PCI Express System Architecture Book Errata

Errata for First Edition, First Book Printing

Table 1 lists errata for the PCI Express System Architecture Book, First Edition, First Printing, dated September 2003.

We appreciate that our readers keep us informed about book errata. Please send us any further errata you discover to ravi@mindshare.com so that we may include it in this active list.

We appreciate the efforts of Rob Elliott, Hewlett-Packard Server Storage group, for making significant contributions to this errata.

Document Update History

Document Created: October 15th, 2003

Document Updated: December 5th, 2003

Document Updated: February 10th, 2004

Document Updated: October 13th, 2004

Document Updated: February 28th, 2005

Document Updated: February 26th, 2006

Errata List

Table 1 is a list of errata with a page number and description for each errata. The errata has been classified with a severity rating. The three levels of errata are: Minor, Medium, and Major. A minor error may be a spelling, grammatical or similar error. A medium error may be a clarification or misleading statement in the book. A major error may include a misinterpretation of the specification or an error that may cause the reader to

make a wrong assumption.

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
11	Medium	In the section on “Looking into the Future”, replace Newcard with ExpressCard
47	Medium	Replace NEWCARD with ExpressCard in two places.
54	Minor	In Figure 1-25, change ‘Infiniband’ to ‘InfiniBand’
54	Minor	Last sentence, replace NEWCARD with ExpressCard. Change ‘Electo’ to ‘Electro’ in bullet 2 and 4. Change ‘Fiber Channel’ to ‘Fibre Channel’ in Figure 1-25.
57	Minor	IO Read, IO Write, Configuration Read and Configuration Write should be followed by ‘Request’.
60	Minor	Figure 2-2. Change ‘Locked normal...’ and ‘Locked error ...’ to ‘Locked...’ Also, change the title from ‘....Locked Read...’ to ‘...Memory Read Lock...’
63	Minor	Last sentence, replace ‘vender’ with ‘vendor’
67	Medium	Add a paragraph that indicates the following: ‘PCI Express endpoints may implement IO space for booting”
83	Minor	In Figure 2-22, change ‘Infiniband’ to ‘InfiniBand’
97	Medium	In Figure 2-31, remove the ‘Data Payload’ field of the TLP shown in the Transaction Layer sub-block
109	Minor	Table 3-1. Change ‘Skip’ to ‘Skip (SKIP)’
110	Minor	Figure 3-2. Change the 3 SKIP (K28.0) fields to SKP. Also, change ‘Local Traffic: Ordered Sets’ to ‘Local Traffic: Ordered Sets or Physical Layer Packets (PLPs)’
112	Medium	In Figure 3-3 change ‘LCRC’ to ‘CRC’ and in the associated label, delete the word ‘Link’
115	Minor	Change Figure 3-4 title to ‘PCI Express Memory Read Request and Completion TLP’

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
125	Medium	Figure 3-9, also circle the Expansion ROM Base Address register
127	Major	Replace the entire section on top of the page titled "Other Notes About ...." with the following: 1. Assume an address-routed packet moves downstream into a switch upstream port: • If the packet's address falls within the upstream port's bridge Base/Limit register range and one of the downstream port's bridge Base/Limit register, the switch will forward the packet downstream that port. • If the packet's address does not map to any switch port's bridge BAR and any port's bridge Base/Limit register range, it will be handled as an unsupported request on the primary link. 2. Assume an address-routed packet moves upstream into one of the switch's downstream ports: • If the packet's address falls outside that port's bridge Base/Limit register range as well as outside the upstream port's bridge Base/Limit register range, the packet is forwarded upstream. • If the packet's address falls outside that downstream port's bridge Base/Limit register range and within the upstream port's bridge Base/Limit register range as well as within another downstream port's bridge Base/Limit register range, the packet is forwarded to that downstream port. • If the packet's address falls within the arriving downstream port's Base/Limit register range, it will be handled as an unsupported request on the secondary link that the packet arrives on. • If the packet's address does not map to any port's bridge BAR and the packet's address falls within the upstream port's bridge Base/Limit register range and outside all other downstream port's bridge Base/Limit register range, it will be handled as an unsupported request on the secondary link that the packet arrives on.

PCI Express System Architecture Book Errata

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
131	Minor	Figure 3-13. In item (2), Change ‘Secondary-Subordinate’ to ‘Secondary and Subordinate’
157	Major	Replace the last sentence with ‘In the event of a NAK (No Acknowledge), the transmitter will re-send those TLPs that are not "implicitly ACKed" by the NAK. For further details, see Chapter 5.’
158	Minor	In Figure 4-2, replace the TLP packet field called ‘CRC’ with ‘LCRC’
164	Minor	In the last row of Table 4-4, change ‘Bit 9 = 1’ to ‘Bit 0 = 1’
165	Minor	Replace the contents of the “Field Use” field with the following: These four high-true bits map one-to-one to the bytes within the last double word of payload. Bit 7 = 1: Byte 3 in last DW is valid; otherwise not Bit 6 = 1: Byte 2 in last DW is valid; otherwise not Bit 5 = 1: Byte 1 in last DW is valid; otherwise not Bit 4 = 1: Byte 0 in last DW is valid; otherwise not See below for details on Byte Enable use.
176	Medium	First bullet. Change ‘are not used’ to ‘are not used and not present’
179	Minor	In Table 4-7, last row, change ‘ Byte 11 Bits 7:2’ to ‘ Byte 11 Bits 7:0’
192	Minor	In Table 4-10, first row, change ‘0111 111xb = Vendor Type 0’ to ‘0111 1110b = Vendor Type 0’
192	Major	In Table 4-10, 2nd to last row, change ‘ Byte 11’ to ‘ Byte 15’. Change ‘Byte 10’ to ‘Byte 14’. Change ‘Byte 9’ to ‘Byte 13’. Change ‘Byte 8’ to ‘Byte 12’.
192	Major	In Table 4-10, last row, change ‘ Byte 15 Bits 7:2’ to ‘ Byte 11 Bits 7:0’. Change ‘Byte 14’ to ‘Byte 10’. Change ‘Byte 13’ to ‘Byte 9’. Change ‘Byte 12’ to ‘Byte 8’.
194	Medium	Table 4-12. Change ‘PM_Turn_Off’ to ‘PME_Turn_Off’

