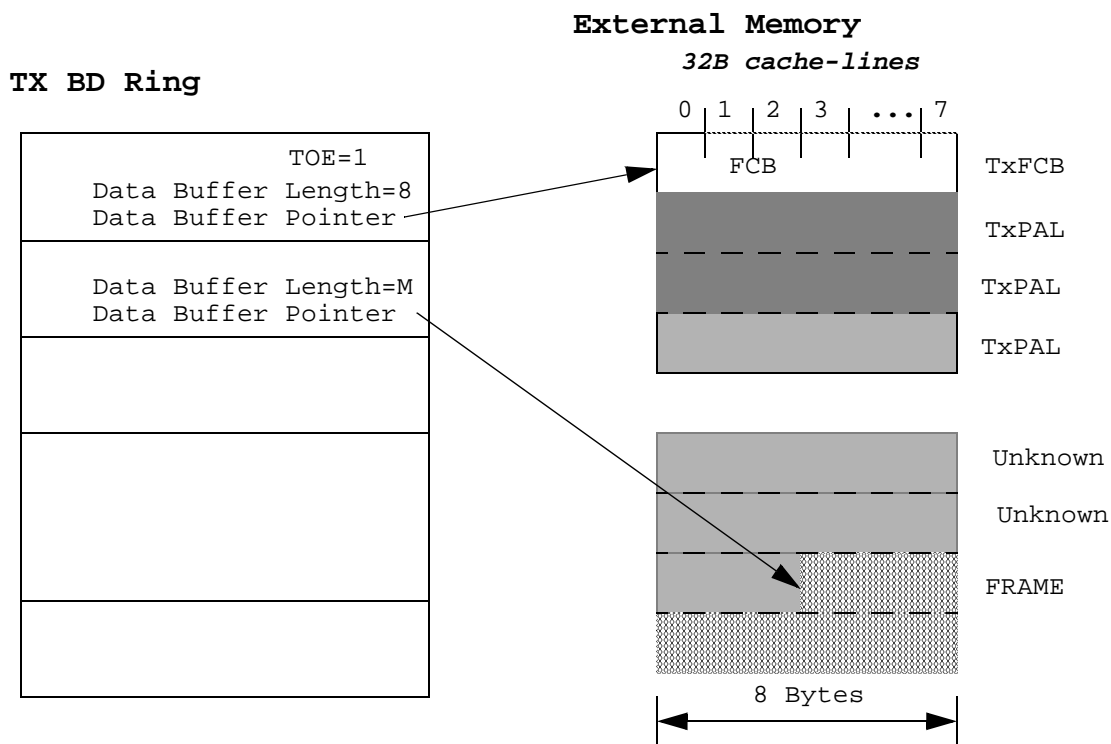


Figure 19-141. Buffer Format for Transmit Time-Stamp Insertion

19.6.6.5.2 Error Condition

When an error is encountered after a PTP packet has begun to be processed, the time-stamp written to the TxPAL is zero. Subsequent frames may be flushed by eTSEC. There will be no time-stamp update to TxPAL for the subsequent flushed frames.

19.6.6.6, 19-177 Add Section 19.6.6.6, “Tx PTP Packet Parsing,” as follows:

19.6.6.6 Tx PTP Packet Parsing

Software instructs the Tx packet to be timestamped via setting bit 15 in the TxFCB to mark a PTP packet. TxFCB[VLCTL] can be translated as the Tx PTP packet identification number. BD[TOE] must be set to enable transmit PTP packet time stamping. TxFCB[PTP] bit takes precedence over TxFCB[VLN] bit. It disables per packet VLAN tag insertion. On a PTP packet, a VLAN tag can be inserted from the DFVLAN register. The TxFCB for the PTP packet is shown in [Figure 19-142](#).

| | | | | | | | | | | | | | | | | | |
|------------|--------------|----|-----|-----|-----|-----|-----|------|---|---|----|----|----|----|----|----|-----|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | |
| Offset + 0 | VLN | IP | IP6 | TUP | UDP | CIP | CTU | NPH | | | | | | | | | PTP |
| Offset + 2 | L4OS | | | | | | | L3OS | | | | | | | | | |
| Offset + 4 | PHCS | | | | | | | | | | | | | | | | |
| Offset + 6 | VLCTL/PTP_ID | | | | | | | | | | | | | | | | |

Figure 19-142. Transmit Frame Control Block

The contents of the Tx FCB are defined in [Table 19-57](#).

Table 19-57. Tx Frame Control Block Description

| Bytes | Bits | Name | Description |
|-------|------|------|---|
| 0–1 | 0 | VLN | VLAN control word valid. This bit is ignored when the PTP bit is set. VLAN tag is read from the DFVLAN register if PTP=1. 0 Ignore VLCTL field. 1 If VLAN tag insertion is enabled for eTSEC, use the VLCTL field as the VLAN control word. |
| | 1 | IP | Layer 3 header is an IP header. 0 Ignore layer 3 and higher headers. 1 Assume that the layer 3 header is an IPv4 or IPv6 header, and take L3OS field as valid. |
| | 2 | IP6 | IP header is IP version 6. Valid only if IP = 1. 0 IP header version is 4. 1 IP header version is 6. |
| | 3 | TUP | Layer 4 header is a TCP or UDP header. 0 Do not process any layer 4 header. 1 Assume that the layer 4 header is either TCP or UDP (see UDP bit), and offload checksumming on the basis that the IP header has no extension headers. |
| | 4 | UDP | UDP protocol at layer 4. 0 Layer 4 protocol is either TCP (if TUP = 1) or undefined. 1 Layer 4 protocol is UDP if TUP = 1. |

