Freescale Semiconductor
Addendum

Document Number: MPC8641DRMAD
Rev. 1.1, 07/2008

Errata to MPC8641D Integrated Host Processor Family Reference Manual, Rev. 1

This errata describes corrections to the *MPC8641D Integrated Host Processor Reference Manual*, revision 1. For convenience, the section number and page number of the errata item in the reference manual are provided.

To locate any published updates for this document, visit our website listed on the back cover of this document.

freescale semiconductor

Section, Page No.	Changes
1.3.1/1-9	Changed “Total latency...” bullet under “On-chip level 2 (L2) cache ...” to 11.5 and 12.5 cycles when ECC is enabled.
1.3.2/1-14	In list of features, changed DDR data rate to 600 MHz and 300 MHz.
1.5.1/1-26	Removed Figure 1-8, “Key Areas for Bandwidth Performance.”
2.2.2/2-5	Added Section 2.2.2, “Precedence of Local Access Windows,” as follows: “If two or more LAWs overlap, the lower numbered window takes precedence. For instance, consider two LAWs, set up as shown in Table 2-5.

Table 2-5. Overlapping Local Access Windows

LAW	Base Address	Size	Target Interface
1	0x0_7FF0_0000	1 Mbyte	0b00100 (local bus controller—LBC)
2	0x0_0000_0000	2 Gbytes	0b01111 (DDR controller 1)

In this case, LAW 1 governs the mapping of the 1-Mbyte region from 0x0_7FF0_0000 to 0x0_7FFF_FFF, even though the window described in LAW 2 also encompasses that memory region.

NOTE

The CCSR mapping, defined by CCSRBAR, supersedes all local access window mappings.”

2.4/2-9	Added a note, as follows: “The CCSR window always takes precedence over all local access windows. However, the CCSR window must not overlap an LAW that maps to the DDR controller. Otherwise, undefined behavior occurs.”
2.4.1/2-10	Added Section 2.4.1, “Accessing CCSR Memory from the Local Processor,” as follows: “When a local e600 core is used to configure CCSR space, the CCSR memory space should typically be marked as cache-inhibited and guarded. In addition, many configuration registers affect accesses to other memory regions or peripherals; therefore writes to these registers must be guaranteed to have taken effect before accesses are made to the associated memory regions or peripherals. To guarantee that the results of any sequence of writes to configuration registers are in effect, the final configuration register write should be chased 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.”
3.2/3-24	In Table 3-3, “MPC8641D Reset Configuration Signals,” changed default setting for cfg_sys_pll[0:3] and cfg_core_pll[0:3] to 1111, to cfg_core_pll[4]to 1.
4.3.1.3.1/4-7	In Figure 4-4, “Boot Page Translation Register (BPTR),” changed offset of BPTR from 0x010 to 0x020.
4.4.2/4-9	Revised first paragraph of note as follows:

NOTE

The common on-chip processor (COP) requires the ability to independently assert HRESET and TRST to fully control the processor. If a JTAG/COP port is used, follow the JTAG/COP interface connection recommendations given in the *MPC8641 and MPC8641D Integrated Host Processor Hardware Specifications*. If the JTAG interface and COP header are not being used, Freescale recommends that TRST be tied to HRESET so that TRST is asserted when the HRESET is asserted, ensuring that the JTAG scan chain is initialized during the power-on reset flow. See the JTAG configuration signals section in the hardware specifications document for more information.

4.4.2/4-9 Revised step 4, as follows:

"HRESET negates after a minimum of 100 µs. Note that POR configuration inputs (except those used during HRESET assertion) should be valid at least 4 SYSCLK cycles prior to HRESET negation and held valid at least 2 SYSCLK cycles after HRESET has been negated."

4.4.3.3/4-12 In Table 4-11, "e600 Core Clock PLL Ratios," changed default setting for cfg_core_pll[0:4] to 11111.

4.4.3.1/4-12 In Table 4-12, "Core 1 Enable," changed default setting for cfg_por1_enable to 1111.

4.4.3.7/4-14 Clarified sentence in first paragraph, as follows: "This POR configuration option only affects the outbound ATMU window 3 of PCI Express controller 1."

5.3.4.2.1/5-30 Added Section 5.3.4.2.1, "Sources of tea_ assertion," as follows:

When the core performs a read transaction to a target on one of the external interfaces, a bus error may occur, which in turn causes the assertion of *tea_* to the core. Table 5-3 summarizes the potential sources of *tea_* assertion and the associated error reporting register fields.

Table 5-3. *tea_* Sources

Read Target	Error Register	Error Register Fields
MCM	EDR	LAE
DDR	ERR_DETECT	MBE MSE
LBIU	LTESR	DM PAR
SRIO	LTLEDCSR	IER PRT

Table 5-3. *tea_* Sources (continued)

Read Target	Error Register	Error Register Fields
PCI Express	PEX_ERR_DR	PCT PCAC CDNSC CRSNC IACA CRST IOIS CIS CIEP IOIEP OAC IOIA
	PEX_PME_MES_DR	HRD LDD

6.3.6.3.2/6-86 Changed the last sentence of the fourth paragraph to say the following: “Note that the broadcast of **tlbie** and **tlbsync** instructions is enabled by the setting of HID1[ABE].”

7.4.1.3/7-6 In Table 7-4, “PCR Field Descriptions,” revised bit description for PORT1_EN to state it is not intended to dynamically enabled and disabled.

In addition, in PORT1_PRI description, "Highest priority level" should have the encoding 10.

7.4.1.4/7-7 In Table 7-5, “EDR Field Descriptions,” revised LAE field description to identify that **tlbie** broadcasts can generate LAEs.

7.4.1.6/7-9 In Table 7-7, “EATR Field Descriptions,” revised bit encodings for PCI Express 2 to now say, “PCI Express 2/RapidIO,” and changed the former RapidIO encoding to reserved.

7.4.1.6/7-9 In Table 7-7, “EATR Field Descriptions,” revised the SRC_ID field description to now use the following encodings:

“10000Core 0 (instruction/data)

10001 Reserved

10010 Core 1 (instruction/data)

10011–10100 Reserved”

8.3.1/8-3 In Table 8-1, “DDR Memory Interface Signal Summary,” added a note clarifying that some devices that implement two DDR controllers may share MDVAL and MSRCID[0:4] signals between both controllers; other devices may offer a set for each controller.

8.3.2.2/8-10 In Table 8-4, “Clock Signals—Detailed Signal Descriptions,” added the following to the MCKE description:

Section, Page No.	Changes
	"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]."
8.4.1.7/8-22	In Table 8-12, "DDR_SDRAM_CFG Field Descriptions," in 8_BE field description, modified note to say the following: "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."
8.4.1.7/8-22	In Table 8-12, "DDR_SDRAM_CFG Field Descriptions," added new programming requirement for DDR_SDRAM_CFG[HSE] such that this bit should not be set if using automatic calibration.
8.4.1.14/8-32	Changed introductory sentence for Figure 8-15 to say the following: "The DDR SDRAM clock control configuration register, shown in Figure 8-15, provides a 1/8-cycle clock adjustment."
8.4.1.28/8-41	In Figure 8-34, "Memory Data Path Read Capture ECC Register (CAPTURE_ECC)," extended bit field for ECE from 24:31 to 16:31. Added detailed bit field description after generic ECE statement, as follows: "Error capture ECC. Captures the ECC bits on the data path whenever errors are detected. 16:23—8-bit ECC code for 1st 32 bits 24:31—8-bit ECC code for 2nd 32 bits **Note:** In 64-bit mode, only 24:31 should be used, although 16:23 shows the 8-bit ECC code replicated."
8.4.1.32/8-44	In Table 8-38, "CAPTURE_ATTRIBUTES Field Descriptions," changed TSIZ field description to say the following: "000 4 double words 001 1 double word 010 2 double words 011 3 double words Others Reserved"
8.4.1.32/8-44	In Table 8-38, "CAPTURE_ATTRIBUTES Field Descriptions," revised bit encodings for TSRC field, as follows: "Transaction source for the error 00000 PCI Express interface 1 00001 PCI Express interface 2/RapidIO 00010–01001 Reserved 01010 Boot sequencer 01011–01111 Reserved

10000 Core 0 (instruction/data)
10001 Reserved
10010 Core 1 (instruction/data)
10011–10100 Reserved
10101 DMA
10110–10111 Reserved
11000 eTSEC 1
11001 eTSEC 2
11010 eTSEC 3
11011 eTSEC 4
11100 RapidIO message unit
11101 RapidIO doorbell unit
11110 RapidIO port-write unit
11111 Reserved”

8.5.6/8-65 Added the following note after the first paragraph:

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.

