



PEX 8533

Quick Start Hardware Design Guide

Version 1.0

September 2007

Website: <http://www.plxtech.com>

Email: apps@plxtech.com

Phone: 408 774-9060
800 759-3735

Fax: 408 774-2169

© 2007 PLX Technology, Inc. All rights reserved.

PLX Technology, Inc. retains the right to make changes to this product at any time, without notice. Products may have minor variations to this publication, known as errata. PLX assumes no liability whatsoever, including infringement of any patent or copyright, for sale and use of PLX products.

PLX Technology and the PLX logo are registered trademarks of PLX Technology, Inc.

Other brands and names are the property of their respective owners.

Order Number: PEX 8533-SIL-DG-P1-1.0

Printed in the USA, September 2007

Table of Contents

1. PEX 8533 Interfaces.....	1
2. PEX 8533 High-Speed Link Interface.....	3
2.1. Programmable SerDes De-Emphasis.....	3
2.2. Programmable SerDes Driver Current.....	3
2.3. Adjustable VTT Settings.....	3
2.4. Layout Considerations.....	4
2.5. Testability Considerations.....	4
3. Reference Clock Circuitry.....	4
3.1. Spread Spectrum Clocking (SSC).....	5
4. Reset Circuitry.....	5
5. EEPROM Interface.....	5
6. I2C Interface.....	6
7. Hot Plug Circuitry.....	7
8. Power Supplies, Power Decoupling, and Power Sequencing.....	9
8.1. Power Supplies.....	9
8.2. Board Level Decoupling.....	10
8.3. Power Sequencing.....	11
9. PEX 8533 Miscellaneous Signal Groups.....	12
9.1. Configuration Strapping Pins.....	12
9.2. PCI Express Lane Good Indicators.....	12
9.3. JTAG Interface Signals.....	12

Figures

Figure 1. PEX 8533 Interfaces.....	2
Figure 2. Return current path utilizing stitching vias.....	4
Figure 3. PEX_REFCLKP/N Implementation Example.....	5
Figure 6. PEX 8533 1-Kbyte SPI EEPROM Interface.....	6
Figure 8. I2C bus connectivity.....	6
Figure 9. PEX 8533 Hot Plug Controller Interface.....	8
Figure 11. Power Plane impedance vs. frequency.....	11
Figure 12. Generic JTAG interface.....	13

Introduction

This design note is intended to provide PLX Technology's PCI Express switch users with an overview of the PEX 8533 from a device interface perspective, and to provide examples of how to connect to the various interfaces.

1. PEX 8533 Interfaces

The PEX 8533 device is a 32-lane, 6-port PCI Express switch designed for high availability, high performance systems. The PEX 8533 signal interfaces can be grouped into several functional blocks. These blocks are as follows:

- PCI Express 2.5Gbps link interface signals (per port)
- Hot plug controller interface signal(s) (per port)
- Reference clock input(s)
- Fundamental Reset input
- EEPROM interface signals
- JTAG controller interface signals
- Strapping pins
- Lane status indicator signals
- Power and ground pins
- I2C Interface

Figure 1 shows the external pin interface for the PEX 8533 device. The PEX 8533 has thirty-two PCI Express lanes, which consist of both transmit and receive differential pairs (four physical signals).

