# PEX 8696 Quick Start Hardware Design Guide

**Version 1.0**

**March 2010**

**Website:** www.plxtech.com
**Technical Support:** www.plxtech.com/support

Document Number: PEX 8696AA-SIL-DG-P1-1.0

# Preface

## Notice

This document contains PLX Confidential and Proprietary information. The contents of this document may not be copied nor duplicated in any form, in whole or in part, without prior written consent from PLX Technology, Inc.

PLX provides the information and data included in this document for your benefit, but it is not possible to entirely verify and test all the information, in all circumstances, particularly information relating to non-PLX manufactured products. PLX makes neither warranty nor representation relating to the quality, content, or adequacy of this information. The information in this document is subject to change without notice. Although every effort has been made to ensure the accuracy of this manual, PLX shall not be liable for any errors, incidental or consequential damages in connection with the furnishing, performance, or use of this manual or examples herein. PLX assumes no responsibility for damage or loss resulting from the use of this manual, for loss or claims by third parties, which may arise through the use of the PEX 8696, or for any damage or loss caused by deletion of data as a result of malfunction or repair.

## Revision History

| Date | Version | Comments |
|---|---|---|
| August 2009 | 0.1 | Initial release. |
| March 2010 | 1.0 | Updated Figure 5 and Figure 6 |

# Contents

# Figures



# Tables

THIS PAGE INTENTIONALLY LEFT BLANK.

# Introduction

This quick start hardware design guide is an overview of PLX Technology's ExpressLane™ PEX 8696 PCI Express Switches and provides examples of how to connect to the various switch interfaces.

## 1 PCI Express Link Interface

PLX's PEX 8696 is a 96-Lane, 24-Port *PCI Express 2.0* (*that is*, Gen2)-compliant switch. *PCI Express 2.0* supports transfer rates of 5.0 GT/s per Lane. The Physical Media Attachment (PMA) Layer for each Lane is implemented as a SerDes transceiver, which is composed of a transmit path and receive path. The transmit path typically contains a serializer, Phase Lock Loop (PLL), and Current Mode Logic (CML) driver. The receive path consists of a CML Receiver buffer, Clock and Data Recovery circuit (CDR), and a de-serializer.

As the *PCI Express Base Specification, Revision 2.0*, continues to mature, so does its description of the Physical Layer Electrical sub-block. A PCI Express serial Link is described in terms of four components – Transmitter, Receiver, Channel, and Reference Clock. The Transmitter and Receiver elements are typically integrated into PCI Express silicon. The channel and Reference Clock are implemented at the system level. The PCI Express interoperability matrix implies that all four elements must support 5.0 GT/s for the Link to successfully run at 5.0 GT/s. If any one element is not 5.0 GT/s-compliant, the Link will not be able to operate beyond 2.5 GT/s. Another important concept is that 2.5 GT/s is ***not*** a subset of 5.0 GT/s. This implies that a design targeted to meet 5.0 GT/s might not successfully run in a 2.5 GT/s environment, if those design criteria are not met, as well.

Figure 1 illustrates a block diagram of a sample PCI Express Link.

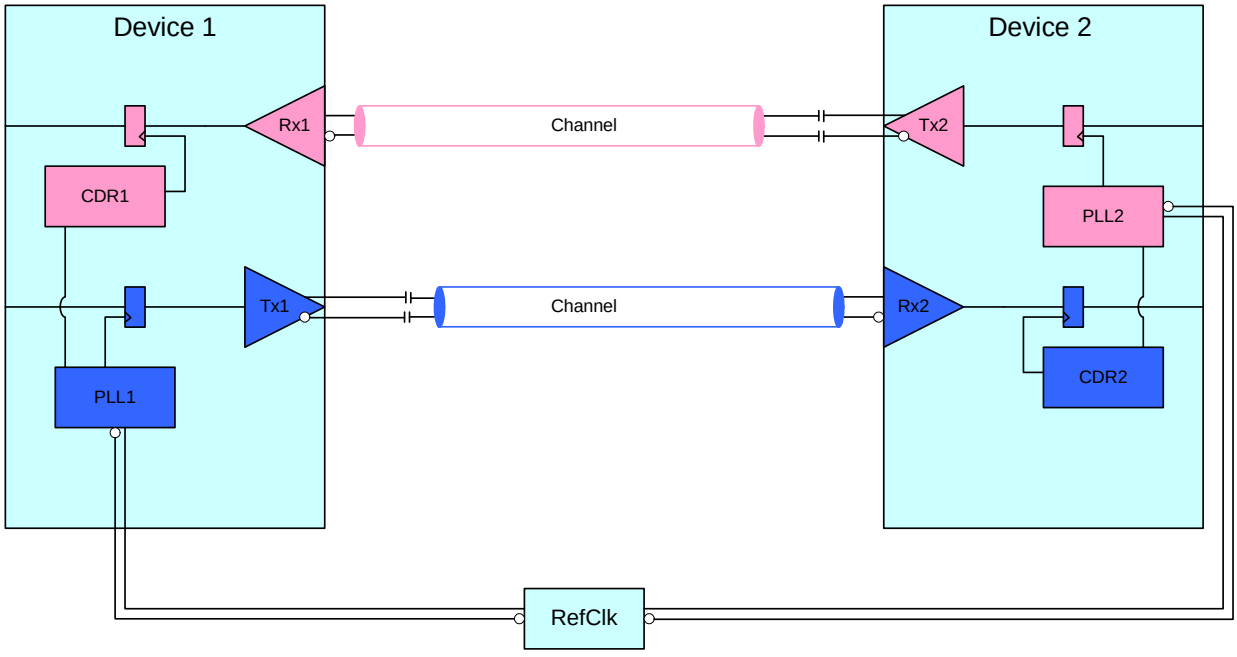


**Figure 1. Sample PCI Express Link Block Diagram**

## 1.1 Transmitter

A PCI Express Transmitter is typically a differential CML driver that transmits an 8b/10b encoded bitstream across the channel to the Receiver. The minimum differential voltage swing ($V_{TX-DIFF-PP}$) of Transmitter is 800 mV at both 2.5 GT/s and 5.0 GT/s. The DC common mode voltage can be anywhere between 0 and 3.6V; hence, AC-coupling capacitors are required to isolate the Transmitter's DC component from the Receiver's fixed 0V DC common mode voltage. The AC-coupling capacitor values must range between 75 nF and 200 nF, to ensure that the lower frequency components of the 8b/10b encoded data are not affected. Figure 2 illustrates what a generic PCI Express differential signal looks like, as compared to a single-ended signal.

***Note:*** *The swing values listed in Figure 2 (400 mV and 800 mV) do not reflect default PLX register values.*

PCI Express Transmitters are required to support de-emphasis. The role of de-emphasis is to reduce the amount of energy used to transmit multiple successive bits of the same polarity (*that is,* non-transition bits), compared to the amount of energy used to transmit a set of transition bits (0 → 1 or 1 → 0). Transition bits have higher frequency components than non-transition bits and are, therefore, more distorted by the low-pass channel. This effect is also known as *Inter-Symbol Interference* (ISI), which is a source of deterministic jitter in the system.

