PEX 8505

PCI Express Switch

Silicon Revisions and Errata List

CONFIDENTIAL PROPRIETARY INFORMATION

NDA REQUIRED

Version 1.1

Oct 2011

History:

Version 0.3, July 6, 2007

- Initial publication of the Errata list
- Added Errata 1 and Errata 2

Version 0.4, August 24, 2007

- Added Errata 3
- Added Errata 4
- Added Errata 5
- Added Errata 6

Version 0.5, August 31, 2007

- Added Caution #1

Version 1.0, April 20, 2010

- Added Errata 7
- Added Errata 8
- Added Errata 9

Version 1.1, October 21, 2010

- Added Errata 10

Errata Table

Errata in AA silicon	Description	Risk Category
1.	Interval Required For Programming a Serial EEPROM With the On-Chip Serial EEPROM Controller	Low
2.	No Unsupported Request Response for Non-Posted TLP Targeted to Non-Existent Downstream Device After an Upstream Port Secondary Bus Reset	Low
3.	Weighted Round Robin Port Arbitration May not Reflect Correct Weights	Low
4.	Dropped SKIP Ordered Set when Exiting L1 Power Management State	Low
5.	SDA and SCL I/O Pins can Load Down I^2C Bus when Chip is Not Powered	Low
6.	AC-JTAG Not Supported	Low
7.	Flow Control Frequency when Extended Sync Bit is Set	Low
8.	No MSI event is generated for pending events when MSI is enabled and Master is enabled after the interrupt event	Low
9.	Blocked Non-Posted TLP Can Block Posted and Completion TLPs if Cut-Thru is Enabled	Low
10.	The PEX8505 Software Lane Status register (Port 0 register 1F4h[4:0]) provides inverted linkup status for each Lane	Low

Erratum #1 – Interval Required for Programming a Serial EEPROM with the On-Chip Serial EEPROM Controller

Risk Category: Low **Silicon Revisions Affected: AA**

Description

When programming an EEPROM through the switch EEPROM controller, there is a read to EEPROM-control-and-status register (Offset-260h[18]) to determine that associated SPI WRITE/READ command is completed. When such a read of Offset 260h[18] returns 0, there is a condition when a ChipSelectHoldTime period still exists before a follow on SPI command access can start. This ChipSelectHoldTime is determined by "Half-Period of SPI_SCLK * CSR_Offfset[268h[10:8]]", and defaults to 500 ns.

Workaround

When programming the EEPROM through switch EEP-controller, space SPI command accesses following the register read to Offset-260h[18] by an interval greater than ChipSelectHoldTime above.

Customer Impact

Programming EEPROM using the switch EEP-controller should restrict consecutive SPI-command accesses to allow for the above interval.

Erratum #2 – No Unsupported Request Response for Non-Posted TLP Targeted to Non-existent Downstream Device After an Upstream Port Secondary Bus Reset

Risk Category: Low **Silicon Revisions Affected: AA**

Description

If the PEX 8505 is strapped in a configuration where the downstream ports are vacant (no device connected) and a Secondary-Bus-Reset is applied to the upstream port of the switch (Reset on the virtual bus, upstream port register 3Ch[22]), subsequent Non-Posted Transaction Layer Packets (TLPs) targeted to the vacant downstream ports are incorrectly silently discarded. The PEX 8505 should respond to such requests with an Unsupported Request (UR). If a device is present on the downstream port of the switch then the transactions on the corresponding down stream port are handled correctly. A Hot Reset received on the upstream port of the switch, or an upstream port link down condition (DL_Down), or a downstream port Hot Reset (Bridge Control 3C[22]) generated from the downstream port of the switch do not lead to this behavior.

Workaround

Rather than apply a Secondary Bus Reset to the upstream port, instead use a Hot Reset from the upstream device, or apply Secondary Bus Resets (3Ch[22]) from the downstream ports of the switch.

Customer Impact

The Device generating a Non-Posted TLP to a Non-Existent downstream device can experience a Completion Timeout, as the switch will not provide a Completion to the Non-Posted TLP.

Erratum #3 – Weighted Round Robin Port Arbitration May Not Reflect Correct Weights

Risk Category: Low **Silicon Revisions Affected: AA**

Description

The PEX 8505 has an ingress queue and an egress queue. Port arbitration moves the TLP from the ingress queue to the egress queue based on Round Robin or Weighted Round Robin arbitration schemes. The default arbitration scheme is Round Robin. When the Weighted Round Robin scheme is enabled on a particular port, it restricts the number of TLPs in the egress queue to provide finer port arbitration. An erratum in the restriction of the TLPs prevents the correct assignment of weights to the traffic flow.