8.5.11/8-71 Added the following text to the end of the section:

“In 32-bit mode, Table 8-54 is split into 2 halves. The first half, consisting of rows 0–31, is used to calculate the ECC bits for the first 32 data bits of any 64-bit granule of data. This always applies to the odd data beats on the DDR data bus. The second half of the table, consisting of rows 32–63, is used to calculate the ECC bits for the second 32 bits of any 64-bit granule of data. This always applies to the even data beats on the DDR data bus.”

8.6.1/8-75 In Table 8-58, Programming Differences Between Memory Types,” for ODT_PD_EXIT, changed it to be set to 0001 for DDR1; for FOUR_ACT, changed it to be set for 00001 for DDR1.

Chapter 9/9-1 Changed all instances of “Message Shared” to “Message Signaled” throughout.

In addition, revised description for Activity bit to reference the respective xxVPRs throughout.

9.1.1/9-2 In Figure 9-1, “Interrupt Sources Block Diagram Features,” changed reference in Internal Interrupts block to Table 9-2, “Internal Interrupt Assignments.”

Section, Page No.	Changes
9.2.2/9-8	In Table 9-3, “Interrupt Signals—Detailed Signal Descriptions,” removed the following note from IRQ*n*: “If IRQ[0:3] or IRQ[4:7] are used to receive INTx signals from PCI Express port 1 and port 2, respectively, as a root complex, the polarity and sense of each of these signals must be set to be active-high and level-sensitive.”
9.3.1/9-18	Added clarification to the proper used of bits in the PIR register. Namely, core reset request should not be cleared until the requested core reset has occurred.
9.3.1.6/9-21	Added sentence in paragraph before Figure 9-8: “However, if one core is used to reset another one, the core being reset can effectively be held off indefinitely from issuing its initial boot vector fetch to the platform by leaving its appropriate PIR[Px] bit asserted.”
9.3.6.3/9-37	In Figure 9-33, “Shared Message Signaled Interrupt Index Register (MSIIR),” changed location of SRS and IBS fields in MSIIR register, as follows. Note that for the MPC8641D Rev 2.0 and earlier, the SRS and IBS fields occupy bits 24–31:

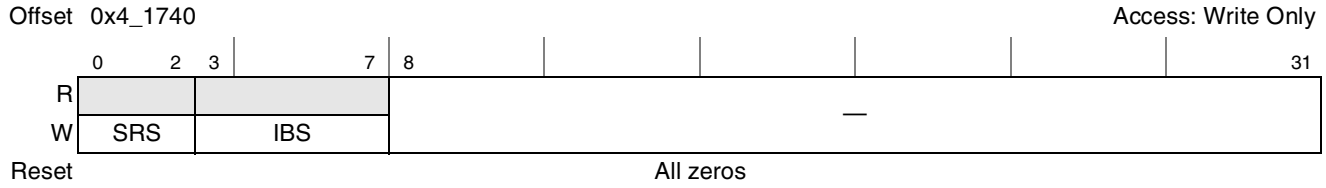

Figure 9-34. Shared Message Signaled Interrupt Index Register (MSIIR)

9.3.7.1/9-41	Removed last sentence regarding polarity and sense for IRQ[0:7] and PEX INTx sharing. This is fully covered in Section 9.4.5, “PCI Express INTx.”
9.3.7.1/9-41	In Table 9-41, “EIVPRn Field Descriptions,” removed note in bit description for P field requiring it to be set to active high for PCI Express INTx.
9.4.3/9-56	Removed Section 9.4.3.1, “Shared Message Signaled Interrupts,” and Section 9.4.3.2, “PCI Express INTx.”
10.3.1.2/10-6	In Figure 10-3, “I^2C Frequency Divider Register (I2CFDR),” changed reset value of I2CFDR to “All zeros.”
10.3.1.2/10-6	In Table 10-5, “I2CFDR Field Descriptions,” revised FDR field description, as follows:

Table 10-5. I2CFDR Field Descriptions

Bits	Name	Description
2–7	FDR	Frequency divider ratio. Used to prescale the clock for bit rate selection. The serial bit clock frequency of SCL is equal to one half the platform (MPX) clock divided by the designated divider. Note that the frequency divider value can be changed at any point in a program. The serial bit clock frequency divider selections are described as follows:

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				

10.3.1.5/10-10 In Table 10-8, “I2CDR Field Descriptions,” in DATA description, modified last sentence to say, “Note that in both master receive and slave receive modes, the very first read is always a dummy read.”

10.4.5/10-17 Revised description of serial bit clock in description, as follows:

“If boot sequencer mode is selected on POR (by the settings on the cfg_boot_seq[0:1] reset configuration signals, as described in Section 4.4.3.11, “Boot Sequencer Configuration”), the I^2C1 module communicates with one or more EEPROMs through the I^2C interface on IIC1_SCL and IIC1_SDA. The boot sequencer accesses the I^2C1 serial ROM device at a serial bit clock frequency equal to the platform (MPX) clock frequency divided by 2560. The EEPROM(s) can be programmed to initialize one or more configuration registers of this integrated device.”

10.5.4/10-22 Removed sentence, “For 1-byte transfers, a dummy read should be performed by the interrupt service routine.”

11.3.1.3/11-7 In Table 11-8, “Baud Rate Examples,” replaced 667 MHz entries with 600 MHz entries.

12.3.1.13/12-27 Added Figure 12-16, “Transfer Error Attributes Register (LTEATR),” as follows:

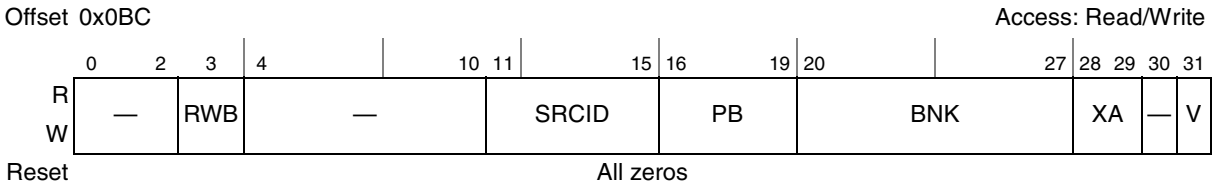

Figure 12-16. Transfer Error Attributes Register (LTEATR)

12.3.1.13/12-27 In Table 12-19, "LTEATR Field Descriptions," added fields XA (bits 28–29) and Reserved (bit 30), and revised bit encodings in SRC_ID field description, as follows:

Table 12-19. LTEATR Field Descriptions

Bits	Name	Description
11–15	SRCID	Transaction source for the error 00000 PCI Express interface 1 00001 PCI Express interface 2/RapidIO 00011–01001 Reserved 01010 Boot sequencer 01011–01111 Reserved 10000 Core 0 (instruction/data) 10001 Reserved 10010 Core 1 (instruction/data) 10011–10100 Reserved 10101 DMA 10110–10111 Reserved 11000 eTSEC 1 11001 eTSEC 2 11010 eTSEC 3 11011 eTSEC 4 11100 RapidIO message unit 11101 RapidIO doorbell unit 11110 RapidIO port-write unit 11111 Reserved
28–29	XA	Extended address for the error. These bits capture the two msbs of the 34-bit address of the transaction resulting in an error.
30	—	Reserved

12.4.4.4.1/12-64 In Table 12-28, "RAM Word Field Descriptions," added the following note to fields LOOP and AMX: "AMX must not change values in any RAM word which begins a loop."

12.4.4.4.7/12-69 Added the following note: "AMX must not change values in any RAM word which begins a loop."

13.3/13-4 Modified first bullet, removing specific maximum data clock frequency ratios, referring reader to the device hardware specificiations document for specific maximum frequencies.

In addition, changed "eTSECs 0 and 1, 2, and 3," to "eTSECs 1, 2, 3, and 4."

13.4/13-6 In Table 13-1, "eTSECn Network Interface Signal Properties," added MII to designated modes during which TXD[7:4] and RXD[7:4] are unused.

13.4/13-6 In Table 13-1, "eTSECn Network Interface Signal Properties," for RGMII and RTBI protocols, changed description of TSEC_RX_ER from "Unused, output driven low" to "Unused."