The *PCI Express Base Specification, Revision 2.0* defines two de-emphasis levels for devices running at 5.0 GT/s: 3.0 to 4.0 dB and 5.5 to 6.5 dB. The desired de-emphasis level for a given Link is advertised by the downstream Ports of a switch during Link recovery. The upstream port of a switch and endpoints connected to the downstream ports capture this value and Set their de-emphasis level, accordingly. Longer Links should use 6.0 dB, whereas shorter Links can use the 3.5 dB level.

The standard de-emphasis level is selectable by way of the PEX 8696 **Link Control 2** register *Selectable De-Emphasis* bit (Configuration register, offset 98h[6]).

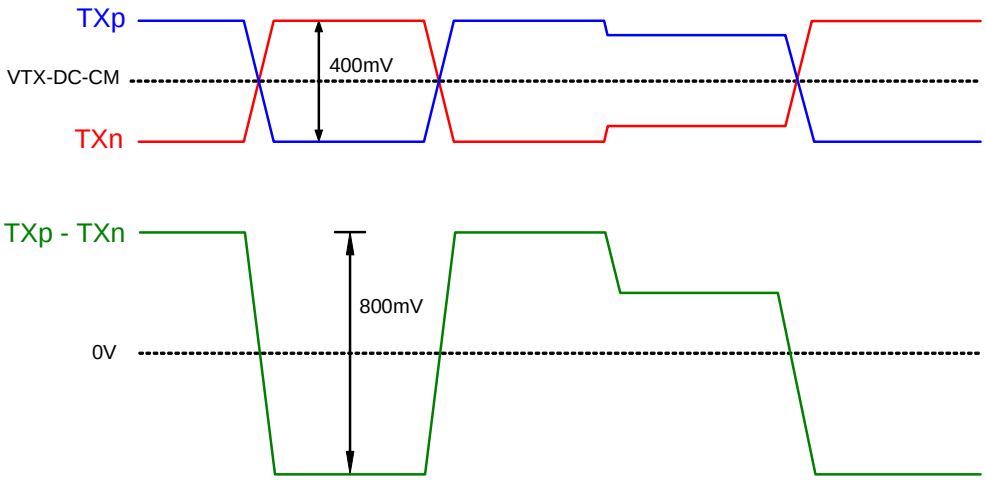


**Figure 2. Single-Ended versus Differential Voltage**

In addition to supporting the standard de-emphasis levels, the PEX 8696 has a number of programmable registers to control the Transmitter's characteristics, *such as* drive level and de-emphasis. The SerDes Transmitter Control registers exist in Station Ports 0, 4, 8, 12, 16, and 20 each controlling a bank of 16 SerDes (Lanes [0-15], [16-31], [32-47], [48- 63], 64-79], and [80-96] respectively). Registers at offsets 0xB8Ch to 0xB94h are the **SerDes Drive Level** registers. Registers at offsets 0xB98h to 0xBA0h are the **Post-Cursor Emphasis Level** registers. The **SerDes Drive Level** and **Post-Cursor Emphasis Level** registers work in conjunction, to determine the transition and non-transition bits driver swing and de-emphasis ratio. The PLX driver is implemented as a two-tap driver. When transition bits are transmitted, the **SerDes Drive Level** and **Post-Cursor Emphasis Level** register levels are added together; for non-transition bits, the two values are subtracted. Using Equation 1, Example 1 presents a calculation of what the drive level and de-emphasis level would be for a given set of register values. For more information on these and other registers, please refer to the PEX8696 Data Book.

Systems with short Links and/or power-sensitive applications (*such as* mobile platforms) can optionally decide to use low-swing output drive levels (400 $mV_{P-P}$). In the PEX 8696, this can be accomplished by setting the **SerDes Drive Level** register for a specific Lane to 01000b (400 $mV_{P-P}$), and the **Post-Cursor Emphasis Level** register to 00000b (no de-emphasis).

**Equation 1. PEX 8696 Transmitter Drive Level**

(a) $V_{TRANS} = V_{DRV\_LVL} + V_{POST\_EMP}$

(b) $V_{NON-TRANS} = V_{DRV\_LVL} - V_{POST\_EMP}$

(c) $V_{TX-DE-RATIO-3.5DB} = 20 \log (V_{POST\_EMP} / V_{DRV\_LVL})$

**Example 1. Setting for Lane 0 Transmitter to 3.5 dB**

Port 0 **SerDes Drive Level** register, offset 0xB84h[4:0] = 01111b (750 mVpp)

Port 0 **Post-Cursor Emphasis Level** register, offset 0xB94h[4:0] = 01101b (162.5 mVpp)

$V_{TRANS} = 750$ mV + 162.5 mV = 912.5 mVpp

$V_{NON-TRANS} = 750$ mV - 162.5 mV = 587.5 mVpp

$V_{TX-DE-RATIO-3.5\ DB} = 20 \log (587.5/912.5) = -3.82$ dB

## 1.2 Receiver

The Receiver's role is to recover the differential bitstream coming across the channel from the Transmitter, and latch it so it can be de-serialized and forwarded to the logical sub-block. The main components of a Receiver are the receive buffer and the CDR circuit.

The PCI Express receive buffer input threshold is 175 mV for 2.5 GT/s data rate and 120 mV for a 5.0 GT/s data rate. PCI Express Receivers are required to have a DC common mode voltage of 0V.

The receive buffer will provide bits to the CDR circuit to be sampled each and forwarded to the de-serializer. Digital-based CDRs must track the edges of the incoming bits and determine the best time to sample each bit, which is typically the center of eye (0.5 UI). The CDR base Reference Clock(s) is provided by the PLL. A CDR must be able to track either a fixed phase offset (common clock system) or small continuous phase offset (non-common clock system) between the incoming data/clock and the CDR base clock. Jitter on the base CDR clock and/or the incoming data stream can cause bit sampling errors to occur.

Although outside of the scope of the PCI Express specification, Receivers may implement some form of Receiver equalization to help compensate for the low-pass characteristics of the channel. In general, Receiver equalization only needs to be used on longer channels.

The PEX 8696 provides a programmable receive equalization function. Ports 0, 4, 8, 12, 16, and 20 each have a set of **Receiver Equalizer** registers, located at offsets 0xBA4h and 0xBA8h, to control a group of 16 SerDes. Each individual SerDes has a 4-bit control word. Table 1 describes the Receiver equalization effects.

**Table 1. Receiver Equalization Settings**

| SerDes *N* Receiver Equalizer[3:0] | Equalization |
|---|---|
| 0000b | Off |
| 0010b | Low |
| 0110b | Medium |
| 1110b | High |

## 1.3 Reference Clock

The Reference Clock is a key component to a Link that was often overlooked by system designers in first generation PCI Express systems. The Reference Clock provides a 100-MHz base frequency for the PLL. The PLL provides a frequency synthesis function, generating the higher speed clocks required to transmit data at a rate of either 2.5 GT/s or 5.0 GT/s. In designs that implement digital CDRs, the PLL output also provides the Reference Clocks to the CDR circuit; hence, jitter on the Reference Clock can affect both the Transmitter and Receiver components.