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
237	Major	Towards the bottom of the page, replace the sentence “The timer resets to 0 and restarts when and ACK or NAK DLLP is scheduled” with the following: “The timer resets to 0 and restarts when all prior received TLPs have been acknowledged and a new TLP is received that is good and forwarded to the receiver Transaction Layer”
252	Minor	At end of page, add this sentence: ‘IEEE 1394 and USB devices are examples of devices that may source isochronous traffic to the PCI Express fabric’
274	Minor	Last paragraph. Add the term ‘(RCRB)’ after ‘Root Complex Register Block’
281	Minor	In Figure 6-21, replace ‘10b = PAT entry is 2-bits (4 ports)’ with ’ 01b = PAT entry is 2-bits (4 ports)’
299	Major	In the section ‘Stage 3 - The Credit Limit Count Rolls Over’ paragraph 1, 3rd line, replace ‘(CR)’ with ‘(PTLP)’
311	Medium	Just above the equation on this page, replace ‘(.4ns)’ with ‘(4ns)’
312	Major	In the section “Error Detection Timer- ...” The middle portion of the paragraph indicates that the Error Detection Timer gets reset any time a Flow Control packet of any type is received. This is not correct. A separate timer is implemented for each FC credit type (P, NP, Cpl) and each timer is reset when a UpdateFC DLLP of the corresponding type is received. Note that the timer associated with FC credit types with infinite credits must not report an error. Note: The specification states that optionally receipt of any DLLP can be used to reset the timer.
322	Minor	Table 8-2. Add to the table title the following: ‘(Incomplete table only used for discussion)’
326	Minor	Table 8-3. Add to the table title the following: ‘(Incomplete table only used for discussion)’

PCI Express System Architecture Book Errata

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
330	Major	1st bullet should say ‘ Native PCI Express device - must support MSI delivery method with 64-bit MSI capability register implementation. Must also support Legacy INTx emulation 2nd bullet should say ‘ Legacy Endpoint Device - must support MSI delivery method with 32- or 64-bit MSI capability register implementation. Must also support Legacy INTx emulation
331	Minor	Figure 9-1, Replace “INTA Assertion Message” with “INTx Assertion Message”
343	Medium	In Figure 9-7, change ‘04h = INTC# used’ to ‘04h = INTD# used’
348	Medium	Replace Figure 9-11 in the book with Figure 4 on page 19 and update the figure heading also with that of Figure 4.
349-350	Major	Replace these two pages with the section “Mapping and Collapsing INTx Messages” on page 20 of this document
339	Major	At the bottom of Figure 9-6, Byte 10, move ‘MSI Message Data’ field to Byte 0 and 1 position and move ‘0000h’ field to Byte 2 and 3 position.
363	Minor	At the end of the section ‘Data Poisoning (Optional)’ add the following paragraph: The EP bit is subject to a single data-bit error because it is not covered by ECRC (EP bit always assumed to be 1 when calculating ECRC). If EP was set to 1 to signal a problem with the data payload, but the bit changed to 0 due to a Link error, the problem would go undetected. It might be advisable to force and ECRC error too if setting EP to 1.
372	Major	In Table 10-2, Change the ‘Status Code’ for Completer Abort (CA) from ‘011b’ to ‘100b’. Change the ‘Status Code’ for Reserved from ‘100b-111b’ to ‘all others’
415	Medium	In Figure 11-12, change ‘STP’ associated with the DLLP packet to ‘SDP’
418	Major	Replace Figure 11-13 of the book with Figure 1 on page 16

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
420	Minor	Second to last line of the page, change ‘performance is degraded by 25%’ to ‘performance is degraded by 20%’
424	Major	Replace the section on ‘8b/10b Encoding Procedure’ up until the start of the ‘Example Encoding’ section with the section titled “8b/10b Encoding Procedure” on page 24 of this document
439	Medium	In Figure 11-21, show the Rx Clock from the Rx Recovery PLL connecting t the Differential Receiver also. Also, in the same figure, move the Lane De-Skew Delay Circuit position to after the Elastic Buffer.
443	Medium	At the bottom of the page is a formula used to determine the maximum number of symbols between SKIP Ordered-Sets. The constant ‘26’ should be changed to ‘28’. Also, replace the last sentence of the page with ‘28 is the number of symbols associated with the header (16 bytes), the optional ECRC (4 bytes), the LCRC (4 bytes), the sequence number (2 bytes) and framing symbols STP/END (2 bytes).
458	Major	Replace Figure 12-3 of the book with Figure 2 on page 17. Small capacitors are added on receiver side
458	Minor	Last sentence. Replace the reference to ‘Table 12-2 on page 480’ with ‘Table 12-1 on page 477’
463	Medium	Figure 12-5, Change the equation ‘V_{CM} = max $\|V_{D+}$ - $V_{D-}\|$/ 2’ to ‘V_{CM} = max $\|V_{D+}$ + $V_{D-}\|$/ 2’
464	Minor	In first and last paragraph, change ‘electrical Idle’ to ‘Electrical Idle’
490	Medium	In Figure 13-1, add an arrow that connects PERST# signal to the Root Complex
483	Medium	Figure 12-18, 4th bullet. Indicate that more realistic numbers appear to be in the 1-2us range rather than 50ns.
499	Minor	Last sentence, replace NEWCARD with ExpressCard