13.4/13-6 In Table 13-1, "eTSECn Network Interface Signal Properties," modified TSEC*n*_TX_CLK signal description, as follows:

	"MII—transmit clock, input TBI—PMA receive clock 1, input RMII—reference transmit and receive clock, input FIFO—transmit clock, input RGMII, RTBI—unused"
13.4.1/13-8	In Table 13-2, "eTSEC Signals—Detailed Signal Descriptions," for signal TSEC*n*_GTX_CLK, removed the following text from all devices: "In RMII mode, TSEC*n*_GTX_CLK feeds back the effective transmit clock according to the interface, 100Base-T is 25 MHz and 10Base-T is 2.5 MHz."
13.4.1/13-8	In Table 13-2, "eTSEC Signals—Detailed Signal Descriptions," changed TSEC*n*_TX_EN signal description from "rising and falling edges of the TSEC*n*_TX_CLK, respectively" to "rising and falling edges of the TSEC*n*_GTX_CLK, respectively.
13.4.1/13-8	In Table 13-2, "eTSEC Signals—Detailed Signal Descriptions," changed the "State Meaning" description for the TSEC*n*_CRS, as follows: "MII—carrier sense, input TBI—signal detect, input FIFO—receive flow control, output
13.4.1/13-8	In Table 13-2, "eTSEC Signals—Detailed Signal Descriptions," modified TSEC*n*_RX_CLK signal description, as follows: GMII, MII, RGMII—receive clock, input TBI—PMA receive clock 0, input FIFO—receive clock, input"
13.5.2/13-13	In Table 13-4, "Module Memory Map," removed eTSEC FIFO Control and Status Registers.
13.5.3.1.2/13-22	In Table 13-6, "TSEC_ID2 Field Descriptions," modifed description for TSE_INT field, as follows:

Table 13-6 TSEC_ID2 Field Descriptions

Bits	Name	Description
10–15	TSEC_INT	Interface mode support. See Table 13-7 for settings. For eTSEC1: 0x30 For eTSEC2: 0x38 For eTSEC3: 0x30 For eTSEC4: 0x38

13.5.3.1/13-21	In Table 13-10, "ECNTRL Field Descriptions," and Table 13-11, "eTSEC Interface Configurations," updated RMM field description, as follows: "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. See Section 4.4.3.13, "eTSECn Width."

0 Non-RMII interface mode

1 RMII interface mode”

13.5.3.3.1/13-39 In Table 13-20, “TCTRL Field Descriptions,” clarified TFC_PAUSE field description, as follows:

“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.”

13.5.3.3.2/13-41 In Table 13-21, “TSTAT Field Descriptions,” replaced sentence in introduction: “Only enabled rings can indicate halt state” with the following: “The halt bit only has meaning for enabled rings.”

13.5.3.3.2/13-41 In Table 13-21, “TSTAT Field Descriptions,” changed THLT*n* descriptions from “This bit is set only on a general error condition (as in IEVENT[TXE]) or...” to “This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[ENn], or...”

13.5.3.4.1/13-49 In Table 13-31, “RCTRL Field Descriptions,” clarified RSF field description, as follows:

“Receive short frame mode. When set, enables the reception of frames shorter than 64 bytes. For packets received over the FIFO packet interface, this bit has no effect (packets shorter than 64 bytes are always accepted).

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.”

13.5.3.4.4/13-55 In Figure 13-30, “RQUEUE Register Definition,” and Table 13-34, “RQUEUE Field Descriptions,” removed EX0–EX7 fields.

13.5.3.4.5/13-56 In Table 13-35, “RBIFX Field Descriptions,” modified B*n*CTL field descriptions to clarify that arbitrary extraction of preamble is not supported in FIFO modes.

13.5.3.4.8/13-59 In Table 13-38, “RQFPR Field Descriptions,” updated ETY field description with the following note:

“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 carring a type designation between 1500 and 1536 (note there are none currently publically 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."

13.5.3.4.8/13-59 In Table 13-38, "RQFPR Field Descriptions," added the following text to IPF field description:

"See the descriptions of receive FCB fields IP and PRO in Section 13.6.4.3, "Receive Path Off-Load," for more information on determining the status of received packets for which IPF is set."

13.5.3.4.8/13-59 In Table 13-38, "RQFPR Field Descriptions," added clarification to ETY field description describing parser/filer behavior with ethertype field (parsed into ETY), as follows:

"Note that the eTSEC filer gets multiple packet attributes as a result of parsing the packet. The behavior of the eTSEC is that it will pull the innermost ethertype found in the packet; this means that in many supported protocols, it is impossible to create a filer rule that will match on the outer ethertype. There are four cases that need to be highlighted:

1. The jumbo ethertype (0x8870)—In this case, the eTSEC assumes that the following header is LLC/SNAP. LLC/SNAP has an associated etherype, and the ETY field will be 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 PID1 VLN bit to indicate that the packet had a VLAN tag.
4. MPLS tagged packets; this is a defect as described in IPGear 435. 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."

13.5.3.6.2/13-69 In Table 13-45, "MACCFG2 Field Descriptions," clarified that TX preamble length may be from 0x3 to 0xF.

13.5.3.6.2/13-69 In Table 13-45, "MACCFG2 Field Descriptions," modified note to PAD/CRC field description to include the following: "This bit must be set when in half-duplex mode (MACCFG2[Full Duplex] is cleared)."

13.5.3.7.26/13-93 Changed Table 13-85, "TPKT Field Descriptions," to show field TPKT as 22 bits (bits 10–31).

13.5.3.9.1/13-110 In Table 13-110, "FIFOCFG Field Descriptions," modified IPG field description, as follows:

"Minimum inter packet gap. This sets the minimum number of cycles inserted between back-to-back frames transmitted over the FIFO interface. The minimum required is 3 cycles if CRCAPP=0, 5 cycles for 16-bit interfaces if CRCAPP=1 and 7 cycles for 8-bit interfaces if CRCAPP=1."

Section, Page No.	Changes
13.5.3.10.1/13-111	In Figure 13-108, “ATTR Register Definition,” and Table 13-111, “ATTR Field Descriptions,” removed ELCWT and BDLWT fields.
13.5.3.10.2/13-112	Removed Section 13.5.3.10.2, “Attribute Extract Length and Extract Index Register (ATTRELI).”
13.5.4.1/13-113	Added Section 13.5.4.1, “Lossless Flow Control Configuration Registers,” as follows:

13.5.4.1 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 13.6.6.1, “Back Pressure Determination through Free Buffers,” for the theory of operation of these registers.

13.5.4.1.1 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 13-104 describes the definition for the RQPRM*n* register.

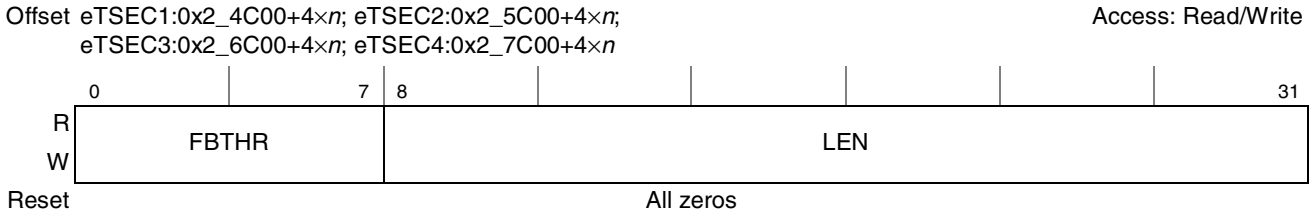

Figure 13-104. RQPRM Register Definition

Table 13-108 describes the fields of the RQPRM register.

Table 13-108. 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.

13.5.4.1.2 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 13.6.6.1, "Back Pressure Determination through Free Buffers."
Figure 13-105 describes the definition for the RFBPTR*n* register.

Offset eTSEC1:0x2_4C40+8×*n*; eTSEC2:0x2_5C40+8×*n*;
eTSEC3:0x2_6C40+8×*n*; eTSEC4:0x2_7C40+8×*n*

Access: Read/Write

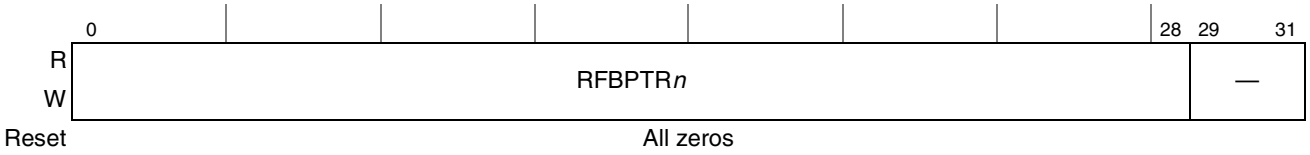

Figure 13-105. RFBPTR0–RFBPTR7 Register Definition

Table 13-109 describes the fields of the RFBPTR*n* registers.