The PLL has a low-pass, filter-jitter transfer function from its reference input to the high speed output clocks; therefore, it is important to minimize the low-frequency jitter in the pass band of the PLL. Low-frequency jitter below the PLL loop bandwidth passes directly to output clocks, which, in turn, drives the Transmitter and CDR circuits. Jitter at the loop bandwidth is especially critical, given most PLLs have some amount of gain at the cut-off frequency. High-frequency jitter on the Reference Clock input above the loop bandwidth is typically attenuated, and is therefore of less concern.

The jitter transfer function of a CDR circuit is modeled as a high-pass filter. Low-frequency jitter, including Spread-Spectrum Clock (SSC) modulation, is tracked by the CDR circuit, whereas higher-frequency jitter content causes eye closure at the Receiver. The cut-off frequency of the CDR high-pass function is usually less than the cut-off frequency of the Transmitter PLL low-pass function. The pass band between these cut-off frequencies is where Reference Clock jitter causes the most problems.

In PCI Express, the cut-off frequency of the PLL is specified to be between 1.5 to 22 MHz for 2.5 GT/s and 8 to 16 MHz for 5.0 GT/s data rates. The purpose of these bandwidth ranges is to limit the difference in PLL bandwidth on the two sides of a Link. This is especially important for common clock systems, where the amount of jitter appearing at the CDR is defined by the difference function between the Tx and Rx PLLs.

Another mechanism that can increase jitter seen by a Receiver in common clocked systems is the fixed phase difference (transport delay delta) between Transmitter data at the CDR input and a Receiver's recovered clock, relative to the 100-MHz Reference Clock source. This delay should not exceed 12 ns per PCI Express specification. The delay budget includes on-chip and off-chip delays. In general terms, all Reference Clock nets in a system should be matched within 38.1 cm (15 in.). Figure 3 illustrates the Reference Clock transport delay delta.

The PEX 8696 PEX_REFCLKn/p signal is the Reference Clock Input buffer. It has an internal DC-biasing circuit, and hence, should be AC-coupled from the RefClk source driver. Use 0.01 to 0.1 µF capacitors (0603 or 0402-size) to AC-couple the Reference Clock input, as illustrated in Figure 4.

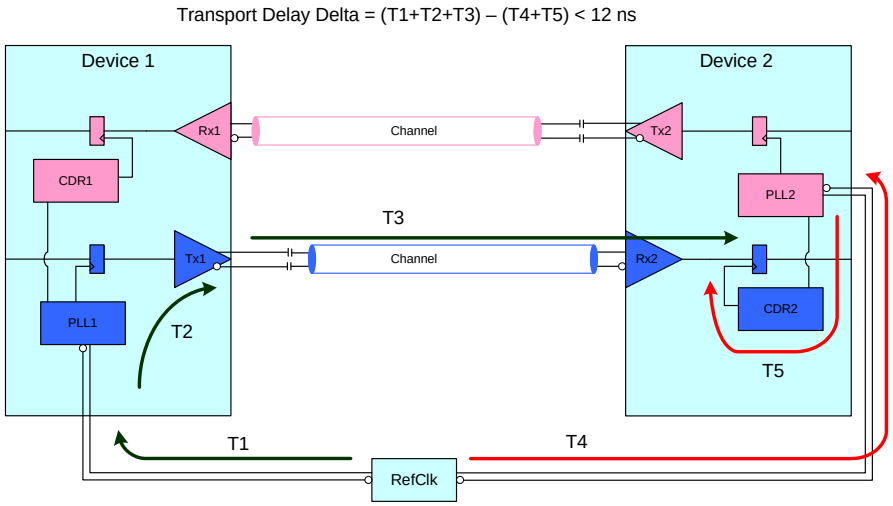


**Figure 3. Transport Delay Delta**

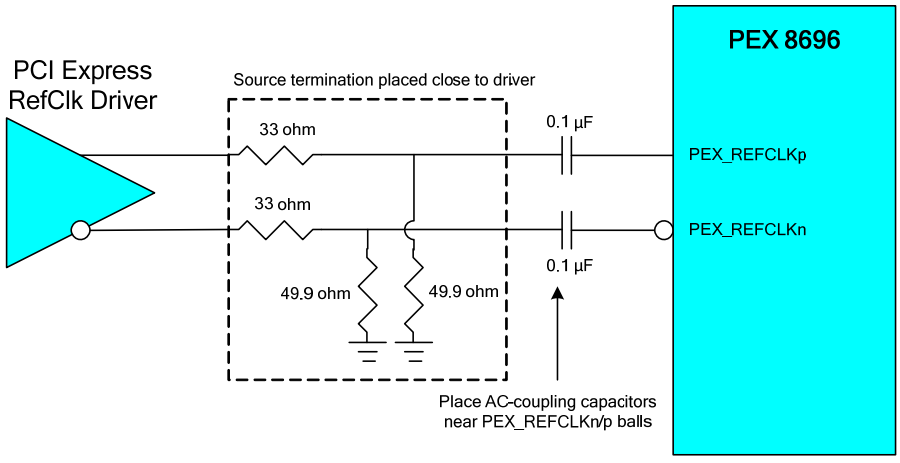


**Figure 4. PEX 8696 RefClk Circuit**

## 1.4 Channel

In PCI Express, the channel refers to the board level copper interconnects (including connectors) that lie between the Transmitter and Receiver balls. The channel is represented as a transmission line, which can be modeled by a distributed series of Resistance Inductance Conductance Capacitance (RLGC) circuits. A transmission line behaves like a low-pass filter due to frequency-dependent dielectric and conductor losses.

In PCI Express, the channel contributes to amplitude loss and deterministic jitter, which is why it is important to minimize discontinuities, *such as* vias and stubs, to minimize channel effects.

A common issue that presents itself to PCI Express system designers is determining allowable channel length. This is a question that does not have a simple answer. The best way to determine if a particular channel length is allowable is to simulate the channel using the HSPICE models for the PEX8696 provided by PLX and available on the product website at www.plxtech.com/8696. The *PCI Express Base Specification, Revision 2.0* provides additional details for simulating a channel.

# 2 PCB Layout and Stackup Considerations

PCB layout is of critical importance for PCI Express systems. Numerous form factor specifications (*PCI Express Base Specification, Revision 2.0* and *PCI Express Card Electromechanical (CEM) Specification, Revisions 1.0a and 1.1*) exist for providing important implementation guidelines for a given form factor. It is important to understand the type of system being designed before starting layout. *For example*, the *PCI Express Card Electromechanical (CEM) Specification* defines two platforms, referred to as *system boards* and *add-in cards (boards)*. Each platform has its own criteria, in terms of jitter and loss budget, trace lengths and length matching, and so forth.