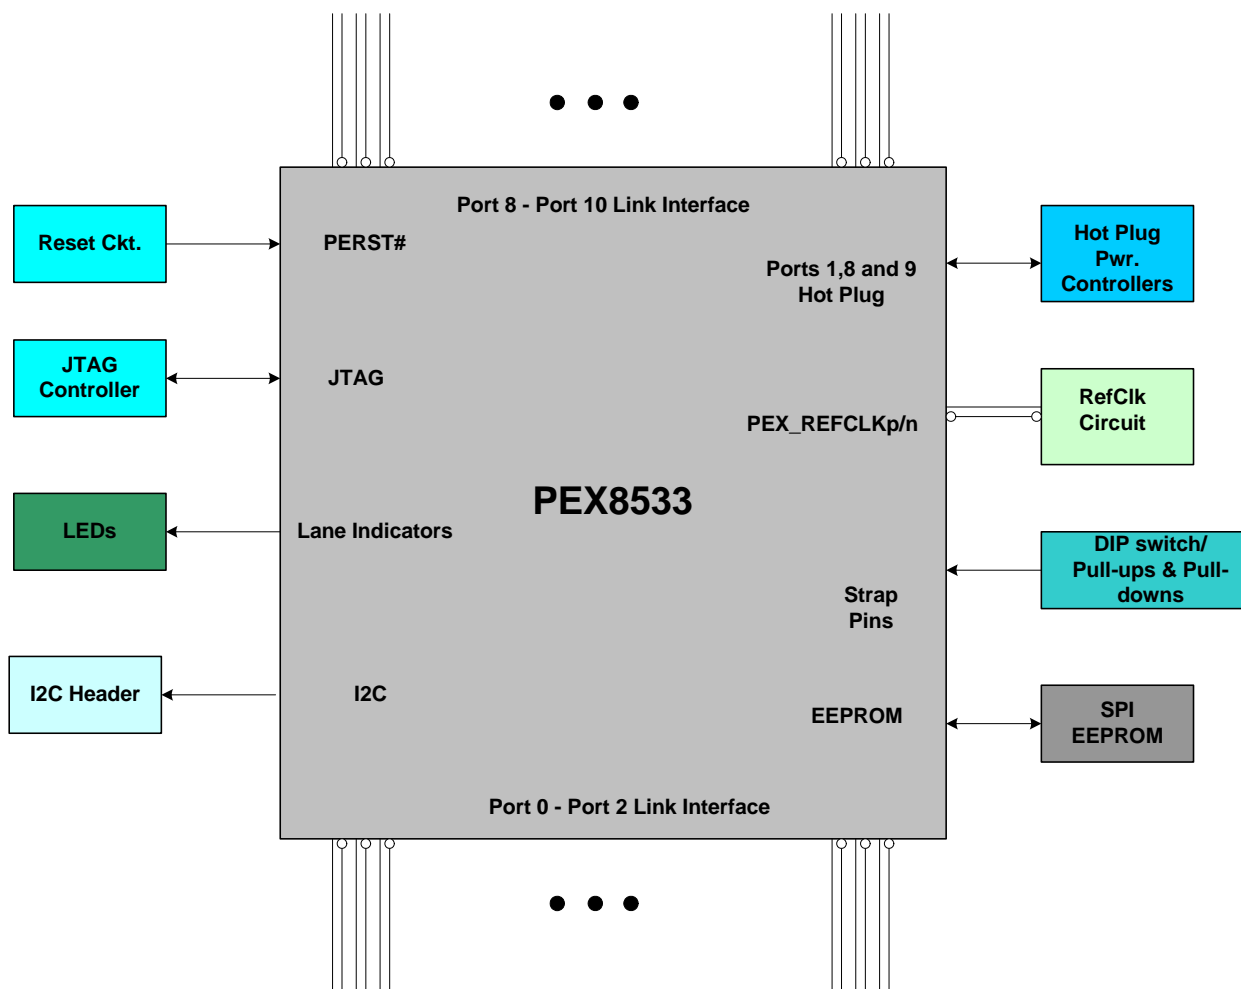


Figure 1. PEX 8533 Interfaces

The PEX 8533 allows multiple lanes to be logically grouped through strapping pins or EEPROM configuration. These lanes constitute a common signal interface between two PCI Express ports (i.e. a PCI Express link).

2. PEX 8533 High-Speed Link Interface

Each lane on the PEX 8533 is implemented using a SerDes I/O (serializer/de-serializer). The line rate of each transmitter is 2.5Gbps; therefore, the bit time or unit interval is 400ps. The rise and fall time of the transmitter is near 100ps for transition bits; therefore, the knee frequency of the signal will be around 3.5 GHz ($f_{3db} = 0.35/t_{rise}$). At these frequencies, PCB transmission lines behave like low-pass filters due to frequency dependant dielectric and conductor losses. It is important to ensure that the physical interconnect between devices produces as little signal loss and jitter as possible.

PCI Express serial links are AC coupled transmission lines. AC coupling capacitors are required on all transmitters to isolate the DC component of a signal between a transmitter and receiver. PCI Express signal transmission utilizes 8b/10b encoding; therefore, the minimum frequency content of a signal is 125MHz. In order to allow the lower frequency components of the signal to pass with minimal loss, AC coupling capacitors must be within the range of 75 nF– 200 nF. For minimizing impedance discontinuities and signal loss, 0402 or 0603-size ceramic capacitors should be used.

The PEX 8533 provides a number of ways to optimize the SerDes driver characteristics to aid the system designer in trying to meet the needs of his/her application.

2.1. Programmable SerDes De-Emphasis

Each SerDes driver contains a first order equalization function that allows for the de-emphasis of non-transition bits within a data pattern. The PEX 8533 provides four 32-bit registers (Memory addresses: BAR0+254h, BAR0+258h, BAR0+8254h, BAR0+8258h), that partition into thirty-two 4-bit fields. These fields allow the equalization ratio to be programmed through EEPROM loading or register writes. The default value coming out of reset is 8h (all 4-bit fields), which sets non-transition bit de-emphasis to - 3.35 dB. Equalization is intended to help reduce Inter Symbol Interference (ISI) and loss by providing a stronger drive for transition bits compared to non-transition bits. This helps to boost the higher frequency components of the transmitted data pattern, in order to counter the low-pass filter effects of FR4 transmission lines.