Table 19-57. Tx Frame Control Block Description (continued)

| Bytes | Bits | Name | Description |
|-------|------|------------------|---|
| 0–1 | 5 | CIP | Checksum IP header enable 0 Do not generate an IP header checksum. 1 Generate an IPv4 header checksum. |
| | 6 | CTU | Checksum TCP or UDP header enable. 0 Do not generate a TCP or UDP header checksum. RFC 768 advises that UDP packets not requiring checksum validation should have their checksum field set to zero. 1 Generate a TCP header checksum if IP = 1 and TUP = 1 and UDP = 0. |
| | 7 | NPH | Disable calculation of TCP or UDP pseudo-header checksum. This bit should be set if IP options need to be consulted in forming the pseudo-header checksum, as eTSEC does not examine IP options or extension headers for TCP/IP offload on transmit. 0 Calculate TCP or UDP pseudo-header checksum as normal, assuming that the IP header has no options. 1 Do not calculate a TCP or UDP pseudo-header checksum, but instead use the value in field PHCS when determining the overall TCP or UDP checksum. |
| | 8–14 | — | Reserved |
| | 15 | PTP | Indication to the transmitter that this is a PTP packet. Enabling PTP disables per packet VLAN tag insertion. Instead, VLAN tag will be read from the DFVLAN when the PTP field is true. 0 Do not attempt to capture transmission event time 1 Valid PTP_ID field. When this packet is transmitted, capture the time of transmission. Must be clear if TMR_CTRL[TE] is clear. |
| 2–3 | 0–7 | L4OS | Layer 4 header offset from start of layer 3 header. The layer 4 header starts L4OS octets after the layer 3 header if it is present. The maximum layer 3 header length supported is thus 255 bytes, which may prevent TCP/IP offload on particularly large IPv6 headers. |
| | 8–15 | L3OS | Layer 3 header offset from start of frame not including the 8 bytes for this FCB. The layer 3 header starts L3OS octets from the start of the frame including any custom preamble header that may be present. The maximum layer 2 header length supported is thus 255 bytes. |
| 4–5 | 0–15 | PHCS | Pseudo-header checksum (16-bit one's complement sum with carry wraparound, but without result inversion) for TCP or UDP packets, calculated by software. Valid only if NPH = 1. |
| 6–7 | 0–15 | VLCTL/ PTP_ID | VLAN control word for insertion in the transmitted VLAN tag. Valid only if VLN = 1. Tx PTP packet identification number. This number will be copied into the Tx PTP packet time stamp identification field. PTP field takes precedence over VLN field. |

19.6.7.3, 19-184

In Table 19-157, “Receive Buffer Descriptor Field Descriptions,” update description of Data Length field (bits 0–15) by replacing the following sentence:

“Data length is the number of octets written by the eTSEC into this BD’s data buffer if L is cleared (the value is equal to MRBLR), or, if L is set, the length of the frame including CRC, FCB (if RCTRL[PRSDEP > 00]) and any padding (RCTRL[PAL])”

with the following:

“Data length is the number of octets written by the eTSEC into this BD’s data buffer if L is cleared (the value is equal to MRBLR), or, if L is set, the length of the frame including CRC, FCB (if RCTRL[PRSDEP > 00]), preamble (if MACCFG2[PreAmRxEn]=1), timestamp (if RCTRL[TS]=1) and any padding (RCTRL[PAL]).”

- 19.7.1.3, 19-193 In Table 19-166, “RMII Mode Register Initialization Steps,” update the third row as follows:

| |
|--|
| Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0001_0000_0101] (I/F Mode = 1, Full Duplex = 1) |
|--|

- 19.7.1.5, 19-199 Add the following note before Table 19-170, “SGMII Mode Register Initialization Steps”:

SGMII mode utilizes the internal TBI PHY. The internal TBI PHY only auto-negotiates at 1 Gbps. However, 10 Mbps and 100 Mbps speeds are supported in SGMII mode. It is recommended that the external PHY inform the MAC if the desired link speed is not 1 Gbps. Software can perform MII management cycles to determine the external PHY link speed and program ECNTRL and MACCFG2 accordingly.

- 19.7.1.5, 19-199 Add Table 19-170, “SGMII Interface Signal Configuration (4-Wire),” to the beginning of the section:

Table 19-170. SGMII Interface Signal Configuration (4-Wire)

| SerDes Signals | | | SGMII Interface | | |
|------------------------|-----|----------------|----------------------|-----|----------------|
| Frequency [MHz] 1250 | | | Frequency [MHz] 1250 | | |
| Voltage [V] LVDS | | | Voltage [V] LVDS | | |
| Signals | I/O | No. of Signals | Signals | I/O | No. of Signals |
| TX_n/\overline{TX}_n | O | 2 | TXD | O | 2 |
| RX_n/\overline{RX}_n | I | 2 | RXD | I | 2 |
| Sum | | 4 | Sum | | 4 |

- 19.7.1.5, 19-200 In Table 19-170, “SGMII Mode Register Initialization Steps,” remove all references to TBICON[Enable Wrap] and TBICON[Comma Detect].
- 19.7.1.5, 19-201 In Table 19-171, “SGMII Mode Register Initialization Steps,” modify the row describing an MII Mgmt write to TBI in order to enable the TBI to restart Auto-Negotiations using the configuration set in the AN Advertisement register. The value for MIIMCON now reads, “0000_0000_0000_0000_0001_0011_0100_0000.”
- 20.3.1.2, 20-5 In second paragraph, change “This ratio is set in SCCR[TSEC2CM],” to “This ratio is set in SCCR[ENCCM].”
 In addition, remove the following:
 “Clock ratios of I2C1 are controllable but clock ratio for I2C2 is not and it is always 1:1 with CSB. Consider this factor when selecting an FDR value.”
- 21.2.2, 21-4 In Table 21-2, “DUART Signals—Detailed Signal Descriptions,” change sentence in $\overline{UART_RTS}[1:2]$ description from: “Can be programmed to be

- 21.3.1.3, 21-8 automatically negated and asserted by either the receiver or transmitter” to: “Can be programmed to be negated and asserted by either the receiver or transmitter.”
 Correct the formula for calculating the percent error value given in step 1 from:

$$AFI = \text{baud rate} \times 16$$

 to read as:

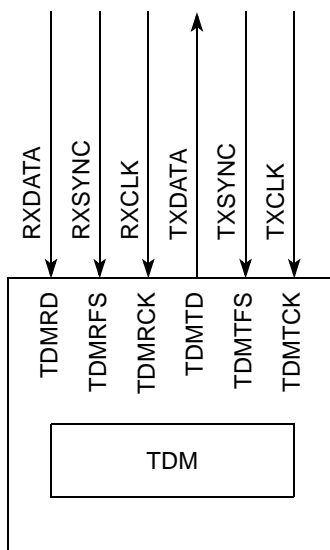
$$AFI = \text{baud rate} \times 16 \times \text{divisor}$$
- 22.2.3.3, 22-6 Change the following in the second paragraph from:
 “The maximum sustained data rate that the SPI supports is input clock/50. However, the SPI can transfer a single character at much higher rates—input clock/4 in master mode and input clock/2 in slave mode.”
 to:
 “The SPI can transfer a single character at much higher rates—input clock/4 in master mode and input clock/2 in slave mode, and subjected to the timing parameters of the interconnected devices, and board trace delays.”
- 22.4.1.2, 22-12 In Table 22-5, “SPIE Field Descriptions,” correct description of bit 17 (LT) from:
 “Last character was transmitted.
 The last character is transmitted and new data can be written to SPID for further transmission.”
 to:
 “Last character was transmitted.
 The last character of the frame was completely transferred. This bit is set only if the transmitted character was the last character of the frame (if SPCOM[LST] is set). New data can be written to SPITD is indicated by bit NF.”
- 25.2.1, 25-2 Modify Table 25-1, “TDM Signal Properties” as follows:

Table 25-1. TDM Signal Properties

| Name | Function | Type | Reset State |
|---------|-------------------------|------|-------------|
| TDM_TCK | TDM transmit clock | I/O | Input |
| TDM_TFS | TDM transmit frame sync | I/O | Input |
| TDM_RCK | TDM receive clock | I/O | Input |
| TDM_RFS | TDM receive frame sync | I/O | Input |
| TDM_TD | TDM transmit data | O | High Z |
| TDM_RD | TDM receive data | I | Input |

25.3.2, 25-6

Remove “CTS” from Figure 25-4 “TDM Modules Shared Mode,” to show as follows:



25.4.2.1.1, 25-9

In Figure 25-5, “TDM General Interface Register,” and in Table 25-5, “TDMGIR Bit Descriptions,” make bit 30 Reserved.

25.4.2.1.1, 25-9

Remove CTS column and last two rows of Table 25-6 to read as follows. In addition, updated the table title as “Pins Utilized for TDM Modules Based on RTS”

| RTS | TDM | | | | Notes |
|-----|--------|--------|--------|--------|---|
| | TCK | RCK | TFS | RFS | |
| 0 | TDMTCK | TDMRCK | TDMTFS | TDMRFS | TDM is independent from other TDMs and Rx and Tx are independent. |
| 1 | TDMTCK | TDMRCK | TDMTFS | TDMTFS | TDM is independent from other TDMs and Rx and Tx are shared. |

25.4.2.1.2, 25-11

Update bit description of RFSE (bit 30) in Table 25-7 “TDMRIR Bit Descriptions” from:

“Determines whether the receive frame sync signal is sampled with the rising or falling edge of the receive clock.”

to:

“Determines whether the receive frame sync signal is driven with the rising or falling edge of the receive clock.”

- 25.4.2.1.3, 25-14 Correct bit description of TFSE (bit 30) in Table 25-9, “TDMTIR Bit Descriptions” from:
 “Determines whether the transmit frame sync signal is sampled with the rising or falling edge of the transmit clock.”
 to:
 “Determines whether the transmit frame sync signal is driven with the rising or falling edge of the transmit clock.”
- 25.4.2.2, 25-17 Update TDM Receive Channel Enable n (TDMRCEN n) register section by combining individual register figures (Figure 25-10, Figure 25-11, Figure 25-12, and Figure 25-13) into one master figure as follows:

25.4.2.2.1 TDM Receive Channel Enable n (TDMRCEN n)

Figure 25-10 shows the TDM receive channel enable registers, 0–3.




Figure 25-10. TDM Receive Channel Enable n

Table 25-13 defines the bit fields of TDMRCEN n .

Table 25-13. TDMRCEN n Bit Descriptions

| Bits | Name | Description |
|------|----------|---|
| 0–31 | RCEN n | Receive channel active enable group n [0–31]. Set when the receive channel m is active. Each bit corresponds to each channel, as follows: <ul style="list-style-type: none"> • $n = 0$, bit 0 is the active bit for channel 0 and bit 1 is the active bit for channel 1. • $n = 1$, bit 0 is the active bit for channel 32 and bit 1 is the active bit for channel 33. • $n = 2$, bit 0 is the active bit for channel 64 and bit 1 is the active bit for channel 65. • $n = 3$, bit 0 is the active bit for channel 96 and bit 1 is the active bit for channel 97. 0 The channel m is not active. 1 The channel m is active. |

- 25.4.2.2, 25-18 Update TDM Transmit Channel Enable n (TDMTCEN n) register section by combining individual register figures (Figure 25-14, Figure 25-15, Figure 25-16, and Figure 25-17) into one master figure as follows:

25.4.2.2.2 TDM Transmit Channel Enable n (TDMTCEN n)

To drive and account for time slots, there is a double layer mask for transmission of data. Two set of registers are defined for each channel—TDMTCEN n and TDMTCMA n (Section 25.4.2.2.3, “TDM n Transmit Channel Mask n (TDM n TCMA n)”).

If all bits of the TDMTCMA n register are cleared, the TDM transmitter will continue to operate based solely on the TDMTCEN n registers. The TDMTCMA n registers are used to output data onto specific

channels. Therefore, TDMTCMA_n register bits should only be set on the channels that have their corresponding TDMTCEN_n bit already cleared; otherwise the channel is disabled.

Once the TDMTCEN_n register is cleared, the TDMTCMA_n register bit is set to 1 in the desired channel bit location for data to be discarded by the TDM. This allows for memory structures to keep channels open that are physically closed on the TDM, which are transferred to the TDM via a DMA, where the TDM discards the data.

If the TDMTCMA_n bit is set on a channel that is enabled, then the data for that channel is discarded by the TDM and the TDM will instead transmit all 1's. Disabling a channel in this manner causes the time slot for that channel to be ignored by the TDM. This means no data is transferred to the transmit shift register.

Figure 25-11 shows the TDM transmit channel enable registers, 0–3.




Figure 25-11. TDM Transmit Channel Enable n

Table 25-14 defines the bit fields of TDMTCEN_n .