## 2.1 PEX 8696 BGA Routing Escape and De-Coupling Capacitor Placement

The PEX 8696 is in a 35x35 $mm^2$ FCBGA package with 1-mm ball pitch. Power and ground pads have small “dog-bone” nets from the pad to a via which will connect it with an internal power or ground plane. The PEX 8696 places all Transmitter differential pairs on the outer two rows of balls and Receiver differential pairs on rows four and five. This means only two signal layers are required in a PCB stackup to escape the differential pairs from the BGA. All Transmitters can escape on the top layer, whereas the Receiver pairs can escape on either the bottom layer or some other internal signal layer. The positive and negative conductors of a pair should be coupled together as quickly as possible, after escaping from the BGA.

Each pair is split between two rows on the package; hence, the pairs start off with a 1-mm (39.4-mil) offset And small serpentines may be necessary to match the lengths within the pair. When implemented, make the serpentines as close to the BGA as possible to allow the differential signal to be tightly coupled as it travels down the channel.

Figure 5 and Figure 6 demonstrate one means of escaping the differential pairs from a typical PLX PCIe switch using two routing layers.

The PEX 8696 is a full matrix, 1-mm pitch BGA. Hence, placing de-coupling capacitors underneath the BGA can be tricky. It is best to use 0201-sized ceramic capacitors under the BGA matrix (bottom layer), so that the capacitors can be placed as close to the power balls as possible. Figure 6 illustrates the typical placement of 0201 de-coupling capacitors underneath the PLX PCIe switch.

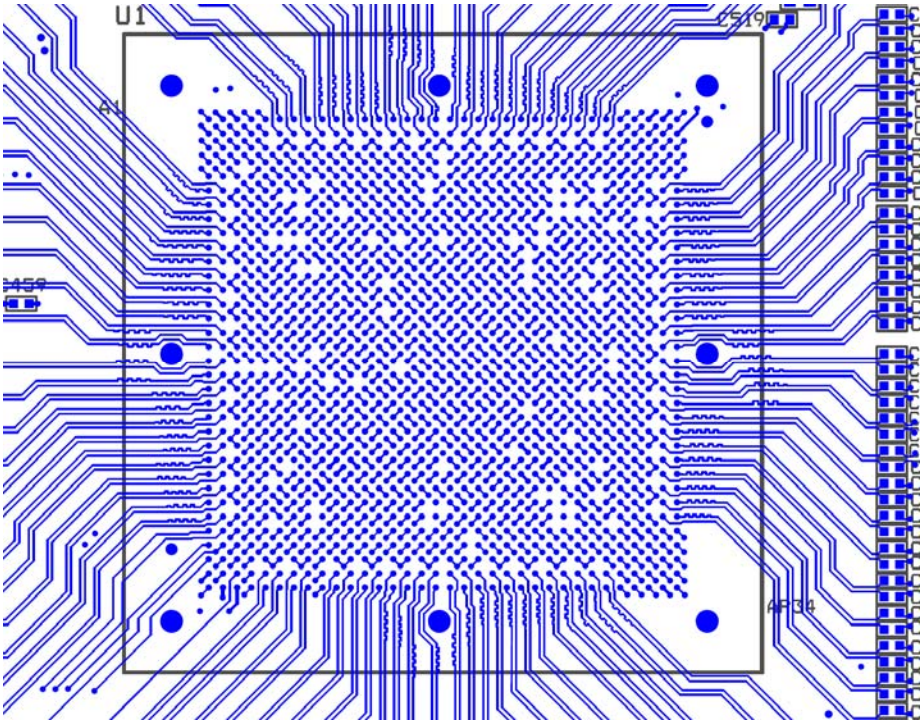


**Figure 5. Top Layer BGA Layout and Routing Escape**

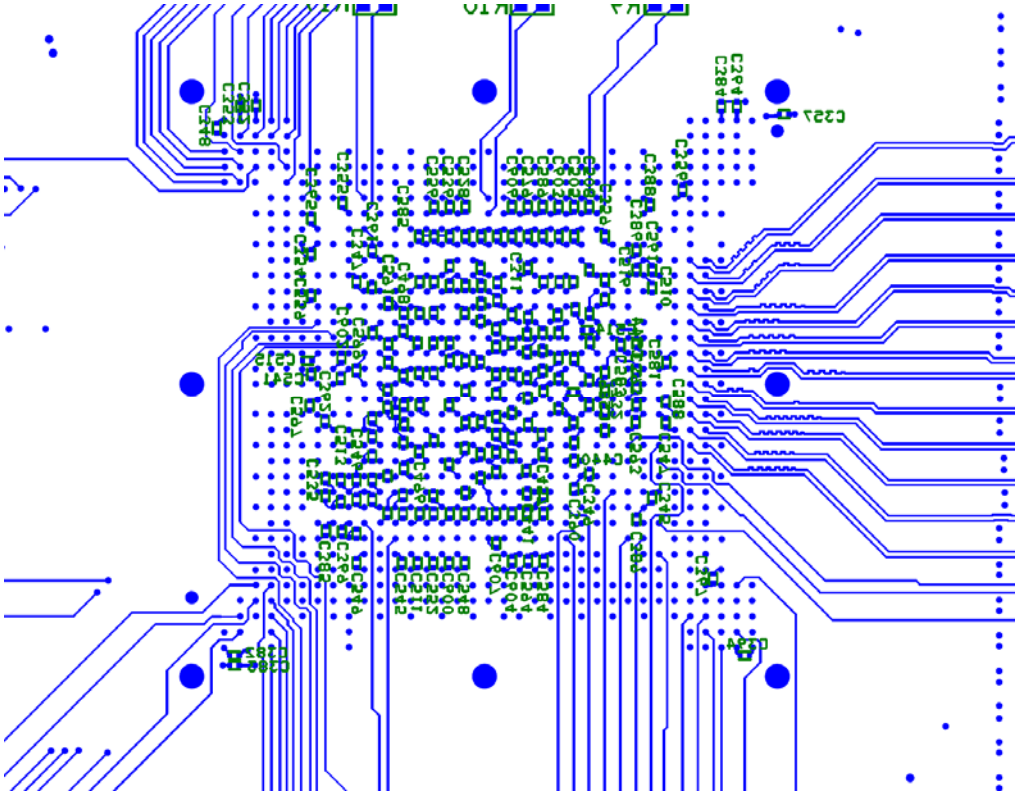

**Figure 6. Bottom Layer BGA Layout, Escape, and De-coupling Capacitor Placement**

## 2.2 Add-in Board Routing

Although the PEX 8696 Transmitter pairs escape on the top layer (as previously mentioned), at some point they must route to the bottom layer to connect to the gold fingers as is the case for an add-in board. If a logic analyzer midbus footprint is placed in the routing path, the layer transition can occur at that point. This works out well, because the midbus footprint will have a significant number of ground vias, which provide effective ground plane stitching for the differential signal's return path. If a midbus footprint is not used, layer changing can occur at the AC-coupling capacitors. Dedicated ground vias can be placed near the capacitors, close to the signal vias, to provide a return path. One ground via per pair is ideal; however, one via per every two pair is acceptable.

Receiver differential pairs must also transition signal layers at least once. Receiver pairs start off at the top layer, from the gold fingers, and into the inner rows of the BGA. The layer transition should occur at the midbus footprint, if one exists, or close to the gold fingers. Either location should have plenty of ground vias.