2.2. Programmable SerDes Driver Current

The PEX 8533 also provides programmable driver current. The PEX 8533 contains two 32-bit register (Memory addresses: BAR0+248h and Memory offset BAR0+8248h) partitioned into thirty-two 2-bit fields. These registers allow the driver's nominal current (I_{nom}) to be set at 10mA, 20mA or 28mA. The default value coming out of power-on reset is 0h (20mA). This default value can be overwritten via EEPROM load or through register writes.

The actual driver current (I_{tx}) is a scaled multiple of the nominal current. The PEX 8533 provides four 32-bit registers, partitioned into thirty-two 4-bit fields that allow for a programmable ratio of actual current to nominal current. These registers allow the actual current to be linearly scaled up to 135% or scaled down to 60% of the nominal current. The default value coming out of reset is 0h ($I_{tx}/I_{nom} = 1.00$). This feature can be used to reduce power consumption in applications where the physical link is short, or to boost driver current for longer transmission lines. The transmitter's output voltage swing is determined by multiplying the fixed lane termination resistance of both the transmitter and receiver (25 Ω) by the actual driver current; hence, the voltage swing can be reduced by lowering the programmable driver current ($V_{TX-DIFFp} = I_{TX} * Z_T/2$). Note that $V_{TX-DIFFp}$ is the single-ended voltage swing.

2.3. Adjustable VTT Settings

In addition to reducing the driver current, the SerDes termination supply voltage (VTT_PEX[15:0]) can range from 1.0V up to 1.8V. This supply voltage determines the driver's common mode voltage and limits the output voltage swing ($V_{TX-CM} = V_{TT} - V_{TX-DIFFp}$). The system designer should simulate their design to find out the best fit for their application.

2.4. Layout Considerations

Differential signaling reduces the amount of return current that will flow on the pair's adjacent reference plane. This helps to reduce impedance discontinuities due to non-ideal return current paths, and to reduce EMI. When laying out differential pairs, pay attention to the pair's common mode return current path. If a signal pair transitions from one signal layer to another signal layer, make sure that there is a low inductance path between the corresponding reference planes for the return current. If the signal is referencing a ground plane, and transitions through a signal via to another layer that references a different ground plane, provide a ground or 'stitching' via (connecting the ground planes) near the signal's via to allow the return current to stay close to the signals. BGA pin field ground connections and connector pin field ground connections provide ground vias for signals that change layers near those devices. Figure 3 shows an example of how return currents can utilize stitching vias to change reference planes.




Figure 2. Return current path utilizing stitching vias

If a signal is referencing a DC power plane and changes layers to reference a DC ground plane, stitching capacitors can be placed near the signal transition points to provide an AC return path from a given reference DC power plane to the reference ground plane.

2.5. Testability Considerations

System designers should consider testability issues when laying out a PCI Express interconnect. For example, add-in card and system board designers might require the ability to use a logic analyzer between two devices. Since, the two devices can be physically separated (one on the system board, the other on the add-in card), a logic analyzer can be placed in-between the two PCI Express ports. Embedded system designers, whose systems integrate both devices on the same board, might consider the addition of a Mid-bus footprint in the routing path. Keep in mind that mid-bus footprints can add additional jitter and loss; therefore, a system designer should account for these effects when assessing their jitter and loss budgets. Additional information regarding Mid-bus footprints and equipment can be found at various PCI Express test equipment manufacturers such as Tektronix, LeCroy and Agilent.

3. Reference Clock Circuitry

PEX_REFCLKp/n is a differential CML (Current Mode logic) clock inputs that receives a 100MHz (+/- 300ppm) clock source. Each receiver input has an internal biasing circuit and an integrated 110 Ω differential termination. The biasing circuit sets the receiver's common mode voltage to ~0.65V. The PEX 8533 reference clock input must be AC coupled to isolate the driver and receiver DC common

mode voltages. The recommended value for AC coupling capacitors on the reference clock input is 0.01 μ F – 0.1 μ F.




Figure 3. PEX_REFCLKP/N Implementation Example

A number of vendors provide PCI Express RefClk generators and buffers. The system designer should carefully select a synthesizer/buffer that meets the PCI Express reference clock jitter requirements.

3.1. Spread Spectrum Clocking (SSC)

Many PCI Express systems implement spread spectrum clocking in order to minimize EMI by spreading the spectral energy of the clock signal over a wide frequency band. In SSC systems, PCI Express components generally need to use a reference clock provided by the same source. This allows a transmitter PLL and receiver clock recovery function (CDR) to track the modulation frequency and stay in sync with each other. If only one side of the link uses a SSC reference clock and the other side does not, the transmitter and receiver circuits will not be able to properly track one another.

Example: If a system designer implements a PCI Express card that has a cable connection to a remote system which utilizes SSC, the card must have a means of using the reference clock from that remote system in order to operate properly.