Table 13-109. 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.

13.5.4.3.10/13-125 In Table 13-124, "TBICON Field Descriptions," updated/clarified the asserted state description for the Clock Select field, as follows:

"Clock select. This bit is cleared by default.

0 Allow the TBI to accept dual split-phase 62.5 MHz receive clocks.

1 Configure the TBI to accept a 125 MHz receive clock from the SerDes/PHY. The 125 MHz clock must be physically connected to 'PMA receive clock 0' if using a parallel (non-SGMII) Ethernet protocol."

13.6.2/13-134 Corrected the last bullet of this section regarding minimum inter-packet gap requirements, as follows:

	"On transmission, the minimum inter-packet gap (set in FIFOCFG[IPG]) is three cycles if CRC is not automatically appended. Each CRC data beat adds to this requirement. For 16-bit FIFO interfaces the minimum Tx IPG is 5 cycles and for 8-bit FIFO interfaces the minimum is 7 cycles."
13.6.3.8/13-152	Clarified last sentence of Magic Packet Mode description to include multicast packets, as follows: "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."
13.6.3.11/13-156	Revised description of inter-frame gap timing, as follows: "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. • 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)."
13.6.3.13/13-156	In Table 13-141, "Reception Errors," removed the following note for Parser error: **"Note:** Any values in the length/type field between 1500 and 1536 will be 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."
13.6.4.3/13-161	In Table 13-143, "Rx Frame Control Block Descriptions," modified description of PRO, as follows: "If IP = 1, PRO is set as follows:

Section, Page No.	Changes
	PRO=0xFF for a fragment header or a back to back route header PRO=0x*nn* 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 If IP = 0, PRO is undefined.
13.6.4.3/13-161	In Table 13-143, "Rx Frame Control Block Descriptions," added the following text to IP field description. "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. RQFPR[IPF] (see Section 13.5.3.4.8, "Receive Queue Filer Table Property Register (RQFPR)") indicates that the packet was fragmented."
13.6.4.3/13-161	In Table 13-143, "Rx Frame Control Block Descriptions," added the following text to PRO field description: "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] will be set. IP header checksumming will operate and perform as intended. Most of the time, the eTSEC will update the RxFCB[PRO] field and RQFPR[L4P] fileds with whatever value was found in the protocol field of the IP header. See Section 13.5.3.3.8, "Receive Queue Filer Table Property Register (RQFPR)," for a description of RQFPR."
13.6.6/13-170	Added Section 13.6.6, "Loss Flow Control," as follows:

13.6.6 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 be not be issued in time to prevent BSY errors.

To allow the eTSEC to generate the pause request automatically, additional information (a pointer the last free BD and ring length) is required.

13.6.6.1 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 (96bit 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 fourBDs 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 RFBPTRn < RBPTR*n*

$$\text{FreeBDs} = \text{RQPRMn[LEN]} - \text{RBPTRn} + \text{RFBPTRn}$$

When RFBPTRn > RBPTR*n*

$$FreeBDs = RFBPTRn - RBPTRn$$

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. 13.6.3.9/13-154) 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.

13.6.6.2 Software Use of Hardware-Initiated Back Pressure

13.6.6.2.1 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.

13.6.6.2.2 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*n*[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.

13.6.6.1/13-170	Changed sentence in first paragraph to say, “Because of pre-fetching, a minimum of four buffer descriptors per ring are required.”
13.6.6.3/13-175	In Table 13-150, “Receive Buffer Descriptor Field Descriptions,” added recommendation to use 64-byte aligned receive buffer pointer addresses to description of Rx Data Buffer Pointer (offset 4–7, bits 0–31).
13.7.1.8/13-204	In Table 13-171, “16-Bit FIFO Interface Mode Signal Configuration (eTSECs1 & 2),” and Table 13-172, “16-Bit FIFO Interface Mode Signal Configuration (eTSECs 3 & 4),” removed support for 3.3-V FIFO interface.
14.3.1.1/14-10	In Table 14-5, “MRn Field Descriptions,” clarified CS field description, as follows: “Channel start. This bit is also automatically set by hardware during single-write start mode and external master start enable mode. Note that in external control mode, deasserting $\overline{\text{DMA\_DREQ}}$ does NOT clear this bit. 0 Halt the DMA process if channel is busy (SR*n*[CB] is set). No effect if the channel is not busy.

Section, Page No.	Changes
	1 Start the DMA process if channel is not busy (CB is cleared). If the channel was halted (CS = 0 and CB = 1), the transfer continues from the point at which it was halted.”
14.4.1.3/14-33	Added clarification to third paragraph, describing external control functionality, as follows: “The external control and the DMA controller use a well defined protocol to communicate. The external control can start or restart a paused DMA transfer. The DMA controller acknowledges a DMA transfer in progress and also indicates a transfer completion. Note that external control cannot cause a channel to enter a paused state.”
14.4.1.3/14-33	Added clarification to first bullet of fifth paragraph, describing the use of the signal, $\overline{\text{DMA\_DREQ}}$, as follows: • $\overline{\text{DMA\_DREQ}}$—Asserting edge triggers a DMA transfer start or restart from a pause request. Sets MR*n*[CS]. (Note that negating $\overline{\text{DMA\_DREQ}}$ does NOT clear MR*n*[CS].)
14.4.1.4/14-34	Modified second paragraph, as follows: “If CC is set by software while the channel is busy with a transfer, the DMA controller finishes all transfers until it reaches the EOLND in basic mode or EOLSD in extended mode. The DMA controller then refetches the last link descriptor in basic mode, or the last list descriptor in extended mode and clears the channel continue bit. If EOLND or EOLSD is still set for their respective modes, the DMA controller remains in the idle state. If EOLND or EOLSD is not set, the DMA controller continues the transfer by refetching the new descriptor. The channel busy (SR*n*[CB]) bit is cleared when the DMA controller reaches EOLND/EOLSD and is set again when it initiates the refetch of the link or list descriptor.”
14.4.2/14-36	Added the following sentence: “Note that a single DMA transfer in any of the direct or chaining modes must not cross a 16GB (34-bit) address boundary.”
Chapter 15/15-1	Change all instances of OMMR[MM] to OMDATR[MM] throughout.
14.5/14-41	Removed bullets also found in Section 14.5.1, “Unusual DMA Scenarios.”
15.6.2.4/15-23	In Table 15-19, “GCCSR Field Descriptions,” added notes clarifying that although the register is R/W it is, in fact, a status register (not a control register). As such, manually changing the value of bitfields does not affect logical behavior.
15.6.4.3/15-37	In Table 15-33, “ECACSR Field Descriptions,” clarified description of ECI field, as follows: “Extended capture information [0:15]. ECI contains the control/data character signal corresponding to each byte of captured data. Each ECI bit reflects the validity of captured data. If a bit is set, then the designated byte of captured data is valid. If a bit is cleared, then the designated

byte of the specified register does not contain valid data and should be disregarded until the bit is set.

ECI[0] reflects validity of PCSECCSR0[0:7]

ECI[1] reflects validity of PCSECCSR0[8:15]

ECI[2] reflects validity of PCSECCSR0[16:23]

ECI[3] reflects validity of PCSECCSR0[24:31]

ECI[4] reflects validity of PECCSR1[0:7]

ECI[5] reflects validity of PECCSR1[8:15]

...

ECI[14] reflects validity of PECCSR3[16:23]

ECI[15] reflects validity of PECCSR3[24:31]"

15.10.2.1.7/15-170 — In Table 15-131, "Outbound Doorbell Programming Errors," replaced three instances of ODMR[EIE] with OM*n*MR[EIE].

16.1.1.1/16-2 — Added note regarding checking the link status before issuing outbound transactions after reset or when recovering from a linkdown event.

16.3.1/16-5 — In Table 16-3, "PCI Express Memory-Mapped Register Map," added footnotes for reset value when alternate boot vector is selected, as follows:

1. If the device is configured to use the alternate boot vector (cfg_boot_vec = 0), the reset value for PCI controller 1 PEXOTAR3 is 0x000F_FFF0.

2. If the device is configured to use the alternate boot vector (cfg_boot_vec = 0), the reset value for PCI controller 1 PEXOWBAR3 is 0x000F_FF00.

3. If the device is configured to use the alternate boot vector (cfg_boot_vec = 0), the reset value for PCI controller 1 PEXOWAR3 is 0x8004_400F.