PCI Express add-in boards must be length-matched within 5-mil. AC-coupling capacitors should be placed close to the gold fingers. Differential pairs for PCI Express Gen2 add-in boards should have a differential impedance of between 68 to 105 ohms (85 ohms, nominal).

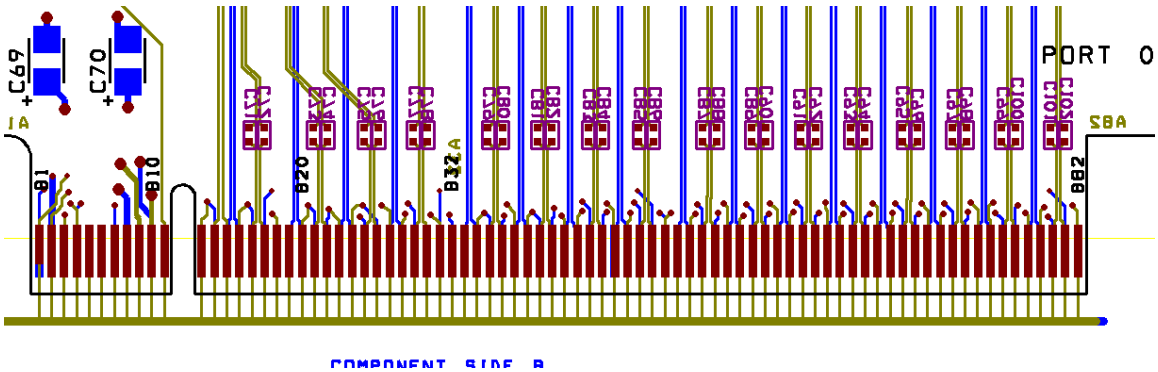


**Figure 7. Add-In Card Routing to PCI Express Gold Fingers**

## 2.3 System Board Routing

System board routing is simplified slightly. Transmitter pairs can escape on the top layer from the BGA and route to AC-coupling capacitors, placed close to the slot. Similarly, Receiver pairs can escape on the bottom layer and directly route to the slot. Transmitter pairs can also transition to the bottom layer after the AC-coupling capacitors, to minimize the stub effects of a through-hole PCI Express slot.

PCI Express system boards must be length-matched within 10-mil and all AC-coupling capacitors should be placed close to the slot. Differential pairs for PCI Express Gen2 system boards should have a differential impedance between 68 to 105 ohms (85 ohms, nominal).

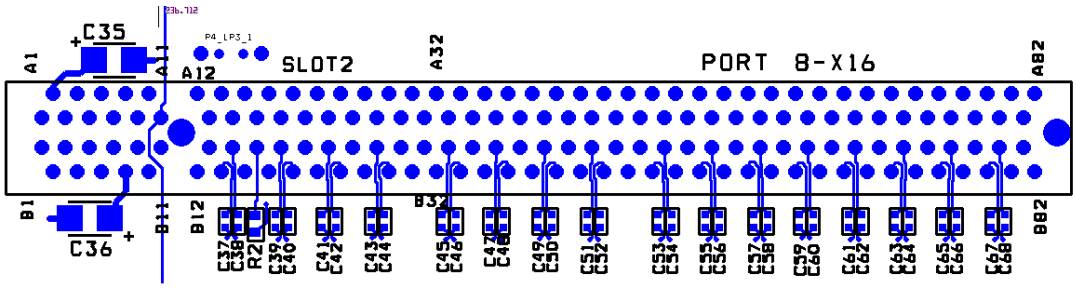


**Figure 8. System Board Routing to PCI Express Slot**

## 2.4 Midbus Routing

Midbus footprints can be placed into the routing path, to provide an interface to various protocol analyzers, as well as provide a location to probe a signal using oscilloscopes. Transmitter pairs route on one side of the footprint, while Receiver signals route through the other side.

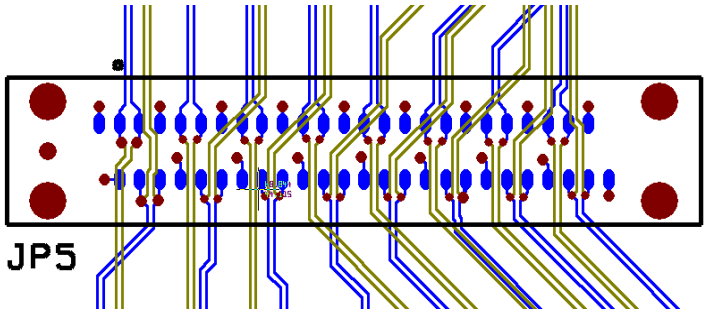


**Figure 9. PCI Express Midbus Routing Example**

## 2.5 PCB Stackup Considerations

Determining the PCB stackup is one of the most important steps in designing and implementing a system. The PCB stackup should be determined prior to board routing, because it will determine the trace width and spacing requirements necessary to achieve a particular characteristic impedance and differential impedance. After the stackup is known, the trace width can be selected. For a single-ended signal, this is enough to determine the characteristic impedance of that trace. For differential signals, the last step is to determine the separation between the positive and negative conductors, to achieve the needed differential impedance.

Additionally, a PCB stackup can determine the power supply de-coupling scheme for a device. Parallel plane capacitance exists between a PCB's DC power and ground planes. PCB reference planes have an insignificant amount of series inductance; therefore, their effective frequency range is much higher than that of discrete capacitors.

PCB traces can be implemented as one of two types of transmission lines – microstrip and stripline. Microstrip traces have only one reference plane, and therefore, represent traces on the outer layers (top and bottom layer) of a PCB. Stripline traces have two reference planes and are implemented using inner routing layers. Typically, stripline traces are only available for PCBs with six or more layers. Microstrip and stripline traces each have their own properties, which must be weighed when determining which type of trace to use.

# 3 Non-Transparent Function

The PEX 8696 supports Non-Transparent (NT) function. There are three ways to enable the NT function and configure the NT Port for the PEX 8696.

**Method 1.** Use of the Strapping balls:

- STRAP_NT_ENABLE#
- STRAP_NT_UPSTREAM_PORTSEL[4:0]

Pull down the STRAP_NT_ENABLE# to logic zero (0) to enable the NT function. Pull up or down the STRAP_NT_UPSTREAM_PORTSEL[4:0] to select the NT Port.

**Method 2.** Enable the NT function and configure the NT Port through the serial EEPROM. The NT configuration settings will be loaded upon power-up and after reset.

**Method 3.** Use the PEX 8696 $I^2C$ Port 0 to enable the NT function and configure the NT Port.

Figure 10 illustrates how to implement the NT functions through the Strapping balls. Figure 11 illustrates how to disable the NT functions, through the PEX 8696's NT Strapping balls.