4. Reset Circuitry

PEX_PERST# serves as a fundamental reset (cold and warm) to the PEX 8533. This signal is typically generated using a power-on reset circuit. The circuit can optionally have a manual reset capability to generate a “warm” reset. Typical PCI Express applications, which provide a system generated PERST#, keep the signal asserted for at least 100 ms after the board’s power is stable. PEX_PERST# does not contain internal debounce circuitry; therefore, the system designer should guarantee that the signal has a smooth monotonic edge while being asserted or de-asserted.

5. EEPROM Interface

The PEX 8533 has an embedded SPI EEPROM controller, which can work with a wide range of SPI Serial EEPROMs, such as the Atmel AT25XXXA and Microchip 25XX080-25XX1024 family of EEPROMs. The PEX 8533 EEPROM controller can communicate with SPI EEPROMs using 8-bit, 16-bit and 24-bit addressing. The EEPROM clock (EESK) rate defaults to 1 MHz at initialization time; however, it can be loaded dynamically by loading register 268h of port 0 (see databook for details), via EEPROM load or programming the register value after boot up.

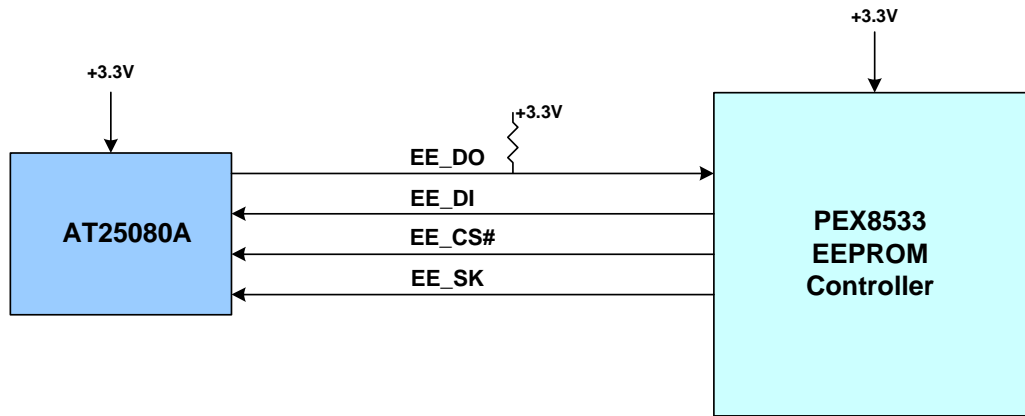


Figure 4. PEX 8533 1-Kbyte SPI EEPROM Interface

6. I2C Interface

The PEX 8533 provides an I2C compatible slave interface. This interface provides an out-of-band method for reading and writing device registers. The I2C interface is implemented as a 2-wire interface providing clock and data, and 3-pins to select the lower 3-bits of the 7-bit slave address. The upper 4-bit portion of the PEX8533's I2C address contains a fixed value of 4'b1110, therefore the 7-bit address can range in values from 58h – 5Fh. The 3-pin address bus can be connected to VDD33 via pull-up resistor, tied to ground, or connected to a switch that allows the system designer to change the address.

The PEX8533's I2C interface must be connected to a supply voltage of 3.3V. According to the I2C specification, the minimum pull-up resistor value for Standard-mode and Fast-mode is limited by the supply voltage. For example, if the supply voltage is 3.3V and there are no series resistors the value of R_{min} is 966Ω ($V_{supply} - V_{ol_max}/I_{ol_min}$). The maximum value of the pull-up resistor is dependant on the bus capacitance and rise times. For example, if a Fast-mode I2C bus has 50pF of capacitance and no series resistors, the maximum pull-up resistor value will be 6.75k Ω . Please see the I2C specification for more details.

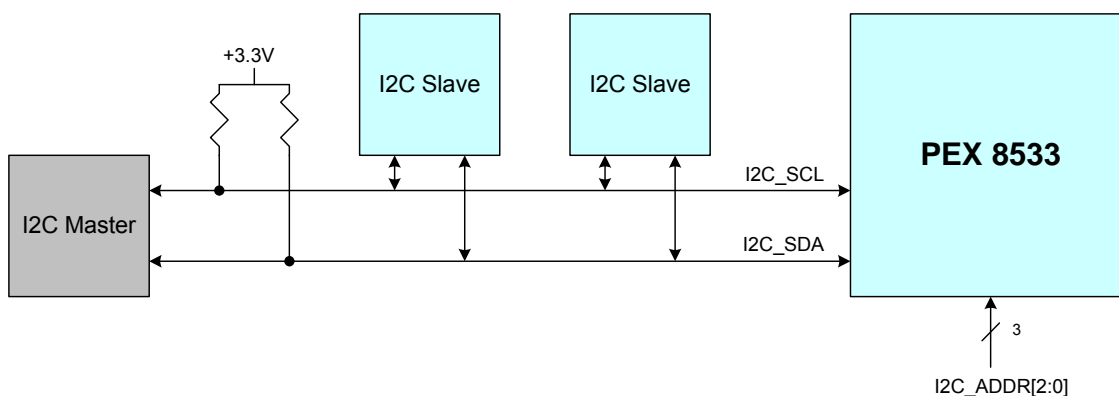


Figure 5. I2C bus connectivity

7. Hot Plug Circuitry

The PEX 8533 has the capability to be a PCI Express Hot Plug master on up to three downstream ports. Each of these three ports has nine Hot Plug pins to control various functions relating to Hot Plug. The PEX 8533 provides the following Hot Plug signals per port:

- HP_PWREN[1,8,9]#
- HP_PWRFLT[1,8,9]#
- HP_ATNLED[1,8,9]#
- HP_PWRLED[1,8,9]#
- HP_MRL[1,8,9]#
- HP_PRSNT[1,8,9]#
- HP_CLKEN[1,8,9]#
- HP_PERST[1,8,9]#
- HP_BUTTON[1,8,9]#

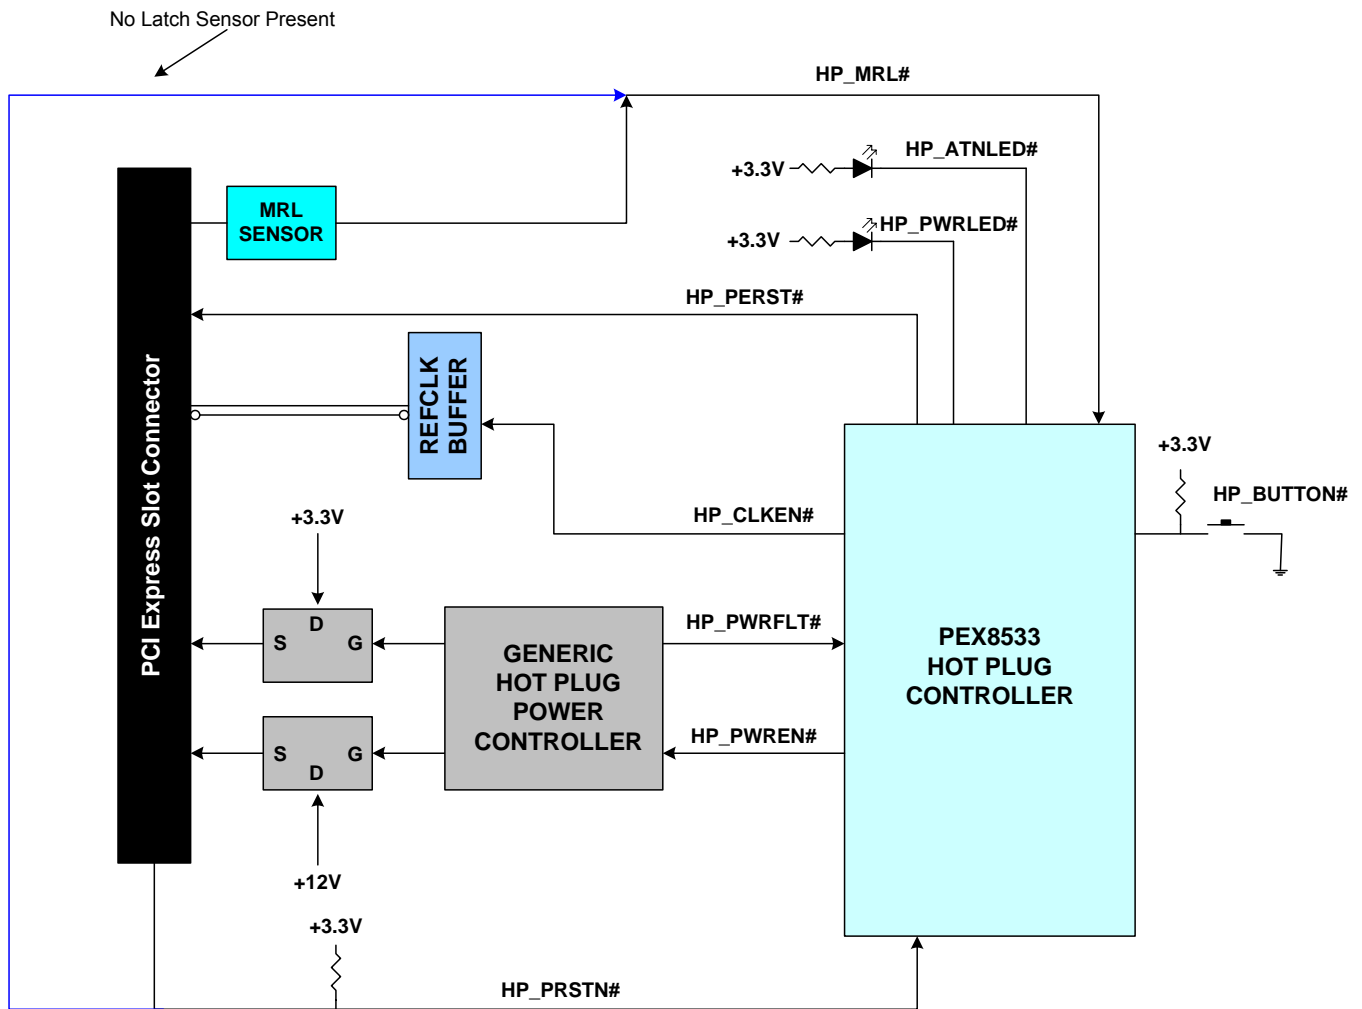


Figure 6. PEX 8533 Hot Plug Controller Interface

The PEX 8533 Hot Plug controller inputs have internal pull-up resistors on-chip; therefore, unused Hot Plug pins can either be left unconnected or connected the VDD33 via 3kohm-10kohm resistors. Numerous vendors such as Texas Instruments, Maxim Integrated Circuits and Micrel Semiconductor provide PCI Express Hot Plug power controller solutions that can easily interface to the PEX 8533 Hot Plug controller.

8. Power Supplies, Power Decoupling, and Power Sequencing

8.1. Power Supplies

The PEX 8533 has the following power pin groups:

- VDD10– Digital core logic supply
- VDD10S – SerDes digital supply
- VDD10A – SerDes analog supply used for PLLs, biasing circuitry, and PEX_REFCLK input buffer
- VDD33 – Hot Plug/EEPROM/JTAG/Lane indicator I/O buffers
- VDD33A – PEX_REFCLK PLL supply
- VTT_PEX – SerDes Transmitter termination supply

At the board level, VDD10 and VDD10S can share a common 1.0V power plane. The current demands for these supplies are relatively high, so make sure that the power plane is large enough to support the specified operating current. For best performance, the 1.0V plane should have an adjacent ground plane, which will provide an interplane capacitor to supply high frequency transient currents. Provide a sufficient number of discrete capacitors for mid and low frequency decoupling. Ceramic capacitors ranging anywhere from 0.001 μF to 0.1 μF , and multilayer ceramic capacitors (MLCC) ranging from 10-22 μF should be used for this supply rail.

VDD10A has a lower noise tolerance than the digital supplies, and therefore may require additional filtering depending on the 1.0V power plane noise. For frequency components less than 600 KHz, the VDD10A can tolerate noise levels up to 100 mVp-p. Between the frequency range 600 KHz – 320 MHz, the noise level must be less than 50 mVp-p. For frequency components above 320 MHz, the VDD10A can tolerate noise up to 100 mVp-p. Make sure that VDD10A power pins do not share vias with other 1.0V power pins going to the power plane.