Table 25-14. TDMTCEN_n Bit Descriptions

| Bits | Name | Description |
|------|-----------------|---|
| 0–31 | TCEN_n | Transmit channel active enable group n [0–31]. Set when the transmit channel m is active. Each bit corresponds to each channel, as follows: <ul style="list-style-type: none"> • $n = 0$, bit 0 is the active bit for channel 0 and bit 1 is the active bit for channel 1. • $n = 1$, bit 0 is the active bit for channel 32 and bit 1 is the active bit for channel 33. • $n = 2$, bit 0 is the active bit for channel 64 and bit 1 is the active bit for channel 65. • $n = 3$, bit 0 is the active bit for channel 96 and bit 1 is the active bit for channel 97. 0 The channel m is not active. 1 The channel m is active. |

25.4.2.2, 25-19 Update TDM Transmit Channel Mask n (TDMTCMA_n) register section by combining individual register figures (Figures 25-18, Figure 25-19, and Figure 25-20) into one master figure as follows:

25.4.2.2.3 TDM Transmit Channel Mask n (TDMTCMA_n)

The transmit channel mask registers 0–3 are shown in Figure 25-12. See Section 25.4.2.2.2, “[TDM \$n\$ Transmit Channel Enable \$n\$ \(\$\text{TDM}_n\text{TCEN}_n\$ \)](#)” for a description on how these registers are used to output data onto specific channels.

Revision History

Offset 0x060, 0x064, 0x068, 0x06C

Access: Read/Write




Figure 25-13. TDM Transmit Channel Mask n

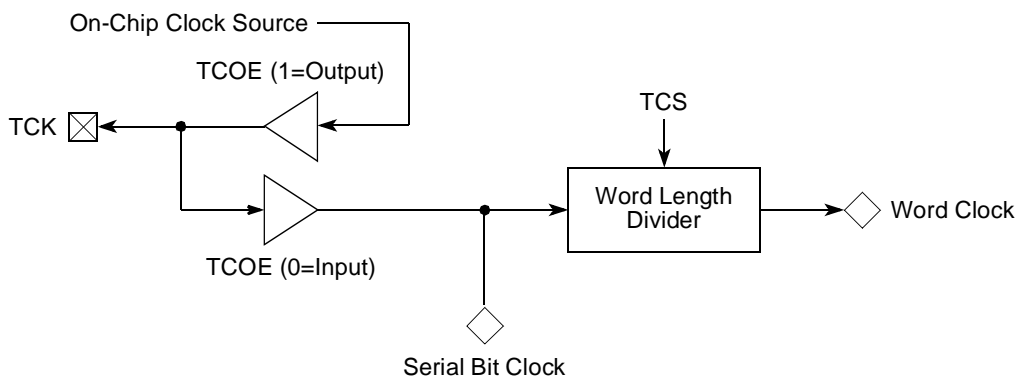
Table 25-15 defines the bit fields of TDMTCA_n.

Table 25-15. TDMTCA_n Bit Descriptions

| Bits | Name | Description |
|------|-------------------|---|
| 0-31 | TCMA _n | <p>Transmit channel mask group 0 [0-31]. Set when the transmit channel m data is to be ignored when received in the transmit data register (TDREG). Each bit corresponds to each channel, as follows:</p> <ul style="list-style-type: none"> $n = 0$, bit 0 is the active bit for channel 0 and bit 1 is the active bit for channel 1. $n = 1$, bit 0 is the active bit for channel 32 and bit 1 is the active bit for channel 33. $n = 2$, bit 0 is the active bit for channel 64 and bit 1 is the active bit for channel 65. $n = 3$, bit 0 is the active bit for channel 96 and bit 1 is the active bit for channel 97. <p>0 The channel m is transmitted according to the corresponding TDMTCE_{N0-3} register. 1 The channel m data is discarded.</p> |

25.5.1, 25-33

In Figure 25-34, “TDM Transmit Clock Generator Block Diagram,” correct “TCOE (1=Output)” from “TCOE (0=Output)” as follows:



25.7.1.1, 25-36

Correct mention of TSOE to TSA and RSOE to RSA. Changed the sentence from: Additionally, the edge the TFS and RFS are driven out is controlled by writing to the TSOE bit in the transmit interface register (TDMTIR) or writing to the RSOE bit in the receive interface register (TDMRIR).

to:

Additionally, the edge the TFS and RFS are driven out is controlled by writing to the TSA bit in the transmit interface register (TDMTIR) or writing to the RSA bit in the receive interface register (TDMRIR)

25.7.2, 25-42

Remove the second paragraph.

25.10.3.14, 25-74 Modify Figure 25-73, “DMA Hardware Request Status Register—High,” and Table 25-50, “DMAGPOR Field Descriptions,” as follows:




Figure 25-73. DMA Hardware Request Status Register—High

Table 25-50. DMAGPOR Field Descriptions

| Bits | Name | Description |
|-------|--------------------|--|
| 31–16 | — | Reserved |
| 15–14 | priority_ctrl_high | Priority control high. Indicates the bump up priority value in case of non-sequential transactions. |
| 13–12 | priority_ctrl | Priority control. Normal priority value. |
| 11–7 | — | Reserved |
| 6 | snoop enable | Snoop enable. 0 Transaction is non-cacheable 1 Transaction is cacheable |
| 5–3 | — | Reserved |
| 2 | rd_safe_enable | Read safe enable. 0 Reading extra bytes from the memory may have side effects. 1 Only indicated amount of data beats should be read. |
| 1–0 | — | Reserved |

B.2 Changes From Revision 0 to Revision 1

Major changes to the *MPC8315E PowerQUICC II Pro Integrated Host Processor Family Reference Manual*, from Revision 0 to Revision 1 are as follows:

Section, Page

Changes

1.2.2, 1-13 Modify Figure 1-3, “Integrated Security Engine Functional Blocks,” as follows:

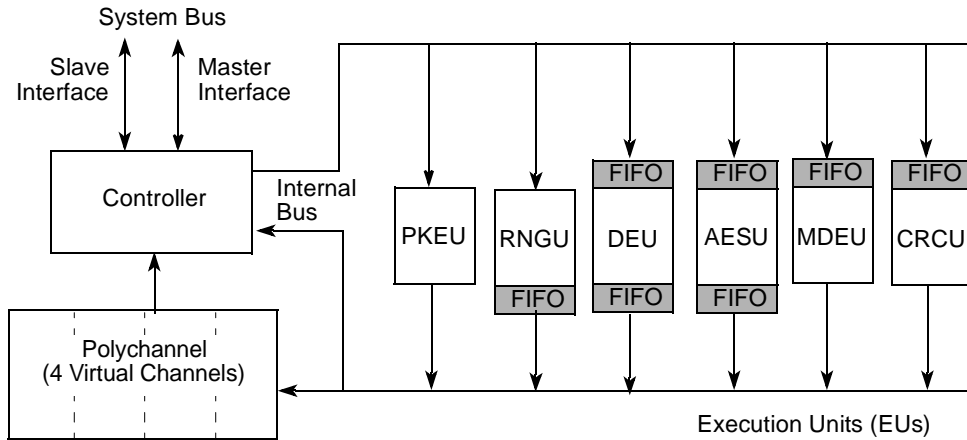


Figure 1-3. Integrated Security Engine Function Blocks

- 2.1, 2-1 In Figure 2-1, “MPC8315E Signal Groupings (1 of 2),” add “LB_POR_CFG_BOOT_ECC” as a muxed signal to “TSEC_MDC.”
- 2.1, 2-15 In Table 2-1, “MPC8315E Signal Reference by Functional Block,” change signal name “RTC_PIT_CLOCK” to say “RTC_CLK.”
- 2.1, 2-17 In Table 2-1, “MPC8315E Signal Reference by Functional Block,” add footnote “In silicon rev 1.0, NAND boot ECC checking is enabled. In silicon rev 1.1, NAND boot ECC checking is configurable; by default it is enabled (pulled down).” to LB_POR_CFG_BOOT_ECC.
- 2.2, 2-17 In Table 2-6, “Output Signal States During System Reset,” add bit ranges and correct “State During Reset” values.
- 4.2.2, 4-6 Remove “There is no need to assert the $\overline{\text{SRESET}}$ signal when $\overline{\text{HRESET}}$ is asserted” from the list of steps.
- 4.3.3.1.1, 4-23 Remove Figure “Loading Reset Configuration Words from Local Bus.”
- 4.4.4, 4-33 Change “SYS_CLK_IN must be 24, 32, or 48 MHz” to say “SYS_CLK_IN must be 24 or 48 MHz.”
- 5.3.2.7.1, 5-29 Change “either” to say “enter,” as follows:
 “The DDR debug configuration enables a DDR memory controller to enter debug mode in which the DDR SDRAM source ID field and data valid strobe...”
- 5.8, 5-75 Add bulleted sentence “A low-power mode (D3Warm) can be achieved using a split supply, which will make...” to bottom of bulleted list.

- 5.8, 5-75 Add beneath bulleted list, “The MPC8315E will support the power states D0, D1, D2, D3Hot, D3Warm, D3Cold. D3Warm is not supported in PCI agent mode. See [Section 5.9.3, “Functional Description,”](#) for details about each state.”
- 5.8.1, 5-76 In Table 5-73, “System Control Signals—Detailed Signal Descriptions,” add signals EXT_PWR_CTRL and PMC_PWR_OK.
- 5.8.2.2, 5-77 In Table 5-58, “Power Management Controller Event Register,” add footnote “Not used for wake-up from D3Warm” to bits 24–26.
- 5.8.2.4, 5-81 In Table 5-78, “PMCCR1 Bit Settings,” add “For D3Warm state, this bit should be set to 0” to bit USE_STATE description.
- 5.8.3, 5-84 Under “There are six power states in the MPC8315E,” change “Wake-on-LAN, GPIO (limited number), external interrupt, and timer only is supported, reset required” to say “Wake-on-LAN, GPIO (limited number), external interrupt, and timer only is supported.”
- 5.8.3, 5-84 Under “These features are not support,” remove the following:
“Power management features of PCI Express and SATA—on in all states except D3Warm and D3Cold.”
- 5.8.3, 5-84 In bulleted list, change “Reception of Magic Packet on eTSEC1” to say “Reception of Magic Packet on eTSEC2.”
- 5.8.3, 5-84 Remove references to 120 W.
- 5.8.3, 5-84 Remove from bulleted list “By programming the I/O register in csb2men, function of eTSEC I/O pads can be programmed.”

5.8.3.4, 5-86 Modify Figure 5-62, “Power Segmentation in Deep Sleep Mode,” as follows:

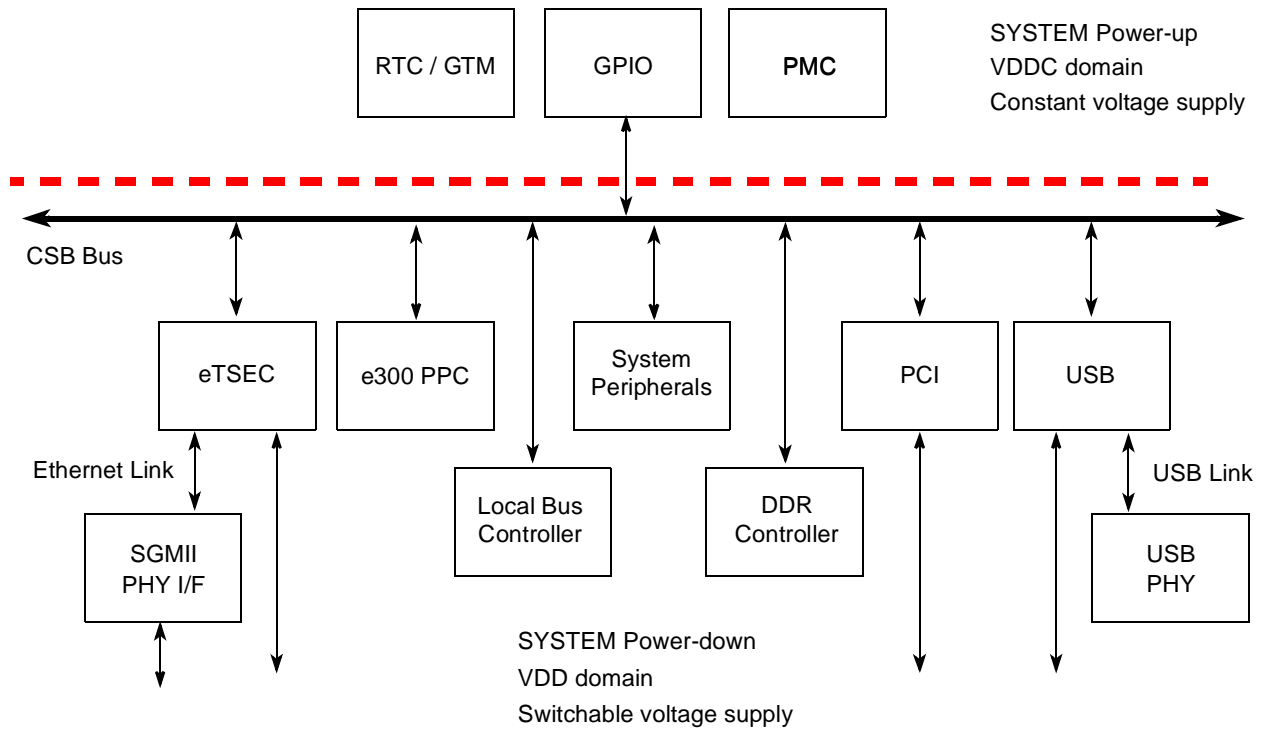


Figure 5-62. Power Segmentation in Deep Sleep Mode

- 5.8.3.5.2, 5-89 In third paragraph, change “The DLL is then shut off and stops driving clocks on MCK n pins” to say “It stops driving clocks on MCK n pins.”
- 5.8.3.7, 5-91 Remove “This applies to D3Warm as well” from second bullet.
- 5.8.3.7.1, 5-92 In Table 5-83, “MPC8315E Agent Mode Wake-Up Support,” remove support for D3Warm for agent mode.
- 5.8.3.7.1, 5-92 In Table 5-83, “MPC8315E Host Mode Wake-Up Support,” remove USB and PCI (PCI_PME) as sources from D3Warm.
- 5.8.3.7.3, 5-98 Remove Figure 5-63, “Example VDD Control of Device as Agent Using the Optional PMC_PMR_OK Signal.”

- 5.8.3.7.3, 5-98 Modify Figure 5-98, “Example VDD Control of Device as Host, Using Optional PMC_PWR_OK Signals,” as follows:

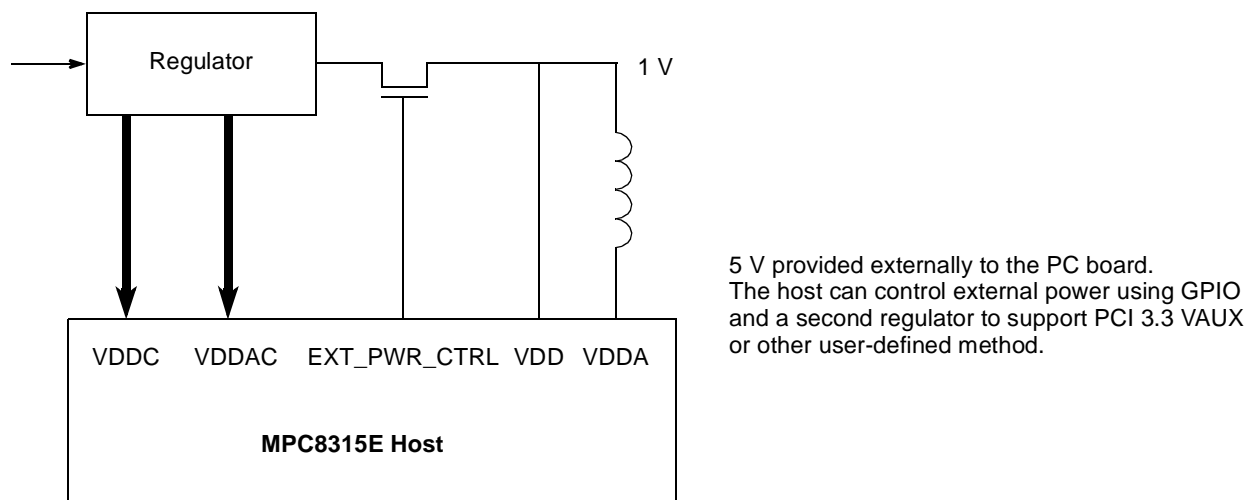


Figure 5-98. Example VDD Control of Device as Host, Using Optional PMC_PWR_OK Signals

- 5.8.3.7.2, 5-99 Add subsection “PMC Power-Down Sequence.”
- 5.8.3.7.2, 5-99 Remove subsections “PMC Power-Down Sequence When the MPC8315E is a PCI Agent” and “PMC Wake-Up When the MPC8315E is a PCI Agent.”
- 5.8.3.7.2, 5-99 Remove references to USB and PCI agent in subsection “PMC Wake-Up When the MPC8315E is a PCI Host.”
- 5.8.3.7.2, 5-101 In Section “PMC Wake-Up When the MPC8315E is a PCI Host,” add Figure 5-66, “PMC Wake-Up Sequence from D3Warm.”
- 5.8.3.7.2, 5-102 In Section “PMC Wake-Up When the MPC8315E is a PCI Host,” add Figure 5-67, “PMC Wake-Up Sequence from D3Warm (continued).”
- 5.8.3.7.4, 5-105 In subsection “Debugger in Low-Power Mode,” change “In D3Warm, the JTAG debug logic is powered down such that JTAG debug accesses will not be possible” to say “In D3Warm, the JTAG debug logic is powered up but JTAG debug accesses will not be possible.”
- 7.2, 7-13 Add Section 7.2, “e300 Processor and System Version Numbers,” and Table 7-1, “Device Revision Level Cross-Reference.”
- 9.1.4.1.7, 9-18 In Figure 9-8, “DDR SDRAM Control Configuration Register (DDR_SDRAM_CFG),” change bits 11–12 from reserved, “—,” to show “DBW.”
- 9.1.4.1.7, 9-19 In Table 9-12, “DDR_SDRAM_CFG Field Descriptions” change bits 11–12 from “Reserved” to show “DBW.”
- 9.6.3.3, 9-55 Modify first paragraph to say, “The DDR controller offers an initialization bypass feature (DDR_SDRAM_CFG[BI]), which system designers may use to prevent re-initialization of main memory during system power-on following an abnormal