Workaround

Use the device-specific source port weight arbitration scheme to assign traffic priority.

Customer Impact

The port arbitration scheme, as defined by the PCI Express Base Specification, may not assign the correct weights to the traffic flows.

Erratum #4 – Dropped SKIP Ordered Set when Exiting L1 Power Management State

Risk Category: Low **Silicon Revisions Affected: AA**

Description

When exiting from the L1 Power Management (PM) state, there is a narrow window (recovery.idle state) during which SKIP ordered sets are not transmitted. The exit from L1 PM state steps through 3 sub-states:

1) rcvry.rcvrlock
2) rcvry.rcvrconfig
3) rcvry.idle.

For the duration in which the link is in Recovery.idle state (waiting for receipt of 8 Idle symbols) there will be no SKIP ordered sets transmitted. Note, SKIP symbols will be transmitted properly when entries into recovery sub-states are from another (L0 or L0s) state.

The Figure below illustrates the behavior.

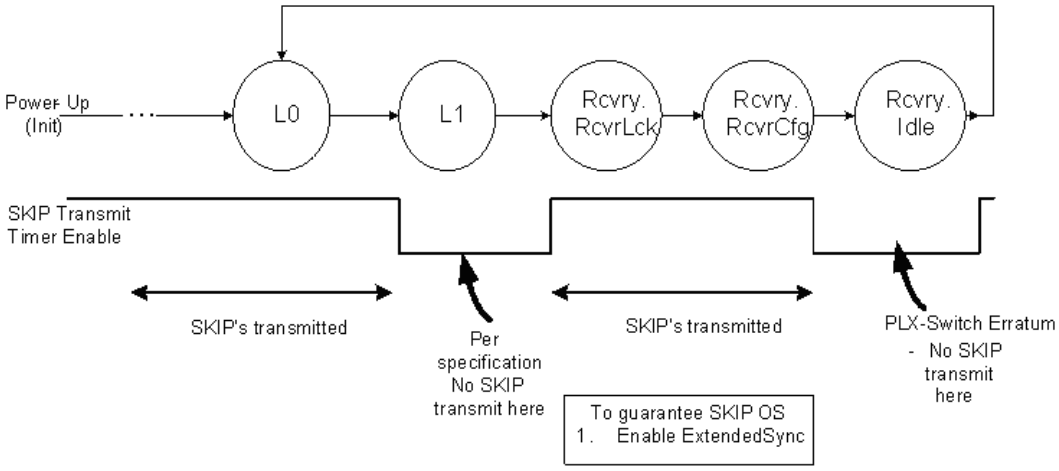

Workaround

One of the following workarounds can be used with upstream devices whose PHYs require SKIP ordered sets to exit its L1 state.

1. Disable the PEX 8505 upstream port's entry to L1 state for ASPM L1 [78h[1:0]] and PCI PM (device specific offset 1E0[7]]) mechanisms.
2. Enable the ExtendedSync mode (Offset 78h[7]), to guarantee SKIP ordered sets in the RcvrLck and RcvrCfg states (shown above)

Customer Impact

PCI Express devices whose PHY require SKIP ordered sets to deskew and also to exit the recovery path shown above could fail L1 entry/exit and reach the L0 state >100ms later; consequently, taking the link down along the way. The workarounds listed above can be used to prevent the failing case. Note that when using Workaround 1, the PEX 8505 will not realize any L1 power savings.

Erratum #5 – SDA and SCL I/O Pins can Load Down I^2C Bus when Chip is Not Powered

Risk Category: Low **Silicon Revisions Affected: AA**

Description

The I^2C specification requires that the open-drain SDA and SCL I/O pins in slave mode be floating and present no load to the wire-AND bus when power is switched off to the PEX 8505.

The PEX I/O implementation contains a clipping diode to the VDD33. When the power is off, this internal diode can turn on when the SDA or SCL line is asserted high by a master or any other slave device hanging on the bus. The resulting forward-bias current can lead to abnormal operation within the PEX 8505, resulting in system hang-up.

Workaround

At the board level, connect the SDA and SCL pins to the bus through a FET device controlled by the local PEX VDD33 supply rail. When the power is off, the two FET devices will isolate the PEX 8505 from the local I^2C bus. The SuperTex TN2501 is an example of a very good, low Vgs transistor.

Customer Impact

If the PEX 8505 is not powered up, the I^2C local bus may not function properly (i.e. it could result in system hang-up) without the two additional FET devices to isolate the PEX 8505 from the local I^2C bus.