VDD33 power is used for signal-ended I/O buffers, such as EEPROM, JTAG, and Hot Plug pins. Although power consumption for this supply is relatively small, the output drivers have fast edge rates and therefore, require that the system designer provide adequate power decoupling to supply transient current to the drivers. It is preferred that VDD33 be implemented as a small plane, either on a signal layer or on a main power plane layer. Provide 0.1 μF and/or 0.01 μF ceramic capacitors, along with one or more 10 μF ceramic capacitors to decouple the VDD33 power pins. The number of capacitors required depends on the number of 3.3V I/O pins utilized in the design and the existence or absence of an interplane capacitance for the VDD33 rail.

VDD33A (and VSSA_PLL) is used to power the internal reference clock PLL. This pin may require additional filtering circuitry if the VDD33 plane has significant noise. The VDD33A supply can tolerate ripple from -100 mV to +100 mV, for frequencies above 10MHz. If additional filtering circuitry is deemed necessary, a wide trace (0.010"-0.015") can be used to power this supply pin. Use a 0 Ω resistor (0603 or 0805) in series with the main VDD33 supply, along with one or more 0.1 μF and/or 0.01 μF capacitors after the resistor, close to the pin. If the VDD33 rail couples significant noise into the VDD33A supply, the resistor can be exchanged for a ferrite bead to attempt to filter additional supply noise. Note that placing ferrites in a power supply path, is not a preferred method of filtering noise for supply rails. Power supplies isolated using ferrite beads, typically have limited access to interplane capacitance, which may have adverse effects on a given supply rail. In designs where VDD33A ties directly to the VDD33 power plane, make sure that VDD33A has its own dedicated via to the plane. Similarly, allow VSSA_PLL to have its own dedicated via to the main ground plane.

VTT_PEX pins are used to set the SerDes transmitter pair's common mode voltage, and limit the output swing. Since the transmitter is implemented as a differential pair, the transient current demands required to be sourced by decoupling capacitors is minimal. Additionally, the supply rail will be less susceptible to noise if the pair is properly length matched. This supply may be implemented as a small plane. Provide at least one 10uF ceramic capacitor, and a sufficient number of 0.1 μ F and 0.01 μ F discrete capacitors for this supply rail. VTT_PEX pins must remain powered while the PEX 8533 is in the D3cold state.