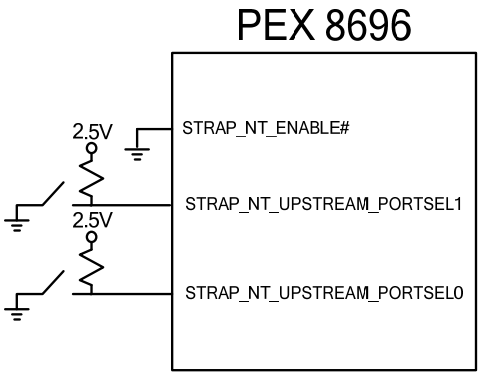


**Figure 10. Enable NT Function with NT Strapping Balls**

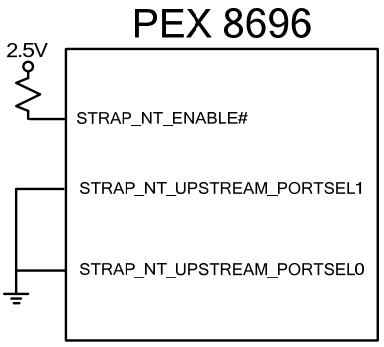


**Figure 11. Disable NT Function**

# 4 I²C Interface

The PEX 8696 implements dual two-wire I²C interface ports. I²C-Port 0 is a slave interface port and I²C-Port 1 is a master interface port used with the SHPC (see Section 5). Through its I2C_SCL0 and I2C_SDA0 balls (I²C-Port 0), the PEX 8696 allows an external I²C Master to read and write device registers through an out-of-band mechanism. The simplest way to implement an I²C interface to the PEX 8696 is illustrated in Figure 12.

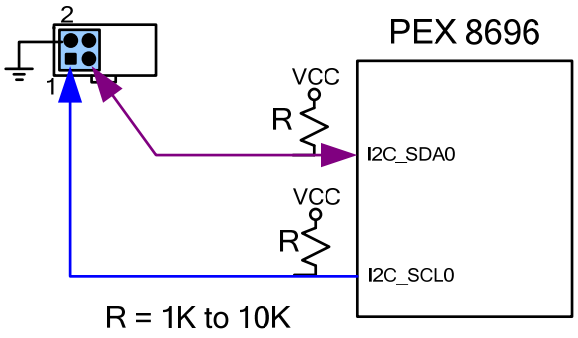


**Figure 12. I²C Interface Block Diagram**

# 5 Hot Plug Circuitry

The PEX 8696 supports four Parallel Hot Plug Controllers (PHPC) and up to twenty-four Serial Hot Plug Controllers (SHPC) to service its downstream Ports. The PHPCs are designated as Hot Plug Ports A, B, C, and D. Each PHPC has 10 Hot Plug signal balls to control various Hot Plug-related functions:

- HP_ATNLED_[n]#
- HP_BUTTON_[n]#
- HP_CLKEN_[n]#
- HP_MRL_[n]#
- HP_PERST_[n]#
- HP_PRSNT_[n]#
- HP_PWER_GOOD_[n]
- HP_PWREN_[n]
- HP_PWRFLT_[n]#
- HP_PWRLED_[n]#

*where "n" is A, B, C and/or D*

Figure 13 provides an example of how to connect the PEX 8696's PHPC to the external circuit, to build a complete PHPC circuit.

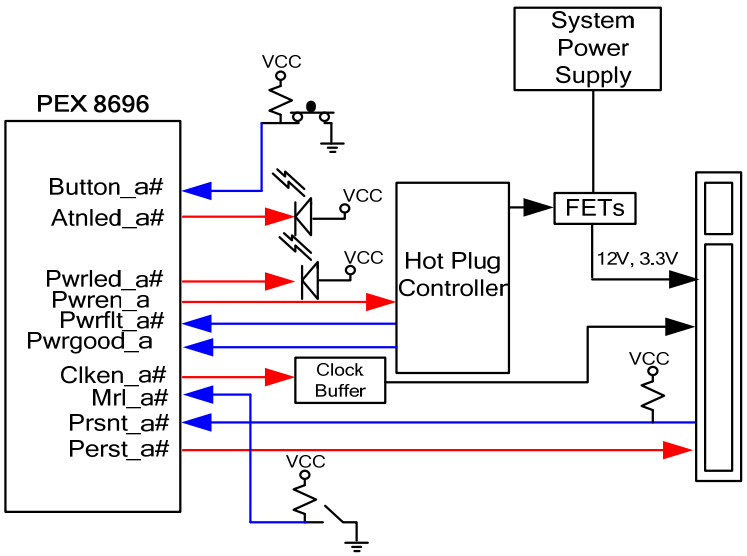


**Figure 13. PHPC Circuit Block Diagram**

When connecting the $I^2C$-Port 1 to multiple I/O expander ICs, the PEX 8696 has the option of having Hot Plug capability on eight of its downstream Ports. The PEX 8696's $I^2C$-Port 1 is the $I^2C$ Master, which is designed to interface to I/O expander ICs, to build SHPCs. One 16-I/O expander connects to $I^2C$-Port 1 for a single SHPC, and one 40 I/O expander connects to the $I^2C$-Port 1 for two SHPCs. Both 16-I/O expander(s) and 40-I/O expander(s) cannot concurrently connect to the $I^2C$ Bus. To use 40-I/O expander(s), a register bit within the PEX 8696 must be Set, and boot with serial EEPROM is essential. After the PEX 8696 is powered up, the state machine inside the PEX 8696 scans the number of I/O expander ICs connecting to the $I^2C$ Bus, starting from Address 000h, in ascending order. If it cannot locate the device with Address 000h, it stops the scan process. After it locates the I/O expander IC, it automatically assigns a valid Port Number for this SHPC. Figure 14 illustrates a block diagram of the SHPC interface to the PEX 8696. The SHPC has more signals than the PHPC. Besides the 10 Hot Plug

signals mentioned at PHPC, INTERLOCK, SLOTID[3:0], and one GPIO are added to the SHPC. Also, the interrupt signal output, INT#, from the I/O expander, should be connected to the PEX 8696's Interrupt Input ball, SHPC_INT#, for the PEX 8696 to service input events at the SHPC. Because the I/O Expander requires some time to be configured upon power-up, it is recommended that GPIO[24-42] signals from the PEX 8696 be used for slot PERST# from the I/O Expander to provide reset to downstream slots.