PCI Express System Architecture Book Errata

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
503	Medium	Replace Figure 14-3 of the book with Figure 3 on page 18
519	Major	Figure 14-7. Replace the sentence for the transition from the Polling.Active state to Detect state with 'no TS1s/TS2s received with Link and Lane number set to PAD on any Lane'
520	Minor	Figure 14-8. Change 'Lane Width' to 'Link Width'
537	Minor	Figure 14-13, on the path from Recovery.RcvrLock to Exit to Detect, it says 'TS1/TX2' when it should say 'TS1/TS2'
527	Minor	Add at the end of the figure title the following: 'for one x4 Link'
529	Minor	Add at the end of the figure title the following: 'for two x2 Link'
532	Minor	Add at the end of the figure title the following: 'for two x2 Link'
549-552	Minor	Update the headings to the following: Maximum Link Speed[3:0] (Link Capability Register[3:0]) Maximum Link Width[5:0] (Link Capability Register[9:4]) Link Speed[3:0] (Link Status Register[3:0]) Negotiated Link Width[5:0] (Link Status Register[9:4]) Training Error[0] (Link Status Register[10]) Link Training[0] (Link Status Register[11]) Link Disable[0] (Link Control Register[4]) Retrain Link[0] (Link Control Register[5]) Extended Sync[0] (Link Control Register[7])
551	Minor	Add ':' to all bullet items same as in the first bullet. Also, remove the ':' after the heading 'Link Speed[3:0];'
562	Medium	In table 15-1, change 40W to 75W
589	Medium	Replace the first bullet under the section 'The function cannot:' with: • initiate TLPs (except PME Message TLP).

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
591	Medium	Replace the two bullets under the section 'The function cannot:' with: • initiate TLPs (except PME Message TLP, if enabled and PME_TO_ACK Message TLP). • act as the target of transactions other than PCI Express configuration transactions and PME_Turn_Off Message TLP. The function's PM registers are implemented in its configuration space and software must be able to access these registers while the device is in the $D3_{hot}$ state.
591	Medium	In the last paragraph, replace: 'When programmed to exit $D3_{hot}$ and return to the D0 PM state, the function performs the equivalent of soft reset and returns to the D0 Uninitialized PM state (but Fundamental Reset is not required to be asserted)' with 'When programmed to exit $D3_{hot}$ and return to the D0 PM state, the function returns to the D0 Uninitialized PM state (but Fundamental Reset is not required to be asserted)'

PCI Express System Architecture Book Errata

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
592	Medium	Replace the entire paragraph associated with the 'D3cold State. Mandatory' with the following: Every PCI Express function enters the $D3_{Cold}$ PM state upon removal of power (Vcc) from the function. When power is restored, a Fundamental Reset must also be asserted or the device must generate an internal reset. The function then transitions from the $D3_{Cold}$ state to the D0 $_{Uninitialized}$ state. A function capable of generating a PME must maintain its PME context while in this state and when transitioning to the D0 state. Since power has been removed from the function, the function must utilize some auxiliary power source to maintain the PME context while in $D3_{Cold}$ state. When the device goes to D0 $_{Uninitialized}$ state, if capable and enabled to do so, it generates a PME message to inform the system of the wake up event. For more information on the auxiliary power source, refer to "Auxiliary Power" on page 645.
597	Minor	In Figure 16-8, change PME# to PME
600	Minor	In Figure 16-9, change the two references to PME# to PME
632	Minor	Figure 16-23, replace the word 'PCIEX Endpoint' with 'PCI-XP Endpoint' in two places
636	Minor	Figure 16-25, box 6, 5th line, add idle after electrical.
639	Minor	Figure 16-27, replace the word 'PCIEX Endpoint' with 'PCI-XP Endpoint'
647 (no page number)	Minor	Change the title from 'Optional Topics' to 'Miscellaneous Topics'
656	Medium	In Table 17-2, replace 'Hot-Plug Device Driver' with 'Hot-Plug System Driver'

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
657	Medium	In Table 17-2, under the descriptions for the 'Card Slot Power Switching Logic', the 'Power Indicator', and the 'Attention Indicator', replace 'Hot-Plug Device Driver' with 'Hot-Plug System Driver'
663	Minor	Last sentence, replace '1.0 revision' with '1.1 revision
693	Medium	Clarification in Table 18-2. The 'Required or Optional' field is in reference to motherboard providing these signals to the connector
693	Minor	Table 18-2. In the Definition column for REFCLK+/-, add the following: 'Add-in cards are not required to use REFCLK from connector'
694	Minor	Add the following sentence to the end of this page: 'See page 918-919 for a description of the Slot Clock Configuration register bit that indicates if an add-in card device uses the REFCLK provided on the PCI Express connector'
699	Minor	Add the following sentence to the end of the first paragraph: 'Chapter 8 of the PCI 2.3 specification defines what SMBus addresses are assumed to be used by add-in cards (C6h-C9h) and other usage guidelines'
699	Minor	Add the following sentence to the end of the JTAG paragraph: 'See the PCI 2.3 specification for a description of how to use the JTAG signals'
701	Minor	Replace last sentence with "At the high end, a x16 graphics card is allowed to consume up to 75W.
707	Medium	Change references to NEWCARD on this page to ExpressCard
728	Medium	In Figure 20-3, device on bottom left should be labeled as Bus3, Device 0
730	Medium	In Figure 20-4, device on bottom left should be labeled as Bus3, Device 0