8.2. Board Level Decoupling

Board-level decoupling requirements for high-speed digital designs are highly dependant on a number of different factors such as: the printed circuit board (PCB) stack-up, differential vs. signal-ended I/O signaling, driver edge rates, number of I/Os utilized, and numerous other factors. For this reason, it is not possible to present a generalized decoupling solution that will work for all designs.

Board level power supply decoupling exists primarily in two forms. The first form of decoupling is the use of discrete capacitors, and the other form is parallel plane capacitance. Parallel plane capacitance exists between a PCBs DC power and ground planes. PCB reference planes have very little series inductance; therefore, their effective frequency range is much higher than that of discrete capacitors. Low valued discrete capacitors can typically be effective for frequencies up to 250 MHz. For frequency components higher than 250 MHz, plane capacitance provides the only effective means for decoupling. Figure 8 shows the results of measuring power plane impedance versus frequency for a PCI Express Emulation board. The plot shows the bare board power to ground impedance (shown in black) compared with the impedance of various power planes after decoupling capacitors were populated. Notice that as frequencies surpass 200 MHz the impedance profile is only affected by the bare board capacitance. Having a power and ground plane separation of 0.010" will result in approximately 100 pF/in², while a separation of 0.004" will provide approximately 200 pF/in². Plane capacitors provide other important benefits such as a low impedance path for AC return currents in cases where a given reference plane has a discontinuity.

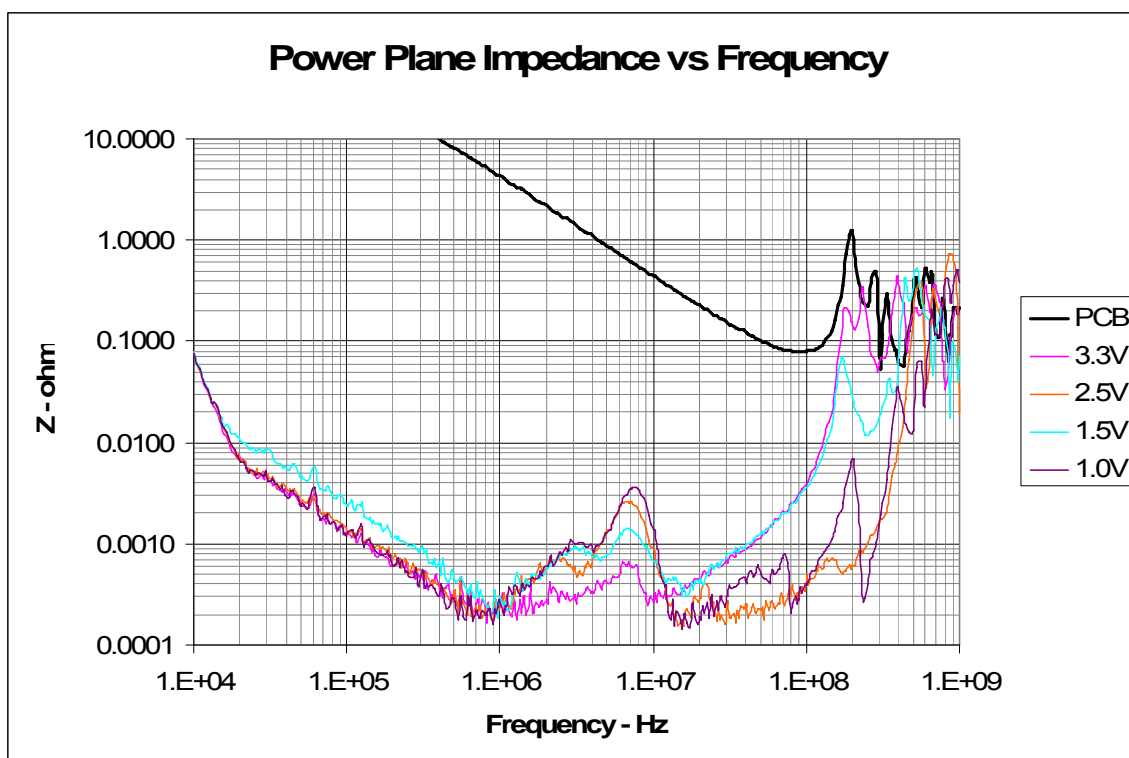


Figure 7. Power Plane impedance vs. frequency

As for discrete capacitors, the footprint and physical size of discrete capacitors has a significant effect on the frequencies where the capacitors provide effective decoupling. In order to minimize series inductance, smaller packaged ceramic capacitors (i.e. 0402 or 0603) should be used for mid-ranged frequency decoupling (20 MHz – 250 MHz). Use a mixed selection of capacitor values, such as 0.1 μ F and 0.01 μ F, in order to try to lower the impedance across a wide frequency range.