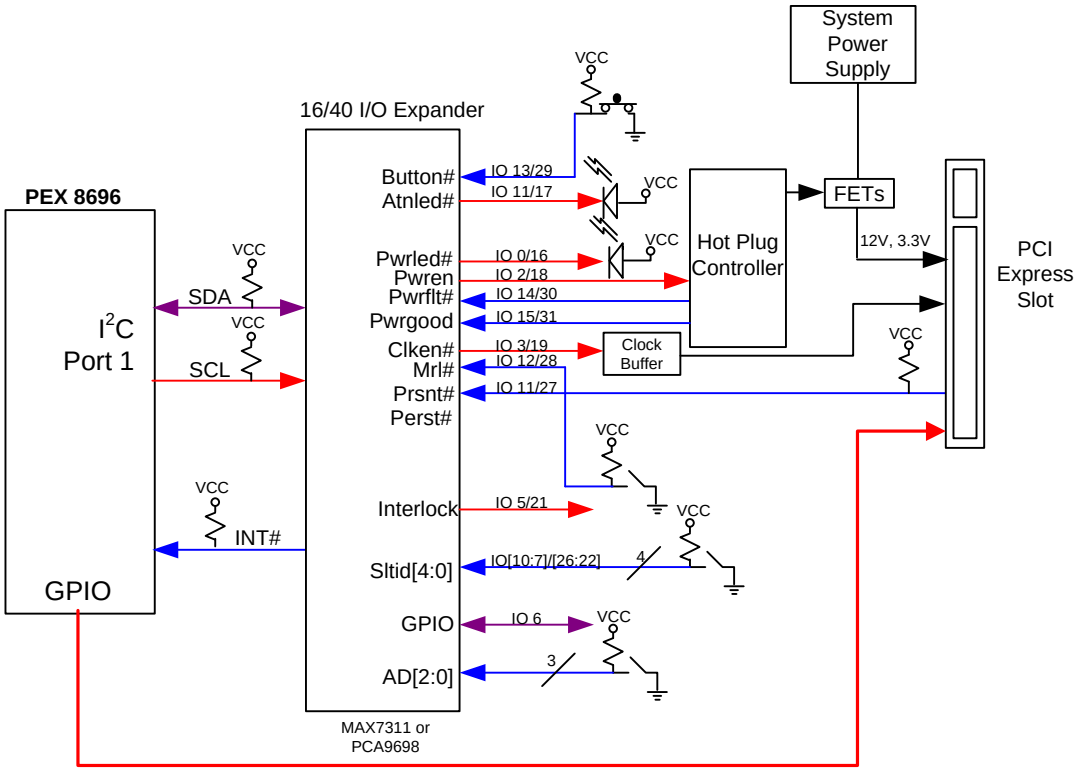


**Figure 14. SHPC Interface to PEX 8696 Block Diagram**

# 6 JTAG Interface

The PEX 8696 supports a five-ball 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, pull JTAG_TDI, JTAG_TMS, and JTAG_TCK up to 2.5V with 1-kohm to 5-kohm resistors. Pull JTAG_TRST# down to VSS with a 1-kohm to 5-kohm resistor. Because the PEX 8696 JTAG clock frequency can be as high as 10 MHz, a 15-ohm series terminator can be added to TCK, TDI, and TDO, to improve signal quality. Figure 15 illustrates a generic JTAG interconnection.

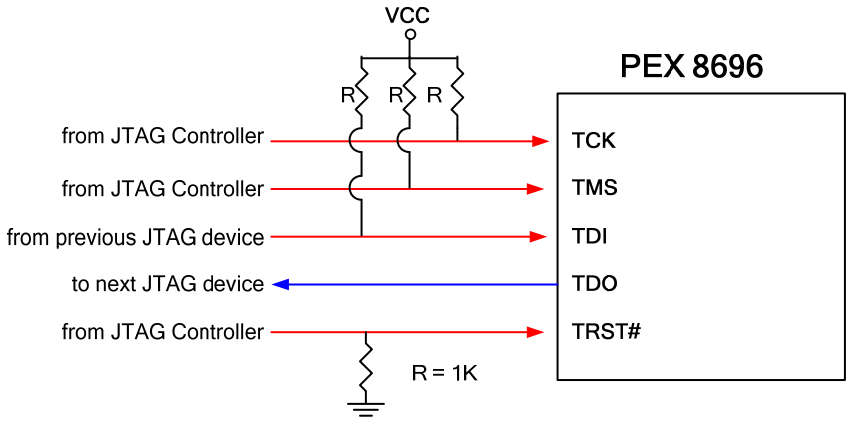


**Figure 15. JTAG Interface Block Diagram**

# 7 PCI Express Port Good Indicators

The PEX 8696 provides up to 24 Active-Low "Port Good" Output balls for the PCI Express Port on the device, PEX_PORT_GOOD[23:0]#. These Output balls can be used to build the Port Status LED circuits, to indicate the status of each PEX 8696 Port. Each Port has five states, which are related to Link Status, Channel Speed, and the Port's Lane width. Table 2 lists the relationship of the LED On/Off patterns to the Port status.

**Table 2. PEX 8696 LED On/Off Patterns, by State**

| State | LED Pattern |
|---|---|
| Link is down | Off |
| Link is up, 5 Gbps, all Lanes are up | On |
| Link is up, 5 Gbps, reduced Lanes are up | Blinking, 0.5 seconds On, 0.5 seconds Off |
| Link is up, 2.5 Gbps, all Lanes are up | Blinking, 1.5 seconds On, 0.5 seconds Off |
| Link is up, 2.5 Gbps, reduced Lanes are up | Blinking, 0.5 seconds On, 1.5 seconds Off |

# 8 Debug Functions

(The optional Debug function is primarily intended for prototyping activities. Its use requires assistance from PLX Technical Support.)

Two major debug functions of the PEX 8696 are External Probe mode (EPM) and SerDes Debug mode (SDM). The EPM function is for viewing the internal state machines and control signals of the three station-based modules and the core-based module. The SDM function is for viewing the 20-bit Receive Bus (elastic buffer exit) and 20-bit Transmit Bus of each Lane of the SerDes, in the PEX 8696. Two Strapping balls are used to enable either Debug mode function. Pulling down the STRAP_PROBE_MODE# ball enables the EPM function. Pulling down the STRAP_SERDES_MODE_EN# ball enables the SDM function. The EPM contains 16 inputs and 39 outputs. The SDM contains 11 inputs and 44 outputs. When either Debug mode is enabled, the parallel Hot Plug, Port Good, General-Purpose I/O, and Spare balls, as well as two other Strapping balls are serviced by EPM and SDM inputs or outputs. cross-references the ball names and their related Debug signal names.

***Notes:*** *Inputs are marked in* ***blue****, outputs are marked in* ***red****.*

*The maximum frequency of Debug mode Output signals, such as PROCMON (N/C, at location W4), is 125 MHz, with fast rise and fall time. When routing these traces to the mictor connector for scope probing,*

*50-ohm, single-ended controlled-impedance traces are recommended. To service normal operation and debug functions, low-capacitive load bus switches can be used to prevent the reflections. For example, a bus switch can be used to separate the LED circuit from the PEX_PORT_GOODx# signals, when they are used as outputs of Debug mode signals.*