PCI Express System Architecture Book Errata

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
737	Medium	In Figure 20-9, device on bottom left should be labeled as Bus3, Device 0
755-756	Major	Replace items 1) through 5) with paragraph below 1) The enumeration software changes both the Bus Number and Subordinate Bus Number register values in the secondary Root Complex's Host/PCI bridge to bus 10 (one greater than the highest-numbered bus beneath the primary Root Complex) 2) The enumeration software then starts searching on bus 10 and discovers the PCI-to-PCI bridge attached to the downstream Root Port. 3) A series of configuration writes are performed to set its bus number registers as follows. • Primary Bus Number Register = 10. • Secondary Bus Number Register = 11. • Subordinate Bus Number Register = 11. The bridge is now aware that the number of the bus directly attached to its downstream side is 11 (Subordinate Bus Number = 11). 4. The Host/PCI's Subordinate Bus Number is updated to 11. 5. A single-function Endpoint device is discovered at bus 11, device 0, function 0.
757	Medium	In Figure 21-8, device on bottom left should be labeled as Bus3, Device 0
761	Major	At the bottom of the page add the following paragraph: Root complex integrated endpoints must exist at the same level as peer virtual PCI-to-PCI bridges representing the root complex ports. In other words, the embedded endpoint within a root complex must reside on the same logical internal bus as the virtual PCI-to-PCI bridges associated with a root complex port

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
762	Major	In Figure 21-12, remove Device 0 virtual P2P bridge. The embedded endpoint Device 0 must exist at the same level as Device 1 virtual P2P bridge. i.e. The embedded endpoint resides on Bus 0 Add an optional Root Complex Event Collector on Bus 0. Embedded endpoints signal their PME Event message and errors to the associated Event Collector
764	Major	Last sentence of this page. Replace the contents in parenthesis with '(only memory write requests and message requests are allowed)'
767	Medium	Add the following sentence to the end of this page: 'PCI Express devices (Roots, Switches, Endpoints) do not support the generation or claiming of Special Cycles'
782	Medium	In Table 22-3, add another row that indicated 11h is the ID for MSI-X support Replace 11h-FFh with 12h-FFh
786	Major	In the Description field of the first table row, replace the description of bit 0 with: – **0**. IO transactions received at the downstream side of a bridge that are moving in the upstream direction are not forwarded and bridge returns completion with 'Unsupported Request' completion status In the Description field of the second table row, replace the description of bit 0 of the first bullet with: – **0**. Memory decoder is disabled and Memory transactions targeting this device results in bridge returning completion with 'Unsupported Request' completion status.
788	Minor	Table 22-9. Change 'Addedin' to 'Added in' and 'Signalled' to 'Signaled' in two places

PCI Express System Architecture Book Errata

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
833	Major	In the Description field of the first table row, replace the description of bit 0 with: – **0**. IO decoder is disabled and IO transactions targeting this device return completion with 'Unsupported Request' completion status. In the Description field of the second table row, replace the description of bit 0 with: – **0**. Memory decoder is disabled and Memory transactions targeting this device return completion with 'Unsupported Request' completion status. In the Description field of the second table row, replace the description of bit 0 of the second bullet with: – **0**. Memory transactions received at the downstream side of a bridge are not forwarded to the upstream side and bridge returns completion with 'Unsupported Request' completion status.
835	Major	Figure 22-22. remove reserved in PCI Express from the descriptions of Bit 2 and Bit 3. Also, remove the shading for these bits

Table 1: PCI Express System Architecture Book Errata, First Printing, September 2003

Page	Severity	Description
836	Major	Under description for ISA Enable Bit 2. Change Attributes from RO to RW. Replace description with 'When set, bridge only recognizes IO addresses within its assigned IO range that do not alias to an ISA range. For more details see page 582 of the MindShare PCI System Architecture book. Reset clears this bit' Also, under description for VGA Enable Bit 3. Change Attributes from RO to RW. Replace description with 'When set, bridge performs positive decode on memory accesses in the range from 000A0000h through 000BFFFFh (the address range of the VGA video frame buffer), and IO addresses associated with the VGA register set (03B0h-03BBh and 03C0h-03DFh—including addresses that alias into these two IO ranges). This bit is qualified by bits zero and one in the Command register. Reset clears this bit. For more information, refer to "Display Configuration" on page 608 of the MindShare PCI System Architecture book'
901	Major	In Figure 24-4, Change the register called 'Endpoint L0 Acceptable Latency' to 'Endpoint L0s Acceptable Latency'
938	Minor	Figure 24-26 heading. Change the heading to 'Advanced Error Uncorrectable and Correctable Error Source ID Registers'
1031	Minor	In the Note section, change 'SIG document PCI IDE Controller Specification' to 'PCI 3.0 specification as well as the T13 standard. Also add to the end of the paragraph the following: 'There is a standard for the ATA chaining/single DMA 05h interface. There is an AHCI interface being defined for Serial ATA (SATA)'
1033	Minor	Under Background, first paragraph, last sentence, change 'Roor' to 'Root'