| | |
|-------------------|--|
| | shutdown. See Section 9.4.1.7, ‘DDR SDRAM Control Configuration (DDR_SDRAM_CFG),’ for information on this bit.” |
| Chapter 10, 10-1 | Add note “The MPC8315E eLBC does not support atomic features. References to this feature should be disregarded.” |
| 10.3.1.2, 10-12 | Move Table 10-5, “Reset Value of OR0 Register,” to the containing section on options registers from the specific description of OR0 in GPCM mode (Section 19.3.1.2.2, “Option Registers (OR n)—GPCM Mode”), because the table applies to OR0 as reset in all machines. |
| 10.4.3.4.2, 10-72 | Change steps 3 and 4 to say the following: <ol style="list-style-type: none"> 3. FCM reads the spare regions of the first two pages of the current block, checking the bad block indication (BI) bytes to validate the block for reading. BI bytes must all hold the value 0xFF for the page to be considered readable. <ul style="list-style-type: none"> — For small-page devices, BI is a single byte read from spare region byte offset 5. — For large-page devices, BI is a single byte read from spare region byte offset 0. <p>If either of the first two pages of the current block are marked invalid, then the boot block index is incremented by 1, and FCM repeats step 3. eLBC will continue searching for a bootable block indefinitely, therefore at least one block must be marked valid for boot loading to proceed. At the conclusion of the boot block search, the value of FBAR[BLK] points to the boot block.</p> 4. If ECC checking is enabled, the FCM recovers from the spare region the stored ECC for each 512-byte block of boot data. The boot block must be prepared with ECC protection. During ECC generation, software should use FMR[ECCM] = 0 for small-page devices, and FMR[ECCM] = 1 for large-page devices. |
| 10.5.1.3, 10-90 | Change example frequency to “133 MHz.” |
| 10.5.4.5, 10-96 | Change first sentence of second paragraph to say “Note that operations specified by OP3 and OP4 (status read) should never be skipped while erasing a NAND Flash device, because, in case that happens, contention may arise on LGPL4.” |
| 10.5.4.6, 10-96 | Change first sentence of second paragraph to say “Note that operations specified by OP5 and OP6 (status read) should never be skipped while programming a NAND Flash device, because, in case that happens, contention may arise on LGPL4.” |
| 12.4.1, 12-5 | In Figure 12-2, “Outbound Message Interrupt Status Register (OMISR),” change bits 0 and 1 to say “w1c.” |
| 12.4.1, 12-5 | In Figure 12-2, “Outbound Message Interrupt Status Register (OMISR),” change access from “User read/write” to say “Mixed.” |
| 12.4.2, 12-6 | Change second sentence of description of OMISR register to say “OMIMR can be read from the CSB or the PCI bus, but it can be written only from the PCI bus.” |
| 14.4.4.7, 14-40 | In Figure 14-49, “PCI Express Next Capabilities Pointer,” add footnote “The reset value of 0b0111_0000 is only true in EP mode. In RC mode, the reset values should be 0b0000_0000.” |
| 14.4.4.7, 14-41 | In Table 14-47, “PCI Express Next Capabilities Pointer Field Descriptions,” add “The reset value is 0x70 in EP mode and 0x00 in RC mode” to description. |

- 14.4.5, 14-51 Update bit field descriptions.
- 14.4.5.1, 14-52 In Figure 14-68, “PCI Express Advanced Error Reporting Capability ID Register,” add bits 31–16.
- 14.4.5.1, 14-52 In Table 14-65, “PCI Express Advanced Error Reporting Capability ID Register Field Descriptions,” add bits 31–16.
- 14.4.5.2, 14-52 In Figure 14-69, “PCI Express Uncorrectable Error Status Register,” change access from “Read/Write” to “w1c.”
- 14.4.5.2, 14-52 In Table 14-69, “PCI Express Uncorrectable Error Status Register Field Descriptions,” update bit field descriptions.
- 14.4.6.2, 14-62 Add description of PEX_GCLK_RATIO register.
- 14.6.1.8, 14-105 In Table 14-131, “Configuration Address Mapping,” change content, as follows:

Table 14-131. Configuration Address Mapping

| CSB Address Bits | PCI Express Address Bits | PCI Express Configuration Space |
|------------------|--------------------------|---------------------------------|
| 0:7 | 31:24 | Bus number |
| 8:12 | 23:19 | Device number |
| 13:15 | 18:16 | Function number |
| 16:19 | 15:12 | Reserved |
| 20:23 | 11:8 | Extended register number |
| 24:29 | 7:2 | Register number |
| 30:31 | 1:0 | Reserved |

- 14.6.5, 14-111 Change first sentence, “When a hot reset condition occurs, the controller (in both RC and EP mode) initiates a cleanup of all outstanding transactions and returns to an idle state” with the following:
 “When a hot reset condition occurs, the controller (in both RC and EP mode) initiates a cleanup of all outstanding transactions and goes into suspend mode. The RC controller then needs to be reset (issuing CBRST in PECR1/PECR2) to bring it back to the idle state followed by a reprogramming of all the CSB bridge registers (offset 0x800–0xFFC, such as the ATMU windows). The EP controller also needs to be reset, followed by a reprogramming of the entire configuration space and CSB bridge registers (offset 0x800–0xFFC, such as the ATMU windows).”
- 14.8.5, 14-195 Remove sentence “Descriptor-based DMA supports debug mode in which the host can directly program the descriptors in the DMA register instead of storing the program in memory.”
- 15.1, 15-1 Add “eSATA” to bulleted list of features.
- 15.3.4.8, 15-26 In Figure 15-24, “PHY Control Configuration Register1 (PhyCtrlCfg1),” show bit field REF_CLK_SEL as bits 6–4.
- 15.7, 15-33 Remove subsection “SerDes Initialization.”