The physical layout of these capacitor footprints is important in determining the frequencies where they are effective. Avoid adding trace segments from the capacitor pads to the vias. These segments will add more series inductance, thereby lowering the LC resonant frequency of the discrete capacitor. Place the vias tangentially to the capacitor pads, and if possible, add multiple vias per pad (refer to “Right the First Time: A Practical Handbook on High Speed PCB and System Design” by Lee Ritchie). If a plane capacitor is used, then the placement of small discrete capacitors is not critical. Place the capacitors on the solder side of the board, underneath the BGA footprint (in the solder ball void area), and directly outside the BGA matrix. If a plane capacitor is not possible (this is typically the case for 4 and 6 layer boards), then place the capacitors as close to the pins as possible. If a PCB stackup is such that plane capacitors are not possible, the layout designer can add power or ground fill areas on signal layers. If a signal layer is referencing a DC ground plane, then fill with power, or if a signal layer is referencing a DC power plane, fill with ground. These copper fill areas tie to the main power and ground planes through component pins.

Multilayer ceramic chip capacitors (ex/ 10 μ F – 22 μ F) should be used for bulk decoupling of lower frequency components. The proximity of these capacitors is not critical; therefore, they can be placed outside the BGA matrix.

8.3. Power Sequencing

The PEX 8533 require three different voltage sources; 3.3V for I/O power and clock PLL power, 1.0V - 1.8V for SerDes transmitter common mode biasing, and 1.0V for SerDes/core power. VDD10, VDD10S,

and VDD10A should power up first, and power down last. If sequenced properly, all supply rails should power up within 50 ms of each other.

9. PEX 8533 Miscellaneous Signal Groups

9.1. Configuration Strapping Pins

The PEX 8533 provides a number of configuration strapping signals that allow the system designer to configure various operating modes of the switch through a hardware-only mechanism. The system designer must use external pull up resistors (value between 3 k Ω – 10 k Ω) on these lines.

Some PEX 8533 strap signals are for PLX factory use only; therefore, customers should provide a means to disable the pin functionality (pull up to VDD33). Factory use only strap signals can be pulled to VDD33 using bussed or isolated resistor networks to save board real estate.

The following PEX 8533 strapping signals are available for customer use:

- STRAP_UPSTRM_PORT_SEL[3:0] – Pull up to VDD33 or pull down to VSS depending on the desired configuration. See Data Book for details.
- STRAP_STN0_PORTCFG[1:0] – Pull up to VDD33 or pull down to VSS depending on the desired configuration. See Data Book for details.
- STRAP_STN1_PORTCFG[1:0] – Pull up to VDD33 or pull down to VSS depending on the desired configuration. See Data Book for details.
- I2C_INTA_ENABLE# - Pull up to VDD33 or pull down to VSS depending on the desired functionality. Pull up to VSS.

The following PEX 8533 signals are PLX factory use strapping signals:

- STRAP_PLL_BYPASS# - Factory use only. Pull up to VDD33.
- STRAP_PROBE_MODE# - Factory use only. Pull up to VDD33.
- STRAP_FAST_BRINGUP# - Factory use only. Pull up to VDD33.
- STRAP_DEBUG_SEL[1:0] - Factory use only. Pull up to VDD33.
- STRAP_TEST_MODE[3:0] - Factory use only. Pull up to VDD33.
- STRAP_SERDES_MODE_ENABLE# - Factory use only. Pull up to VDD33.

9.2. PCI Express Port Good Indicators

PEX 8533 provides a “port good” indicator for each PCI Express port on the device. The PEX_PORT_GOOD[10:8,2:0]# signals are used for ports 0,1,2,8,9 and 10. These signal groups can be used to drive discrete LEDs, thereby giving a visual indicator that a given port’s link is up. These output buffers can sink up to 8mA; therefore the system designer should either use low current LEDs or size the series resistor to current limit the circuit.

If the design is extremely cost sensitive and/or real estate constrained, the PEX_PORT_GOOD[10:8,2:0]# signals can be brought out the test points. The system designer can probe these signals with an oscilloscope to determine if the link is up.

9.3. JTAG Interface Signals

The PEX 8533 support a five pin JTAG Boundary scan interface. The JTAG interface consists of the following signals:

- JTAG_TCK
- JTAG_TMS
- JTAG_TDI

- JTAG_TDO
- JTAG_TRST#

At the board level, JTAG_TDI, JTAG_TMS, and JTAG_TCK should be pulled up to VDD33 with 1 k Ω -5 k Ω resistors, while JTAG_TRST# should be pulled down to VSS with a 1 k Ω -5 k Ω resistor. A 33 Ω series terminator can be added to JTAG_TDO to improve signal quality. Figure 12 shows a generic JTAG interconnection.




Figure 8. Generic JTAG interface