Figure 1: Scrambler

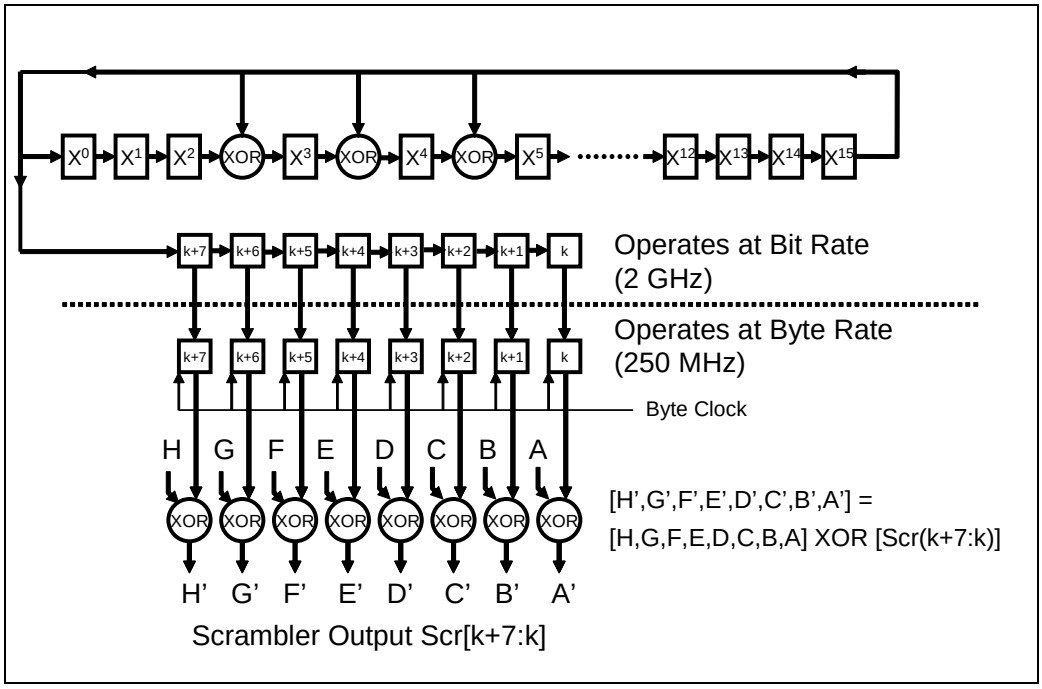

Figure 2: Receiver DC Common Mode Voltage Requirement

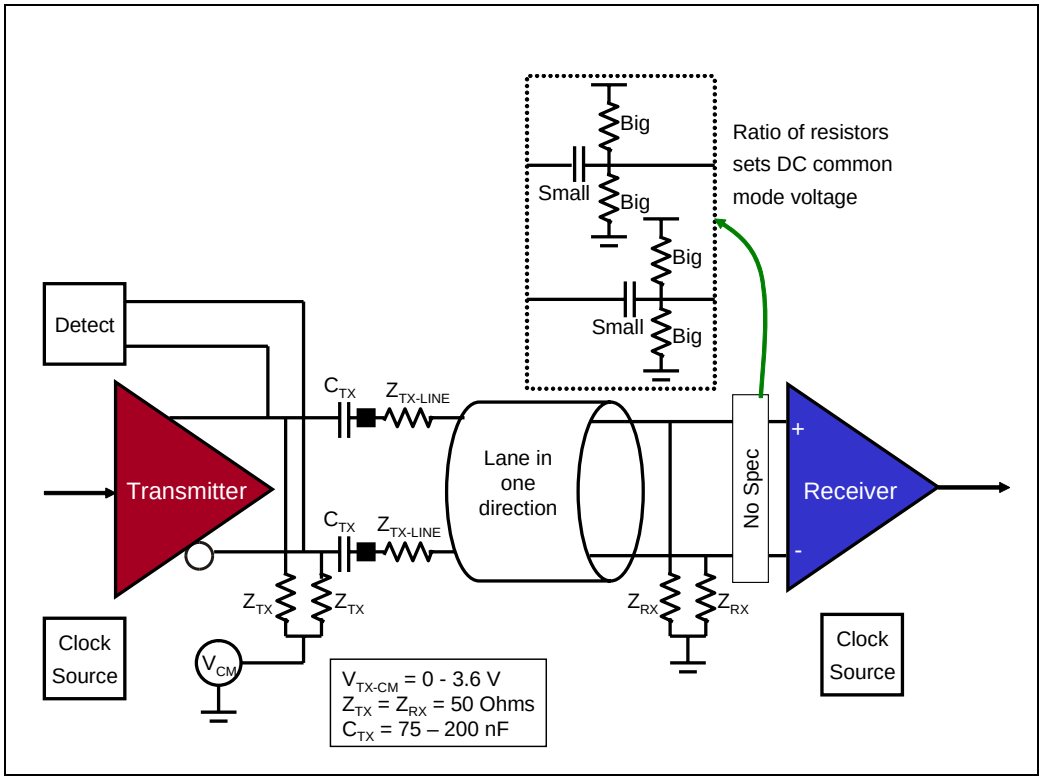

Figure 3: Example Showing Polarity Inversion

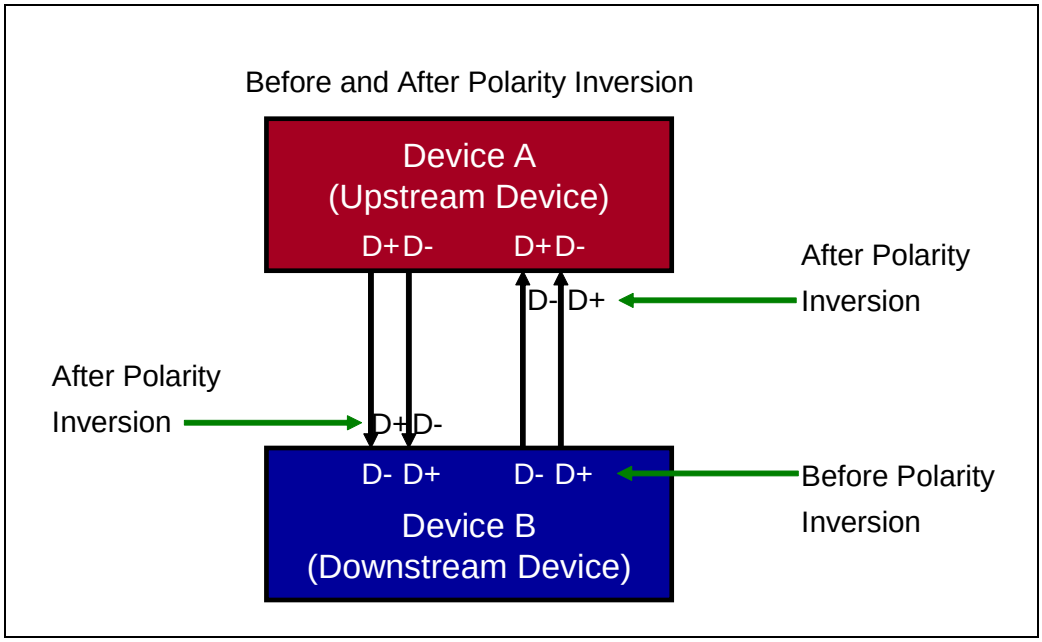

Figure 4: Legacy Devices Use INTx Messages to Virtualize INTA#-INTD# Signal Transitions