Revision History

| | |
|-----------------------|--|
| 15.3.9, 15-36 | In numbered list, change “Modifies the protocol converter control register to place the PHY into loopback mode” to say “Modifies the SerDes control register 1 to place the PHY into analog loopback mode.” |
| Chapter 16, 16-1 | Add note, “On the MPC8315E there are two SerDes lanes, A and B. This chapter describes a SerDes with lanes A and E. Whenever lane E is mentioned in this chapter, the MPC8315E lane B is intended.” |
| 16.2, 16-3 | In Table 16-1, “SerDes External Signals—Detailed Signal Descriptions,” change “200 W” and “100 W” to “200 Ω” and “100 Ω” in descriptions for “SD_IMP_CAL_RX/” and “SD_IMP_CAL_TX/.” |
| 16.3.1, 16-6 | In Table 16-8, “SRD _n CR0 Field Descriptions,” for bits 16 and 20, change bit setting for “1” to “Reserved.” |
| 17.3, 17-6 | In Table 17-3, “USB Interface Memory Map,” replace USBDR register reset values, as follows: Address offset= 0x23104: Value=0x1100_0100 Address offset= 0x23124: Value= 0x8301_0000 Address offset= 0x23140: Value= 0x0000_0800 Address offset= 0x23164: Value= 0x0000_0000 Address offset= 0x23184: Value= 0x0000_0010 Address offset= 0x231A4: Value= 0x200C_0000 |
| 17.5.6, 17-61 | In Figure 17-41, “Queue Head Layout,” show field “R”L as 4 bits (31–28). |
| 17.6.1, 17-67 | Remove Table 17-64, “Default Values of Operational Register Space.” |
| 17.6.1, 17-67 | Change the first paragraph to say, “After initial power-on or host controller reset (hardware or through USBCMD[RST]), all of the operational registers will be at their default values. After a hardware reset, only the operational registers will be at their default values.” |
| Chapter 19 | Add overbars to active-low signal “ \overline{LCS}_n .” |
| Chapter 19 | Remove statements that refer to supporting IEEE Std 1588™ on a per-eTSEC basis. |
| Chapter 19 | Remove section “Rx FCB Change,” which states that there are no changes to RxFCB due to IEEE 1588. |
| 19.4.1, 19-8 | In Table 19-2, “eTSEC Signals—Detailed Signal Descriptions,” for GTX_CLK125 description, add SGMII to list of protocols for which GTX_CLK125 is not used to for. |
| 19.4.1, 19-8 | In Table 19-2, “eTSEC Signals—Detailed Signal Descriptions,” in State Meaning description for TSEC _n _CRS, change “TSEC _n _TX_CLK” to say “TSEC _n _CRS.” |
| 19.5.3.1.3, 19-25, 27 | In Figure 19-4, “IEVENT Register Definition,” and Table 19-7, “IEVENT Field Descriptions,” add field IEVENT[FGPI] field, bit 27. |
| 19.5.3.3.1, 19-49 | In Table 19-25, “RCTRL Field Descriptions,” change RTCRL[RSF] field description to say: |

- Receive short frame mode. When set, enables the reception of frames shorter than 64 bytes.
- 0 Ethernet frames less than 64B in length are silently dropped.
- 1 Frames more than 16B and less than 64B in length are accepted upon a DA match.
- 19.5.3.3.7, 19-57 In Figure 19-27, “Receive Queue Filer Table Control Register Definition” and Table 19-31, “RQFCR Field Descriptions,” add field RQFCR[GPI], bit 0.
- 19.5.3.6.26, 19-91 In Table 19-78, “TPKT Field Descriptions,” make bits 0–9 reserved and bits 10–31 TPKT.
- 19.5.3.10.1, 19-111 In Table 19-107, “TMR_CTRL Register,” add “Inverted frequency tuned timer input clock (NOTE: this setting is reserved if CKSEL=01.)” to TMR_CTRL[CIPH] field description.
- Note that frames less than or equal to 16B in length are always silently dropped.
- 19.5.3.10.1, 19-111 In Table 19-107, “TMR_CTRL Register Field Descriptions,” updated bitfield definitions.
- 19.6.2.8, 19-152 Change last sentence of second paragraph to say “Only frames addressed specifically to the MAC’s station address or a valid multicast or broadcast address can be examined for the Magic Packet sequence.”
- 19.6.4.1.1, 19-165 Add following statement to bottom of bulleted list:
- “The GPI field offers the user the ability to interrupt the core upon matching a rule that causes a frame to be filed to memory. Once the last RxBD corresponding to that frame is written to memory, the IEVENT[FGPI] event will be asserted. This bit will be set regardless of any interrupt coalescing that may be set.”
- 19.6.4.1.4, 19-166 Add the following paragraphs to the end of the section:
- “A functional interrupt is provided via use of the general purpose interrupt (GPI) bit in the filer table. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the filer will set IEVENT[FGPI] when the corresponding receive frame is written to memory. This allows the user to set up a filer rule where the core will be interrupted upon the reception of ‘special’ frames.
- If the timer is enabled (TMR_CTRL[TE] = 1), then the interrupt dedicated for timer events (in addition to the usual receive, transmit and error interrupts) will be asserted.”
- 19.6.5.2.1, 19-173 Change last sentence of second paragraph to say:
- “As soon as the hardware consumes a BD (by writing it back to memory), RBPTRn will advance and the free BD count will reflect the correct number of available free BDs.”
- 19.6.6.4.1, 19-177 Change paragraph to say:
- “The eTSEC receive filer has been enhanced with the addition of a general-purpose event bit. This event bit can be used in conjunction with filing

Revision History

| | |
|------------------|--|
| | <p>table rules to identify 1588 packets and indicate these packets by setting special timer status register bits (TMR_STAT). Additionally, 1588 packets can be easily identified by upper-layer software by using the filer to queue all PTP packets to one or more predefined virtual queues. See Section 19.6.5.1.1, “Filing Rules,” for further information.”</p> |
| 19.6.7.3, 19-184 | <p>In Table 19-157, “Receive Buffer Descriptor Field Descriptions,” for offset 4–7 and bits 0–31, add sentence “For best performance, use 64-byte aligned receive buffer pointer addresses. The buffer must reside in memory external to the eTSEC.”</p> |
| 20.3.1.5, 20-9 | <p>In Table 20-8, “I2CDR Field Description,” modify last sentence to say, “Note that in both master receive and slave receive modes, the very first read is always a dummy read.”</p> |
| 20.5.5, 20-23 | <p>Remove sentence, “For 1-byte transfers, a dummy read should be performed by the interrupt service routine (see Figure 20-11).”</p> |
| 21.3.1.3, 21-24 | <p>Replace Table 21-7, “UDLB Field Description.”</p> |