16.3.1/16-5 — In Table 16-3, "PCI Express Memory-Mapped Register Map," changed reset value for PEXOWAR3 to 0x0000_0000.

16.3.2.3/16-11 — In Table 16-6, "PEX_OTB_CPL_TOR Field Descriptions," clarified description of TC field, as follows:

"Timeout counter. This is the value that is used to load the response counter of the completion timeout.

One TC unit is 8× the PCI Express controller clock period; that is, one TC unit is 20 ns at 400 MHz, and 30 ns at 266.66 MHz.

The following are examples of timeout periods based on different TC settings:

0x00_0000Reserved

0x10_FFFF22.28 ms at 400 MHz controller clock; 33.34 ms at 266.66 MHz controller clock

0xFF_FFFF335.54 ms at 400 MHz controller clock; 503.31 ms at 266.66 MHz controller clock

Section, Page No.	Changes
16.3.2.4/16-11	In Table 16-7, “PEX_CONF_RTY_TOR Field Descriptions,” clarified description of TC field, as follows: “Timeout counter. This is the value that is used to load the CRS response counter. One TC unit is 8× the PCI Express controller clock period; that is, one TC unit is 20 ns at 400 MHz and 30 ns at 266.66 MHz. Timeout period based on different TC settings: 0x000_0000Reserved 0x400_FFFF1.34 s at 400 MHz controller clock, 2.02 s at 266.66 MHz controller clock 0xFFF_FFFF5.37 s at 400 MHz controller clock, 8.05 s at 266.66 MHz controller clock”
16.3.5.1/16-19	In Table 16-14, “PCI Express Outbound Translation Address Registers (PEXOTARn),” Table 16-15, “PCI Express Outbound Translation Extended Address Registers (PEXOTEARn),” and Table 16-16, “PCI Express Outbound Window Base Address Registers (PEXOWBARn),” added footnote for reset value when alternate boot vector is selected.
16.3.5.1.4/16-22	In Figure 16-18, “PCI Express Outbound Window Attributes Registers 1–4 (PEXOWARn),” revised footnote for reset value of PEXOWAR3 when alternate boot vector is not selected, as follows: “If the device is configured to use the alternate boot vector (cfg_boot_vec = 0), the reset value for PCI controller 1 PEXOWAR3 is 0x8004_400F. If the device is not configured to use the alternate boot vector (cfg_boot_vec = 1), the reset value for PCI controller 1 PEXOWAR3 is 0x0000_0000.”
16.3.6.1/16-29	In Table 16-23, “PCI Express Error Detect Register Field Descriptions,” added note recommending hot reset after a completion time-out is detected to bit description for PCT.
16.3.6.4/16-35	In Table 16-26, “PCI Express Error Capture Status Register Field Descriptions,” revised bit encodings for GSID field, as follows: “Global source ID. This field indicates the internal platform global source ID that the error transaction originates. This field only applies to non PEX_CONFIG_ADDR/PEX_CONFIG_DATA transactions. 00000 PCI Express 1 00001 PCI Express 2/RapidIO 00010 Reserved 01010 Boot sequencer 10000 Core 0 instruction/data 10010 Core 1 instruction/data 10101 DMA 11000 eTSEC1

Section, Page No.	Changes
	11001 eTSEC2 11010 eTSEC3 11011 eTSEC4 11100 RapidIO message/doorbell/port write with responses and reads All other settings reserved.”
16.3.6.7/16-39	In Table 16-23, “PCI Express Error Detect Register Field Descriptions,” Table 16-24, “PCI Express Error Interrupt Enable Register Field Descriptions,” and Table 16-25, “PCI Express Error Disable Register Field Descriptions,” changed descriptions for OD0, OD1, and OD2 to “Internal platform transaction information. Reserved for factory debug.”
16.3.7.2/16-43	In the first paragraph, added statement that the PCI Express Controller Internal CSR registers are not accessible by inbound PCI Express configuration transactions.
16.3.9.11/16-73	In Table 16-84, “PCI Express Link Control Register Field Description,” augmented field description for RL (bit 5), as follows: “Retrain link (Reserved for EP devices). In RC mode, setting this bit initiates link retraining by directing the Physical Layer LTSSM to the Recovery state; reads of this bit always return 0.”
16.3.10/16-80	Added note to Figure 16-101, “PCI Express Extended Configuration Space,” stating that the PCI Express Controller Internal CSRs are not accessable by inbound PCI Express configuration transactions and any attempts to access these registers will return all 0s.
16.3.10/16-80	In Figure 16-101, “PCI Express Extended Configuration Space,” revised range for Internal CSR space to 0x400–0x6FF.
16.3.10.2/16-81	In Table 16-97, “PCI Express Uncorrectable Error Status Register Field Description,” added note recommending hot reset after a completion time-out is detected to bit description for CTO.
16.3.10.18/16-93	Clarified that the reset value of the Configuration Ready Register is defined during POR configuration (host/agent mode and CPU boot configuration).

16.4.1.10/16-98 Added Section 16.4.1.10, “Error Handling,” as follows:

16.4.1.10 Error Handling

The PCI Express specification classifies errors as correctable and uncorrectable. Correctable errors result in degraded performance, but uncorrectable errors generally result in functional failures. As shown in Figure 16-130 uncorrectable errors can further be classified as fatal or non-fatal.

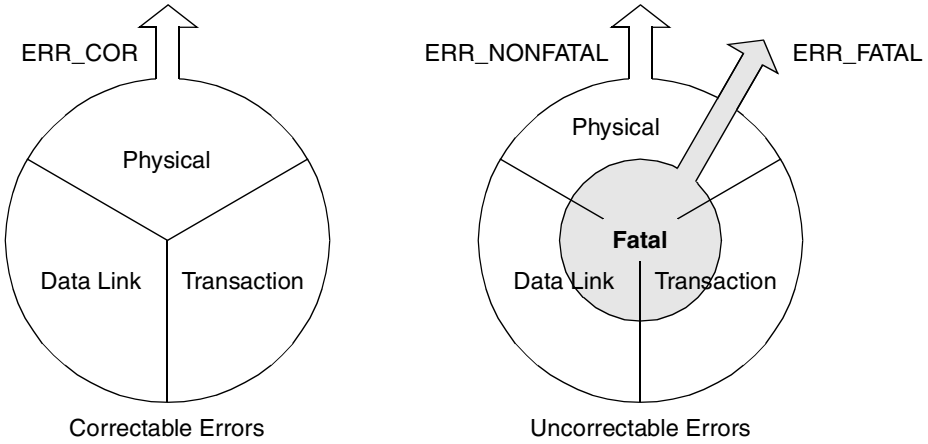

Figure 16-30. PCI Express Error Classification

16.4.1.10.1 PCI Express Error Logging and Signaling

Figure 16-131 shows the PCI Express-defined sequence of operations related to signaling and logging of errors detected by a device. Note that the PCI Express controller on this device supports the advanced error handling capabilities shown within the dotted lines.