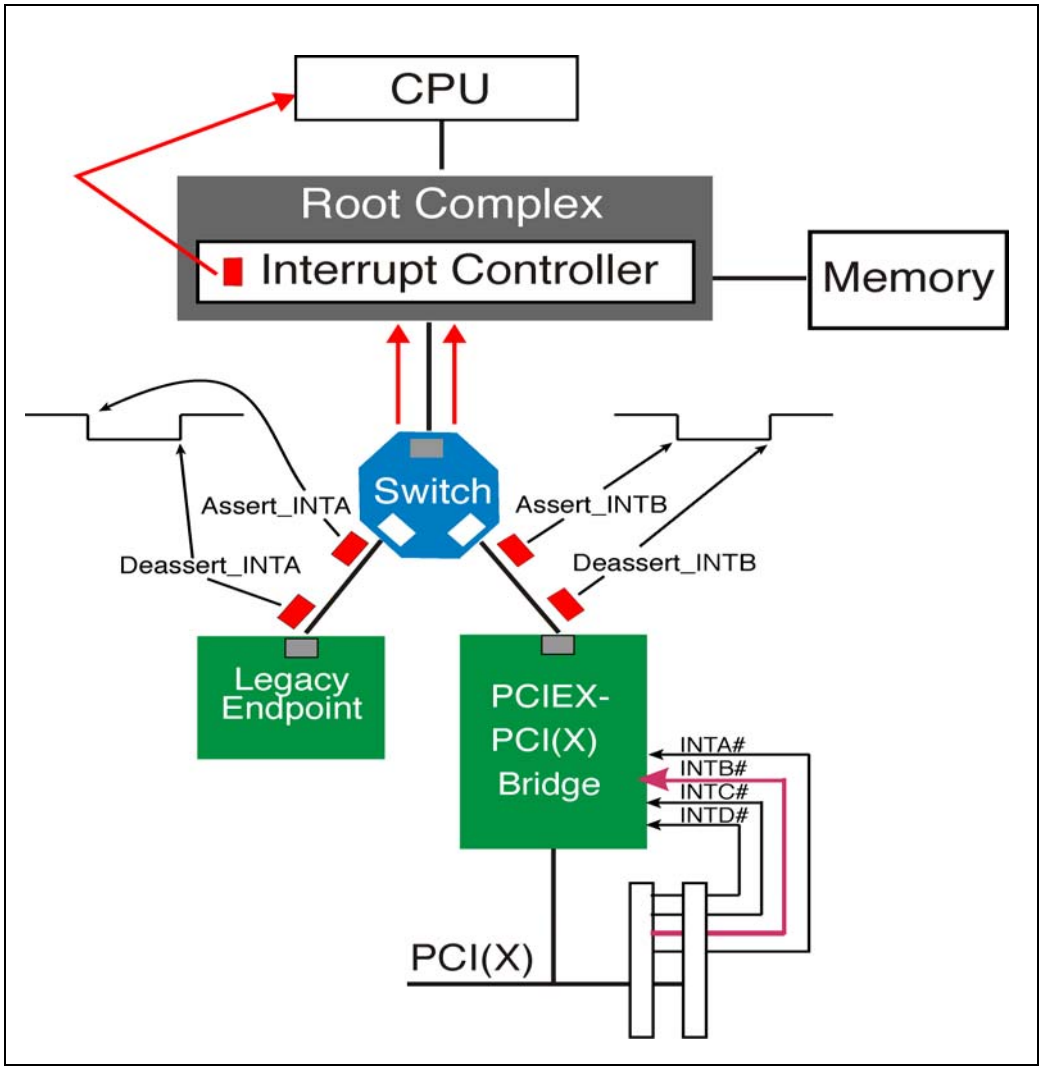

Mapping and Collapsing INTx Messages

INTx Message Mapping and Delivery

Switches that have multiple downstream ports to which legacy devices attach must ensure that INTx transactions are delivered upstream in the correct fashion. The specific requirement is to ensure that the interrupt controller receives INTx messages that represent the wire-ORed behavior and interrupt routing of legacy PCI implementations and that software can determine at which interrupt controller input an interrupt is routed. As illustrated in Figure 9-8 on page 344 (in the book), INTx lines may be shared when one or more INTx lines are tied together (wire-ORed), which results in these interrupts being routed to the same IRQ input to the interrupt controller. Consequently, when multiple devices signal interrupts that overlap only the first assertion is seen by the interrupt controller. Similarly, when one of these devices deasserts its INTx line, the line remains asserted until the last deassertion will be seen by the interrupt controller. These same principles are applied to PCI Express INTx messages.

INTx Mapping. Switches must adhere to the INTx mapping defined by the PCI specification (See Table 1-1 on page 21 in this document). This mapping defines the virtual connection that exists when interrupts are routed across a PCI-to-PCI bridge. The routing (or mapping) is based on the INTx message type and the source device number. For example, the messages being received on ports 1 and 2 (in Figure 5 on page 22) are both INTA messages. The virtual PCI-to-PCI bridge at each of the ingress ports will map both INTA message to INTA. This is because the configuration device number of both originating Endpoint devices is zero. Table 1-1 shows that interrupts messages coming from Device 0 map to the same INTx message on the other side of the bridge (i.e., internal to the bridge both INTA messages are mapped to INTA). Consequently, the egress port's PCI-to-PCI virtual bridge will also receive these interrupts as INTA messages. However, when the egress port maps the interrupts for delivery on the upstream link, the mapping changes. Assume that the virtual PCI bridge at port 1 is designated device 1 and that port 2 is designated device 2). When consulting the INTx mapping table the INTA interrupt message from Device 1 will be mapped to the INTB message and the INTA message arrive on port 2 will be mapped to the INTC message.