Erratum #6 – AC-JTAG Not Supported

Risk Category: Low Silicon Revisions Affected: AA

Description

The SerDes receiver circuitry contains a circuit defect that causes the scan flip-flop to be falsely cleared just before it is to be scanned. The receiver fails to operate reliably at conditions that would normally be used for system level AC-JTAG operations. The conclusion is that the SerDes AC-JTAG model is not reliable for typical system environments and is no longer included in the BSDL file. DC-JTAG (1149.1) is unaffected.

Workaround

None. AC-JTAG has been removed from the BSDL file.

Customer Impact

Boundary-Scan instructions *EXTEST_PULSE* and *EXTEST_TRAIN* are not supported.

Erratum #7 – Flow Control Frequency when Extended Synch Bit Set

Risk Category: Low Silicon Revisions Affected: AA

Description

When the extended sync bit is set on a particular port, flow control update timers and subsequently flow control updates seen on the PCI express bus on that port violate the PCI Express specified values in the PEX 8505. The PCI Express specification requires flow control frequency with extended sync bit set to be 120us +50%, which calculates to a maximum of 180us gap between flow control update DLLPs. In the PEX 8505, however, the flow control update timer is set to 256us.

Workaround

Do not set the extended synch bit.

Customer Impact

If the extended synch bit is set, the device flow control update timer violates the specification.

Erratum #8 – No MSI Event Generated for Pending Events when MSI Enabled and Master is Enabled after Interrupt Event

Risk Category: Low **Silicon Revisions Affected: AA**

Description

This erratum is applicable only in systems that use the MSI. It is only applicable when Legacy interrupts are disabled in the system and MSI is not yet enabled (MSI Enable bit in the MSI Control register is not set or the Bus Master Enable bit in the Command register is not set). In this case, if there is an event like a Hot Plug event which can cause an MSI event to be saved, when the MSI/Bus Master is later enabled, an MSI notification of the prior event is not sent.

Workaround

There are two options:

1) When the MSI enable bit is set by the host in the system, the host should scan all pending status bits that can generate an MSI event. The same is true for Bus master enable bit.

2) Keep INTx interrupt signaling mechanism enabled until MSI enable bit is set.

Customer Impact

If MSI (rather than Legacy INTx) is to be used such (i.e. Hot Plug event notification), and MSI is enabled after Hot Plug event(s) occur (i.e. Attention Button pressed), the host will not be notified of the prior event when the host later enables MSI. Therefore, when the host enables MSI, the host should scan the interrupt status bits to determine whether a prior event occurred. Subsequent events will trigger MSI writes.

Erratum #9 - Blocked Non-Posted TLP Can Block Posted and Completion TLPs if Cut-Thru is Enabled

Risk Category: Low **Silicon Revisions Affected: AA**

Description

If a Non-Posted TLP is blocked at the switch egress port due to lack of Non-Posted flow control credits from the connected device, the Non-Posted TLP remains queued at the switch egress port. If Cut-Thru is enabled, a Non-Posted TLP that is blocked (by insufficient credits at the switch egress port) can block

pending Posted and Completion TLP(s) that are queued (within the switch) behind the blocked Non-Posted TLP.

Solution/Workaround

1. Disable Cut-Thru by clearing the Cut-Thru Enable bit (Port 0 register 1DCh[21] = 0). This will allow pending Posted and Completion TLPs to pass the stalled Non-Posted Request and prevent a possible cause of system deadlock.
2. Disable Completion Timeout in the requester to prevent Completion Timeout error, if a Completion TLP being blocked behind a Non-Posted TLP can be tolerated..
3. Ensure that any deficiency of Non-Posted flow control credits from the connected device is a transitory condition that is resolved before any resulting timeout error condition can occur

Impact

If Cut-Thru is enabled, Posted or Completion TLP(s) can be blocked behind a Non-Posted TLP that is blocked at the switch egress port due to insufficient Non-Posted credits from the connected device. In some systems, such blockage could potentially lead to a deadlock condition.

Erratum#10 - The PEX8505 Software Lane Status register (Port 0 register 1F4h[4:0]) provides inverted linkup status for each Lane

Risk Category: Low **Silicon Revisions Affected: AA**

Description

The Data Book description states that:
0 = Lane is down
1 = Lane is up

Since this register provides an inverted status value, the register description should instead state that
1 = Lane is down
0 = Lane is up

List of Cautions

Caution #1

Description

The PCI Express Base Specification 1.1 requires switch ports to forward Configuration TLPs and Messages while in D3-Hot State. Recently, the PCI SIG released an erratum to the PCI Express Base Specification 1.1 which describes the requirement for switch ports to forward Completion TLPs as well while in D3-Hot State.

Customer Impact

Many PCI Express devices do not support the forwarding of completion TLPs while in D3-Hot states since they were released prior to the PCI Express Base Specification 1.1 erratum.