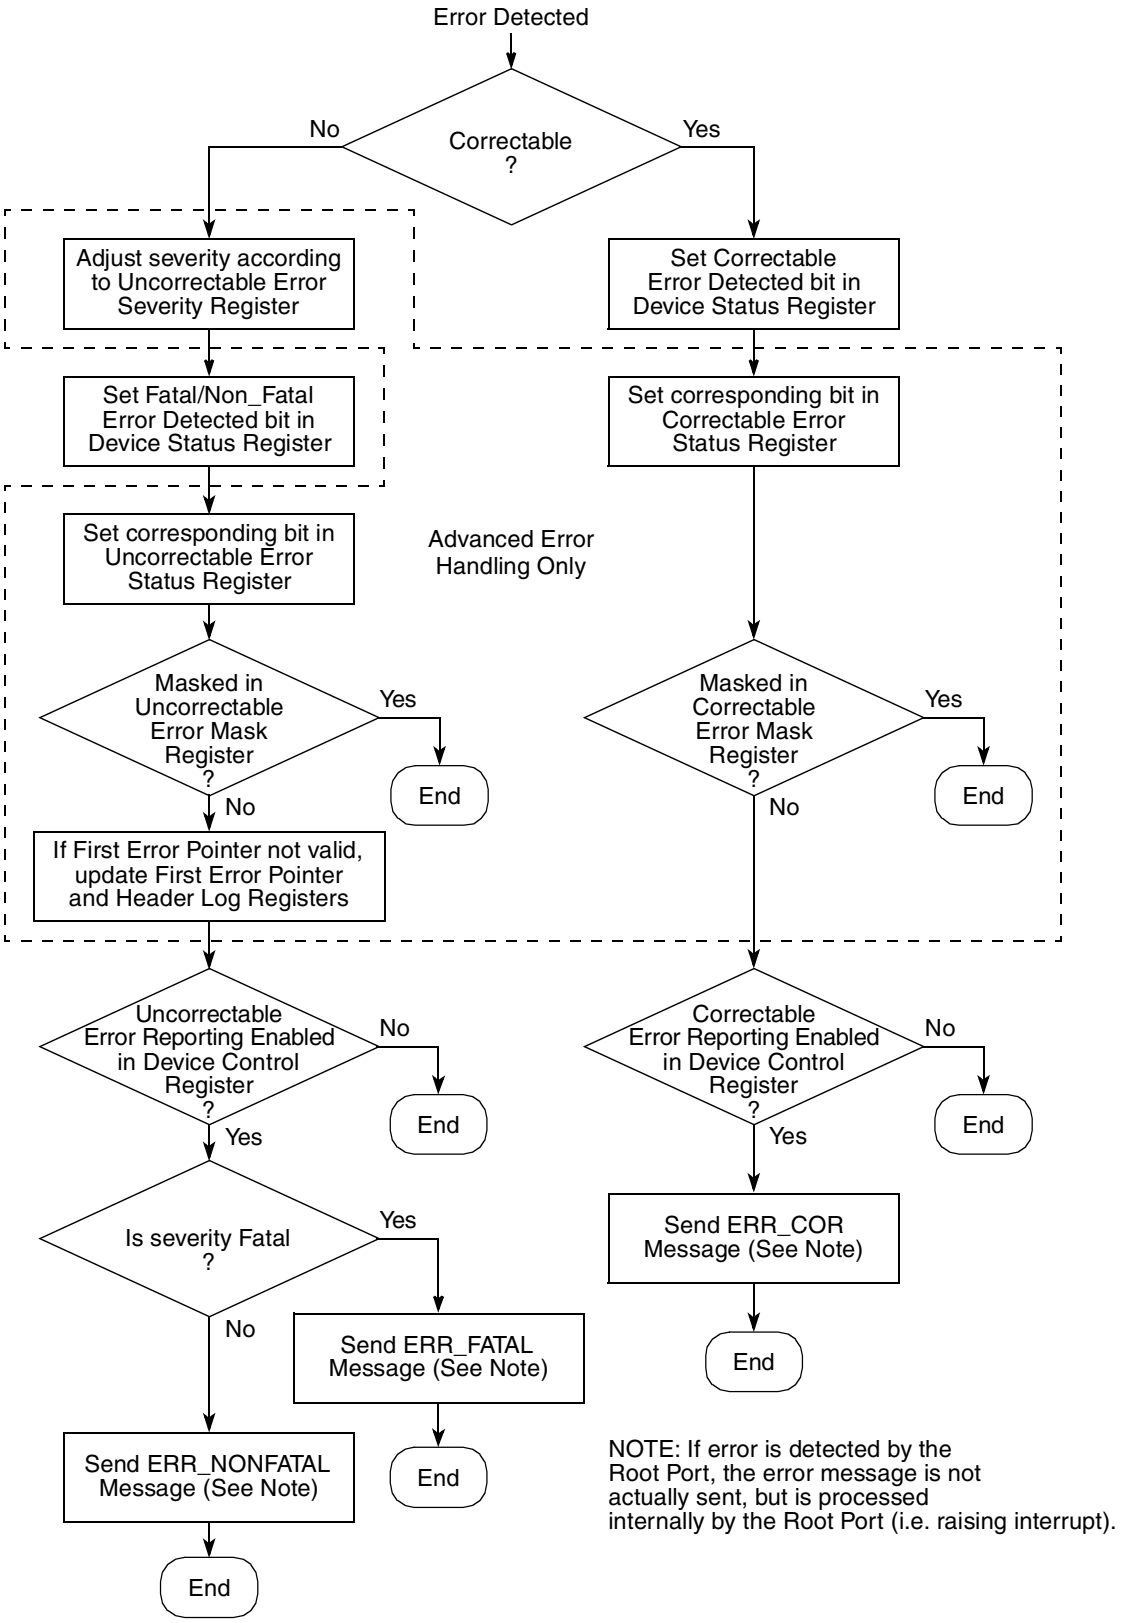

Figure 16-131. PCI Express Device Error Signaling Flowchart

16.4.1.10.2 PCI Express Controller Internal Interrupt Sources

Table 16-125 describes the sources of the PCI Express controller internal interrupt to the PIC and the preconditions for signaling the interrupt.

Table 16-125. PCI Express Internal Controller Interrupt Sources

Status Register Bit	Preconditions
Any bit in PEX_PME_MES_DR set	The corresponding interrupt enable bits must be set in PEX_PME_MES_IER
Any bit in PEX_ERR_DR set	The corresponding interrupt enable bits must be set in PEX_ERR_EN.
PCI Express Root Status Register[16] (PME status) is set	PCI Express Root Control Register [3] (PME interrupt enable) is set
PCI Express Root Error Status Register[6] (fatal error messages received) is set	PCI Express Root Error Command Register [2] (fatal error reporting enable) is set or PCI Express Root Control Register [2] (system error on fatal error enable) is set
PCI Express Root Error Status Register [5] (non-fatal error messages received) is set	PCI Express Root Error Command Register [1] (non-fatal error reporting enable) is set or PCI Express Root Control Register [1] (system error on non-fatal error enable) is set
PCI Express Root Error Status Register[0] (correctable error messages received) is set	PCI Express Root Error Command Register[0] (correctable error reporting enable) is set or PCI Express Root Control Register[0] (system error on correctable error enable) is set.
Any correctable error status bit in PCI Express Correctable Error Status Register is set	The corresponding error mask bit in PCI Express Correctable Error Mask Register is clear and PCI Express Root Error Command Register[0] (correctable error reporting enable) is set
Any fatal uncorrectable error status bit in PCI Express Uncorrectable Error Status Register is set. (The corresponding error is classified as fatal based on the severity setting in PCI Express Uncorrectable Error Severity Register.)	The corresponding error mask bit in PCI Express Uncorrectable Error Mask Register is clear and either PCI Express Device Control Register[2] (fatal error reporting) is set or PCI Express Command Register[8] (SERR) is set.
Any non-fatal uncorrectable error status bit in PCI Express Uncorrectable Error Status Register is set. (The corresponding error is classified as non-fatal based on the severity setting in PCI Express Uncorrectable Error Severity Register.)	The corresponding error mask bit in PCI Express Uncorrectable Error Mask Register is clear and either PCI Express Device Control Register[1] (non-fatal error reporting) is set or PCI Express Command Register[8] (SERR) is set.
PCI Express Secondary Status Register[8] (master data parity error) is set.	PCI Express Secondary Status Interrupt Mask Register[0] (mask master data parity error) is cleared and PCI Express Command Register[6] (parity error response) is set.
PCI Express Secondary Status Register[11] (signaled target abort) is set	PCI Express Secondary Status Interrupt Mask Register[1] (mask signaled target abort) is cleared.
PCI Express Secondary Status Register[12] (received target abort) is set	PCI Express Secondary Status Interrupt Mask Register[2] (mask received target abort) is cleared.

Table 16-125. PCI Express Internal Controller Interrupt Sources (continued)

Status Register Bit	Preconditions
PCI Express Secondary Status Register[13] (received master abort) is set	PCI Express Secondary Status Interrupt Mask Register[3] (mask received master abort) is cleared.
PCI Express Secondary Status Register[14] (signaled system error) is set.	PCI Express Secondary Status Interrupt Mask Register[4] (mask signaled system error) is cleared.
PCI Express Secondary Status Register[15] (detected parity error) is set	PCI Express Secondary Status Interrupt Mask Register[5] (mask detected parity error) is cleared.
PCI Express Slot Status Register[0] (attention button pressed) is set	PCI Express Slot Control Register[0] (attention button pressed enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set and either PCI Express PM Control Register[1–0] = 00 (the function power state is D0) or PCI Express PM Control Register[8] (PME enable) is set.
PCI Express Slot Status Register[1] (power fault detected) is set	PCI Express Slot Control Register[1] (power fault detected enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set and either PCI Express PM Control Register[1–0] = 00 (the function power state is D0) or PCI Express PM Control Register[8] (PME enable) is set.
PCI Express Slot Status Register[2] (MRL sensor changed) is set	PCI Express Slot Control Register[2] (MRL sensor changed enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set and either PCI Express PM Control Register[1–0] = 00 (the function power state is D0) or PCI Express PM Control Register[8] (PME enable) is set.
PCI Express Slot Status Register[3] (presence detect changed) is set	PCI Express Slot Control Register[3] (presence detect changed enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set and either PCI Express PM Control Register[1–0] = 00 (the function power state is D0) or PCI Express PM Control Register[8] (PME enable) is set.
PCI Express Slot Status Register[4] (command completed) is set	PCI Express Slot Control Register[4] (command completed interrupt enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set.