Table 1-1: INTx Message Mapping Across Virtual PCI-to-PCI Bridges

Device Number of Delivering INTx	INTx Message Type at Input	INTx Message Type at Output
0, 4, 8, 12 etc.	INTA	INTA#
	INTB	INTB#
	INTC	INTC#
	INTD	INTD#
1, 5, 9, 13 etc.	INTA	INTB#
	INTB	INTC#
	INTC	INTD#
	INTD	INTA#
2, 6, 10, 14 etc.	INTA	INTC#
	INTB	INTD#
	INTC	INTA#
	INTD	INTB#
3, 7, 11, 15 etc.	INTA	INTD#
	INTB	INTA#
	INTC	INTB#
	INTD	INTC#

Figure 5: Example of INTx Mapping

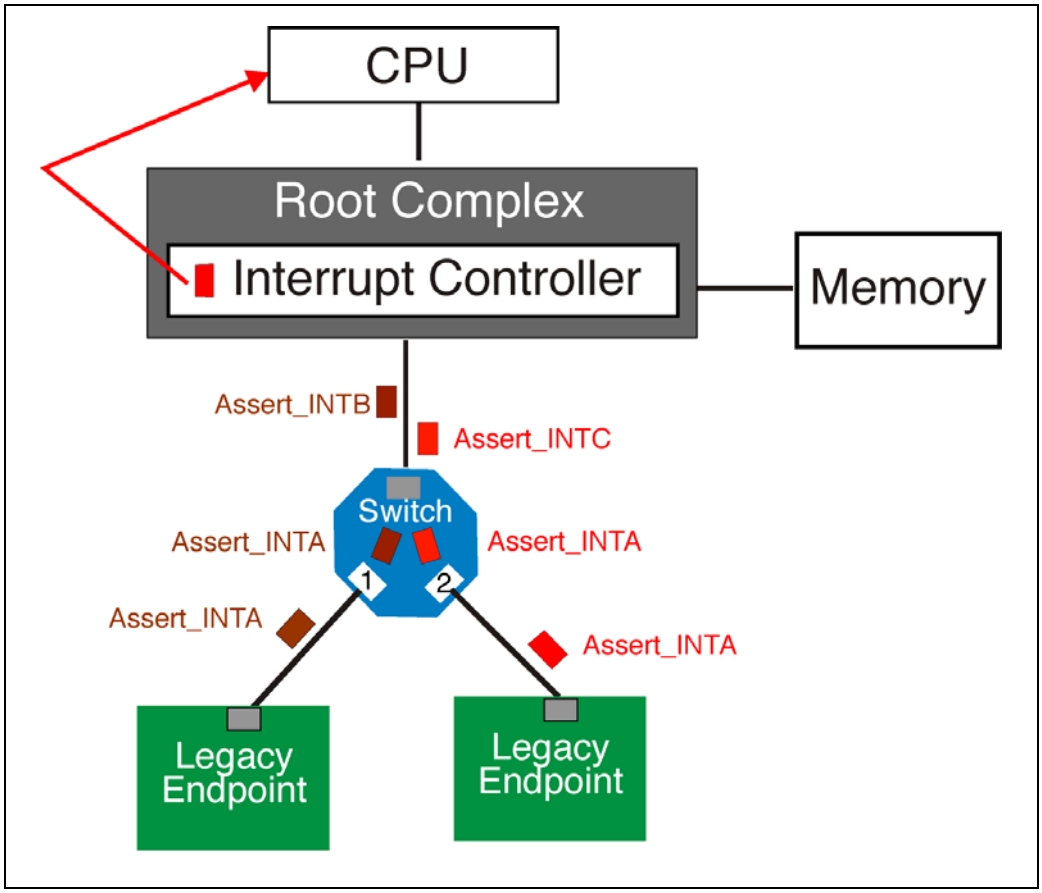

INTx Collapsing. In some cases two overlapping INTx messages may be mapped to the same INTx message by the virtual PCI bridge at the egress port, requiring the messages to be collapsed. Consider the following example illustrated in Figure 6 on page 23.

Once again the incoming interrupt message mapping does not change at the ingress port. However, when the egress port maps the interrupt messages for delivery on the upstream link, both interrupts will be mapped as INTB (assuming that the virtual PCI bridge at port 1 is designated device 1 and port 5 is designated device 5). Note that because these two overlapping messages are of the same type they must be collapsed.

Collapsing ensures that the interrupt controller will never receive two consecutive Assert_INTx messages of the same type nor two consecutive Deassert_INTX messages of the same type. This is equivalent to INTx signals being wire-ORed as illustrated in Figure 6.

Figure 6: Switch Collapses INTx Message to Achieve Wired-OR Characteristics

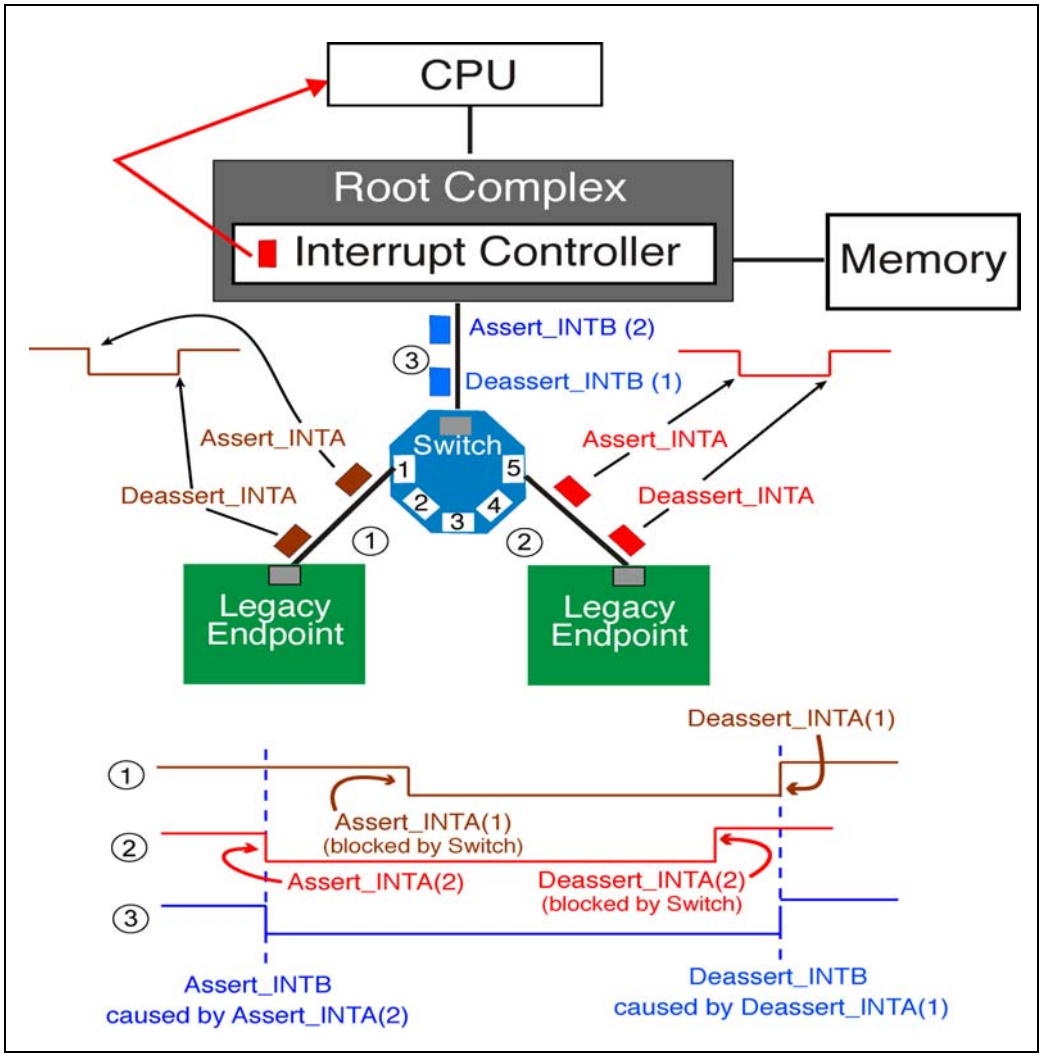

8b/10b Encoding Procedure

Refer to Figure 11-16 on page 425. There are two procedures for 8b/10b encoding.

First Procedure. The encode is accomplished by performing two table lookups in parallel (not shown separately in the illustration). The 8-bit character is split into two portions, an upper 5-bit portion (bits ABCDE) and a lower 3-bit portion (bits FGH). The first table lookup converts the upper 5-bits of the 8-bit character to the upper 6-bits of the 10-bit symbol (bits *abcdei*). The second table lookup converts the lower 3-bits of the 8-bit character to the lower 4-bits of the 10-bit symbol (bits *fghj*). :

- First Table Lookup: Three elements are submitted to a 5-bit to 6-bit table for a lookup (see Table 11-1 on page 427 and Table 11-2 on page 429):
  - The 5-bit portion of the 8-bit character (bits A through E).
  - The Data/Control (D/K#) indicator.
  - The current state of the CRD (positive or negative). This CRD value is the result of disparity calculation from the lower 4 bits of the previous 10-bit symbol (bits *fghj*).
  - The table lookup yields the upper 6-bits of the 10-bit symbol (bits *abcdei*). The disparity for this 6-bit portion is calculated and fed to the second table lookup described in the bullet below.
- Second Table Lookup: Three elements are submitted to a 3-bit to 4-bit table for a lookup (see Table 11-3 on page 429 and Table 11-4 on page 430):
  - The 3-bit portion of the 8-bit character (bits F through H).
  - The same Data/Control (D/K#) indicator.
  - The current state of the CRD (positive or negative). This CRD value is derived from the CRD calculation in the first table lookup described in the bullet above. i.e. it is the CRD value derived from the upper 6-bits of the current symbol (bits *abcdei*).
  - The table lookup yields the lower 4-bits of the 10-bit symbol (bits *fghj*). The disparity for this 4-bit portion is calculated and used for the encoding of the 5-bit portion of the next character.

Second Procedure. This procedure uses only one table lookup:

- Table Lookup: Three elements are submitted to a combined 8-bit to 10-bit table for a lookup (The 8b/10b encoding table is not shown in this textbook but is shown in Appendix B of the PCI Express Base Specification):
  - The 8-bit character to be encoded (bits A through H).
  - The Data/Control (D/K#) indicator.
  - The current state of the CRD (positive or negative). This CRD value is the result of disparity calculation from the previous 10-bit symbol (bits *abcdeifghj*) encoding.

— The table lookup yields the 10-bit symbol (bits *abcdeifghj*). The disparity for this symbol is calculated and used for the encoding of the next 8-bit character.

In both procedures described above, if the calculated CRD is neutral (i.e., it has an equal number of 1s and 0s) for the current encode, the polarity of the CRD used for the next encode is assumed to remain unchanged from the last time the CRD was either + or -.

CRD Property Used in Disparity Error Detection . If the CRD for the current encode is either + or –, the next encode results in the CRD either flipping to its opposite state or becoming neutral.

At a receiver device where the 10b symbols are converted back to 8b characters, an error is detected if two consecutive symbols are received with the same CRD state. i.e. two consecutive symbols have either + disparity or - disparity. This error is referred to as a Disparity Error.

Serializer. The 8b/10b encoder feeds a Parallel-to-Serial converter (Serializer) which clocks 10-bit symbols out in the bit order ‘abcdeifghj’ (shown in Figure 11-16 on page 425).