16.4.1.10.3 Error Conditions

Table 16-126 describes specific error types and the action taken for various transaction types.

Table 16-126. Error Conditions

Transaction Type	Error Type	Action
Inbound response	PEX response time out. This case happens when the internal platform sends a non-posted request that did not get a response back after a specific amount of time specified in the outbound completion timeout register (PEX_OTB_CPL_TOR)	Log error (PEX_ERR_DR[PCT]) and send interrupt to PIC, if enabled.
Inbound response	Unexpected PEX response. This can happen if, after the response times out and the internal queue entry is deallocated, the response comes back.	Log unexpected completion error (PCI Express Uncorrectable Status Register[16]) and send interrupt to PIC, if enabled.
Inbound response	Unsupported request (UR) response status	Depending upon whether the initial internal request was broken up, the error is not sent until all responses come back for all portions of the internal request. Log the error (PEX_ERR_DR[CDNSC] and PCI Express Uncorrectable Status Register[20]) and send interrupt to PIC, if enabled.
Inbound response	Completer abort (CA) response status	Depending upon whether the initial internal request was broken up, the error is not sent until all responses come back for all portions of the internal request. Log the error (PEX_ERR_DR[PCAC, CDNSC] and PCI Express Uncorrectable Status Register[15] and send interrupt to PIC, if enabled.
Inbound response	Poisoned TLP (EP=1)	Depending upon whether the initial internal request was broken up, the error is not sent until all responses come back for all portions of the internal request. Log the error (PCI Express Uncorrectable Status Register[12]) and send interrupt to PIC, if enabled.
Inbound response	ECRC error	Depending upon whether the initial internal request was broken up, the error is not sent until all responses come back for all portions of the internal request. Log the error (PCI Express Uncorrectable Status Register[19]) and send interrupt to PIC, if enabled.
Inbound response	Configuration Request Retry Status (CRS) timeout for a configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	1.The controller always retries the transaction as soon as possible until a status other than CRS is returned. However, if a CRS status is returned after the configuration retry timeout (PEXCONF_RTY_TOR) timer expires, then the controller aborts the transaction and sends all 1s (0xFFFF_FFFF) data back to requester. 2. Log the error (PEX_ERR_DR[PCT]) and send interrupt to the PIC, if enabled.
Inbound response	UR response for configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	1. Send back all 1s (0xFFFF_FFFF) data. 2. Log the error (PEX_ERR_DR[CDNSC] and PCI Express Uncorrectable Status Register[20]) and send interrupt to PIC, if enabled.

Table 16-126. Error Conditions (continued)

Transaction Type	Error Type	Action
Inbound response	CA response for Configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	1. Send back all 1s (0xFFFF_FFFF) data. 2. Log the error (PEX_ERR_DR[PCAC, CDNSC] and PCI Express Uncorrectable Status Register[15]) and send interrupt to PIC, if enabled.
Inbound response	Poisoned TLP (EP=1) response for Configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	1. Send back all 1s (0xFFFF_FFFF) data. 2. Log the error (PCI Express Uncorrectable Status Register[12]) and send interrupt to PIC, if enabled.
Inbound response	ECRC error response for Configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	1. Send back all 1s (0xFFFF_FFFF) data. 2. Log the error (PCI Express Uncorrectable Status Register[19]) and send interrupt to PIC, if enabled.
Inbound response	Configuration Request Retry Status (CRS) response for Configuration transaction that originates from ATMU	1.The controller always retries the transaction as soon as possible until a status other than CRS is returned. However, if a CRS status is returned after the configuration retry timeout (PEXCONF_RTY_TOR) timer expires, then the controller aborts the transaction. 2. Log the error (PEX_ERR_DR[CRST]) and send interrupt to the PIC, if enabled.
Inbound response	UR response for Configuration transaction that originates from ATMU	Log the error (PEX_ERR_DR[CDNSC] and PCI Express Uncorrectable Status Register[20]) and send interrupt to PIC, if enabled.
Inbound response	CA response for Configuration transaction that originates from ATMU	Log the error (PEX_ERR_DR[PCAC, CDNSC] and PCI Express Uncorrectable Status Register[15]) and send interrupt to PIC, if enabled.
Inbound response	Malformed TLP response	PCI Express controller does not pass the response back to the core. Therefore, a completion timeout error eventually occurs.
Inbound request	Poisoned TLP (EP=1)	1. If it is a posted transaction, the controller drops it. 2. If it is a non-posted transaction, the controller returns a completion with UR status. 3. Release the proper credits
Inbound request	ECRC error	1. If it is a posted transaction, the controller drops it. 2. If it is a non-posted transaction, the controller returns a completion with UR status. 3. Release the proper credits
Inbound request	PCI Express nullified request	The packet is dropped.
Outbound request	Outbound ATMU crossing	Log the error (PEX_ERR_DR[OAC]). The transaction is not sent out on the link.
Outbound request	Illegal message size	Log the error (PEX_ERR_DR[MIS]). The transaction is not sent out on the link.
Outbound request	Illegal I/O size	Log the error (PEX_ERR_DR[IOIS]). The transaction is not sent out on the link.
Outbound request	Illegal I/O address	Log the error (PEX_ERR_DR[IOIA]). The transaction is not sent out on the link.
Outbound request	Illegal configuration size	Log the error (PEX_ERR_DR[CIS]). The transaction is not sent out on the link.

Table 16-126. Error Conditions (continued)

Transaction Type	Error Type	Action
Outbound response	Internal platform response with error (for example, an ECC error on a DDR read or the transaction maps to unknown address space).	Send poisoned TLP (EP=1) completion(s) for data that are known bad. If the poison data happens in the middle of the packet, the rest of the response packet(s) is also poisoned.

16.4.6/16-110 Added Section 16.4.6, "Link Down," as follows:

"Typically, a link down condition occurs after a hot reset event; however, it is possible for the link to go down unexpectedly without a hot reset event. When this occurs, a link down condition is detected (PEX_PME_MSG_DR[LDD]=1). Link down is treated similarly to a hot reset condition.

Subsequently, while the link is down, all new posted outbound transactions are discarded. All new non-posted ATMU transactions are errored out. Non-posted configuration transactions issued using PEX_CONFIG_ADDR/PEX_CONFIG_DATA toward the link returns 0xFFFF_FFFF (all 1s). As soon as the link is up again, the sending of transaction resumes.

Note that in EP mode, a link down condition causes the controller to reset all non-sticky bits in its PCI Express configuration registers as if it had been hot reset."

17.4.1.2/17-7 In Figure 17-2, "POR Boot Mode Status Register (PORBMSR)," changed default value for ROM_LOC to *nnnn*.

17.4.1.4/17-9 In Figure 17-4, "POR Device Status Register (PORDEVSR)," in entry for ECW3 = 1 changed "8-bit FIFO" to "16-bit FIFO."

17.4.1.7/17-14 Revised section, as follows:

"The device can be configured, using registers described in the next sections, to have 16 general-purpose input and 16 general-purpose output signals; however eTSEC1 and eTSEC2 must be disabled for their signals to be used for general purpose input and output. In brief, this configuration may be achieved as follows:

1. GPIOCR[RXIN] must be set to enable general-purpose input and GPIOCR[TXOUT] must be set to enable general-purpose output. See Section 17.4.1.7.1, "General-Purpose I/O Control Register (GPIOCR)," for a description of GPIOCR.
2. DEVDISR[eTSEC1] and DEVDISR[eTSEC2] must be set to disable eTSEC1 and eTSEC2. See Section 17.4.1.9, "Device Disable Register (DEVDISR)," for more complete information about DEVDISR."

17.4.1.7.2/17-15 Revised section, as follows:

"GPOUTDR, shown in Figure 17-8, contains the data driven on TSEC1_TXD[0:7] and/or TSEC2_TXD[0:7] when these signals are configured

for use as general-purpose outputs in GPIOCR, as described in Section 17.4.1.7.1, “General-Purpose I/O Control Register (GPIOCR).” Writes to GPOUTDR only affect signals enabled as general-purpose outputs. GPOUTDR may be accessed using single-byte writes (using big-endian addressing) so that writes to one byte do not affect outputs controlled by the other. Reads return the data currently being driven on the outputs only if the associated bits are configured as general-purpose output signals rather than their primary function as eTSEC signals.

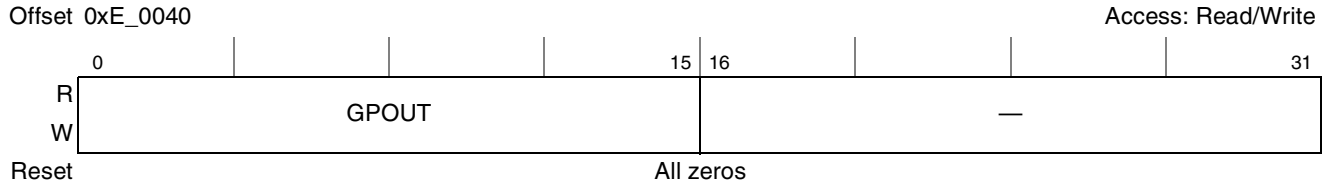

Figure 17-8. General-Purpose Output Data Register (GPOUTDR)

Table 7-11 describes the fields of GPOUTDR.

Table 17-11. GPOUTDR Field Descriptions

Bits	Name	Description
0–15	GPOUT	General-purpose output data. When the corresponding signals are configured to be general-purpose output signals, the values of the bits of GPOUT are driven onto those pins. GPOUT[7:0] corresponds to TSEC1_TXD[7:0] as follows: GPOUT[7] ↔ TSEC1_TXD[7] GPOUT[6] ↔ TSEC1_TXD[6] GPOUT[5] ↔ TSEC1_TXD[5] GPOUT[4] ↔ TSEC1_TXD[4] GPOUT[3] ↔ TSEC1_TXD[3] GPOUT[2] ↔ TSEC1_TXD[2] GPOUT[1] ↔ TSEC1_TXD[1] GPOUT[0] ↔ TSEC1_TXD[0] GPOUT[15:8] corresponds to TSEC2_TXD[7:0] as follows: GPOUT[15] ↔ TSEC2_TXD[7] GPOUT[14] ↔ TSEC2_TXD[6] GPOUT[13] ↔ TSEC2_TXD[5] GPOUT[12] ↔ TSEC2_TXD[4] GPOUT[11] ↔ TSEC2_TXD[3] GPOUT[10] ↔ TSEC2_TXD[2] GPOUT[9] ↔ TSEC2_TXD1] GPOUT[8] ↔ TSEC2_TXD[0]
16–31	—	Reserved, should be cleared

17.4.1.9/17-17 In Table 17-14, “DEVDISR Field Descriptions,” added note instructing software to place a core in sleep mode prior to disabling it with the DEVDISR.

17.4.1.10/17-20 Added Figure 17-12, “Power Management Control & Status Register (POWMGTCSR),” as follows:

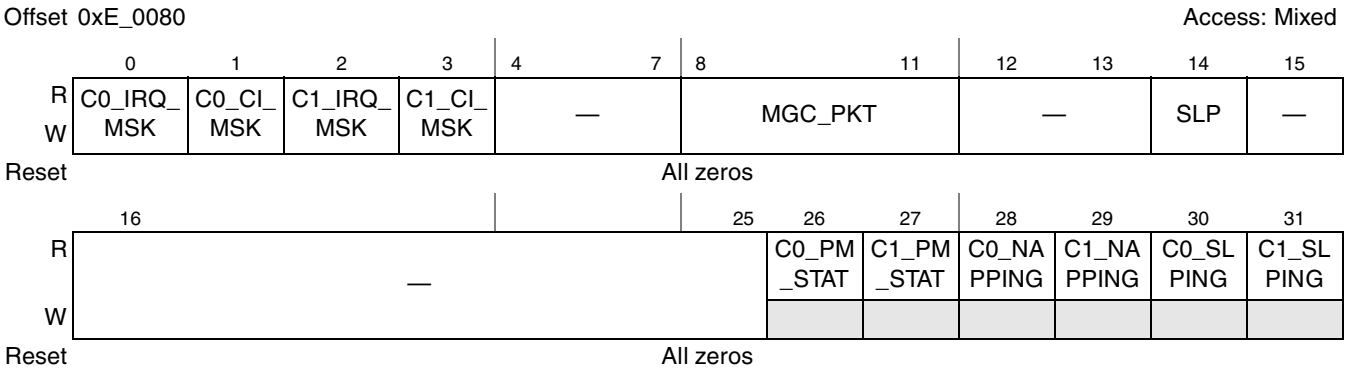

Figure 17-12. Power Management Control & Status Register (POWMGTCSR)

17.4.1.20/17-29 In Figure 17-21, "SerDes 1 Control Register 1 (SRDS1CR1)," Table 17-26, "SRDS1CR1 Field Descriptions," Figure 17-23, "SerDes 2 Control Register 1 (SRDS2CR1)," and Table 17-28, "SRDS2CR1 Field Descriptions," added LBSELTYPE field, as follows:

Bits	Name	Description
21–24	LBSELTYPE	Select type loop-back 0000 Application mode (normal operating mode) 0001 Digital loopback mode All other settings reserved.

18.3.1/18-3 Exchanged address offsets for PMC0 upper and PMC0 lower in Table 18-1, "Control Register Memory Map," and Figure 18-7, "Performance Monitor Counter Register 0 (PMC0)," as follows:

Table 18-1. Control Register Memory Map

Performance Monitor—Block Base Address 0xE_1000				
Offset	**Register**	**Access**	**Reset**	**Section/Page**
0x018	PMC0 (lower)—Performance monitor counter 0 lower	R/W	0x0000_0000	18.3.3.1/18-9
0x01C	PMC0 (upper)—Performance monitor counter 0 upper	R/W	0x0000_0000	18.3.3.1/18-9

18.4.7/18-14 In Table 18-10, "Performance Monitor Events Assignment," revised entries for DDR1 event Ref:13 and DDR2 event Ref:24, as follows:

Table 18-10. Performance Monitor Events Assignment

Event	Event Assignment	Description
DDR Memory Controller 1 Events		

Table 18-10. Performance Monitor Events Assignment (continued)

Event	Event Assignment	Description
Total number of reads or writes from core.	Ref:13	Total number of reads or writes from Core 0. Memory select errors must be enabled (ERR_DISABLE[MSED] = 0) for accurate counting of this event.
Total number of reads or writes from core.	Ref:24	Total number of reads or writes from Core 0. Memory select errors must be enabled (ERR_DISABLE[MSED] = 0) for accurate counting of this event.

18.4.7/18-14 In Table 18-10, “Performance Monitor Events Assignment,” removed entries for MCM events C6:17, C8:15, C4:22, C4:17, C6:18, and C7:15

19.4.1/19-29 Revised source and target ID encodings in Table 19-30, “Source and Target ID Values,” as follows:

Table 19-30. Source and Target ID Values

Source/ Target ID	Source/Target	Source/ Target ID	Source/Target
00	PCI Express 1	10	Core 0 (instruction)
01	PCI Express 2/RapidIO	11	Core 0 (data)
02	Reserved	12	Core 1 (instruction)
03	Reserved	13	Core 1 (data)
04	Local bus controller	14	Reserved
05	Reserved	15	DMA
06	Reserved	16	DDR port 2
07	Reserved	17	System access port (SAP)
08	Configuration space	18	eTSEC1
09	Reserved	19	eTSEC2
0A	Boot sequencer	1A	eTSEC3
0B	Interleaved DDR memory controllers 1 & 2	1B	eTSEC4
0C	Reserved	1C	RapidIO message unit
0D	Reserved	1D	RapidIO doorbell unit
0E	Reserved	1E	RapidIO port-write unit
0F	DDR port 1	1F	Reserved

Section, Page No.	Changes

THIS PAGE INTENTIONALLY LEFT BLANK

Section, Page No.	Changes

THIS PAGE INTENTIONALLY LEFT BLANK

How to Reach Us:

Home Page:
www.freescale.com

email:
support@freescale.com

USA/Europe or Locations Not Listed:
Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
1-800-521-6274
480-768-2130
support@freescale.com

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)
support@freescale.com

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

Asia/Pacific:
Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate,
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:
Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
1-800-441-2447
303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor
@hibbertgroup.com

Document Number: MPC8641DRMAD
Rev. 1.1
07/2008