**Table 3. Cross-Reference of Ball Names and Related Debug Signal Names**

| Ball Name | Probe Mode Name | SerDes Debug Mode Name |
|---|---|---|
| HP_PWRFLT_C# | stn_sel2 | stn_sel2 |
| HP_BUTTON_C# | stn_sel1 | stn_sel1 |
| HP_MRL_C# | stn_sel0 | stn_sel0 |
| HP_PWRFLT_B# | mod_sel4 | |
| HP_BUTTON_B# | mod_sel3 | |
| HP_PRSNT_B# | mod_sel2 | |
| HP_PWR_GOOD_D | mod_sel1 | |
| HP_PRSNT_D# | mod_sel0 | rcvr_polarity |
| HP_PWRFLT_D# | port_sel1 | ln2_add2 |
| HP_MRL_D# | port_sel0 | ln2_add1 |
| HP_MRL_A# | outA_sel3 | |
| HP_BUTTON_A# | outA_sel2 | |
| HP_PWR_GOOD_C | outA_sel1 | ln_sel1 |
| HP_PRSNT_C# | outA_sel0 | ln_sel0 |
| HP_BUTTON_D# | outB_sel3 | ln2_add0 |
| HP_PWR_GOOD_A | outB_sel2 | |
| HP_PRSNT_A# | outB_sel1 | |
| HP_PWRFLT_A# | outB_sel0 | |
| STRAP_DEBUG_SEL1 | | ln_sel3 |
| STRAP_DEBUG_SEL0 | | ln_sel2 |
| PEX_PORT_GOOD1# | prb_outB17 | rcvr_dat17 |
| PEX_PORT_GOOD0# | prb_outB16 | rcvr_dat16 |
| SPARE0 | prb_outB15 | rcvr_dat15 |
| PEX_PORT_GOOD7# | prb_outB14 | rcvr_dat14 |
| PEX_PORT_GOOD6# | prb_outB13 | rcvr_dat13 |
| PEX_PORT_GOOD5# | prb_outB12 | rcvr_dat12 |
| PEX_PORT_GOOD4# | prb_outB11 | rcvr_dat11 |
| PEX_PORT_GOOD3# | prb_outB10 | rcvr_dat10 |
| HP_ATNLED_B# | prb_outB9 | rcvr_dat9 |
| PEX_PORT_GOOD2# | prb_outB8 | rcvr_dat8 |
| HP_PWREN_B | prb_outB7 | rcvr_dat7 |
| HP_PWRLED_B# | prb_outB6 | rcvr_dat6 |
| HP_PERST_B# | prb_outB5 | rcvr_dat5 |
| HP_CLKEN_B# | prb_outB4 | rcvr_dat4 |
| GPIO7 | prb_outB3 | rcvr_dat3 |
| GPIO6 | prb_outB2 | rcvr_dat2 |
| GPIO5 | prb_outB1 | rcvr_dat1 |

| Ball Name | Probe Mode Name | SerDes Debug Mode Name |
|---|---|---|
| GPIO4 | prb_outB0 | rcvr_dat0 |
| GPIO0 | prb_outA17 | xmit_dat17 |
| GPIO3 | prb_outA16 | xmit_dat16 |
| GPIO2 | prb_outA15 | xmit_dat15 |
| HP_CLKEN_D# | prb_outA14 | xmit_dat14 |
| HP_PWRLED_D# | prb_outA13 | xmit_dat13 |
| HP_PERST_D# | prb_outA12 | xmit_dat12 |
| HP_PWREN_D | prb_outA11 | xmit_dat11 |
| HP_CLKEN_C# | prb_outA10 | xmit_dat10 |
| HP_PWRLED_C# | prb_outA9 | xmit_dat9 |
| HP_ATNLED_D# | prb_outA8 | xmit_dat8 |
| HP_CLKEN_A# | prb_outA7 | xmit_dat7 |
| HP_PWREN_C | prb_outA6 | xmit_dat6 |
| HP_ATNLED_C# | prb_outA5 | xmit_dat5 |
| HP_PERST_C# | prb_outA4 | xmit_dat4 |
| HP_PWRLED_A# | prb_outA3 | xmit_dat3 |
| HP_PERST_A# | prb_outA2 | xmit_dat2 |
| HP_PWREN_A | prb_outA1 | xmit_dat1 |
| HP_ATNLED_A# | prb_outA0 | xmit_dat0 |
| NC_on_W9 | PROCMON | rclk/2 |
| SPARE1 | ext_trig_in | |
| SPARE2 | trig_out | trig_out |
| PEX_PORT_GOOD12# | | xmit_dat19 |
| GPIO1 | | xmit_dat18 |
| GPIO8 | | rcvr_dat18 |
| GPIO9 | | rcvr_dat19 |
| GPIO10 | | rx_sts0 |
| GPIO11 | | rx_sts1 |
| GPIO12 | | rx_sts2 |

# 9 PEX 8696 Configuration Strapping Balls

The PEX 8696 has a total of 47 Strapping balls. Twenty-six of them service different configuration functions (STRAP_NT_ENABLE#, STRAP_NT_UPSTRM_PORTSEL[4:0], STRAP_UPSTRM_PORTSEL[4:0], STRAP_STNx_PORTCFG[1:0] for each of the six Stations, and STRAP_VS_MODE[2:0]). For the PEX 8696, internal pull-up and pull-down resistors set the default configuration. External pull-up and pull-down resistors are not required unless the strapping signals are connected to circuit traces (the internal resistors are relatively weak, and may not be strong enough to hold the circuit traces to the default input states). Table 4 lists the names, functions and default values of the PEX 8696 Configuration Strapping balls.

**Table 4. Configuration Strapping Balls**

| Ball/Signal Name | Functions | Default |
|---|---|---|
| STRAP_FAST_BRINGUP# | ***Factory Test Only.*** | Internal pull-up |
| STRAP_G1_COMPATIBLE# | When this ball is tied High, the Data Rate Identifier symbol in the TS Ordered-Sets always advertises support for both the Gen 2 data rate and Autonomous Change.<br>When this ball is tied Low, if the Link training sequence fails during configuration, the next time the LTSSM exits the detect state, TS Ordered-Sets advertise only the Gen 1 data rate, and no Autonomous Change support. | Internal pull-up |
| STRAP_I2C_CFG_EN# | When this ball is tied Low, the I2C Bus configures the device. The Links do not start training until I2C Sets the *Configuration Release* register *Initiate Configuration* bit.<br>When this ball is tied High, the device operates normally. | Internal pull-up |
| STRAP_SMBUS_EN# | When this ball is tied Low, SMBus Slave protocol is implemented on the I2C_SCL0 and I2C_SDA0 2-wire bus.<br>When this ball is tied High, I2C Slave protocol is implemented on the I2C_SCL0 and I2C_SDA0 2-wire bus | Internal pull-up |
| STRAP_NT_P2P_EN# | This input should be pulled Low, unless the NT PCI-to-PCI bridge between the internal Virtual PCI Bus and the NT Port Virtual Interface must be disabled for software compatibility to earlier NT mode switches. | Internal pull-up |
| STRAP_PLL_BYPASS# | ***Factory Test Only. When pulled Low, PLL that generates internal clock is bypassed.*** | Internal pull-up |
| STRAP_PROBE_MODE# | ***Factory Test Only.*** | Internal pull-up |
| STRAP_SERDES_MODE_EN# | ***Factory Test Only.*** | Internal pull-up |
| STRAP_RESERVED[1:0] | ***Factory Test Only. NOTE: These pins MUST be pulled low for normal operation. External pull-downs must be used to overcome the internal pull-ups.*** | Internal pull-up |
| STRAP_RESERVED[3:2] | ***Factory Test Only.*** | Internal pull-downs |
| STRAP_RESERVED16 | ***Factory Test Only. Must be tied directly to Ground.*** | No internal pull-up or pull-down. |
| STRAP_RESERVED17# | ***Factory Test Only.*** | Internal pull-up |
| STRAP_NT_ENABLE# | Enable NT function | Internal pull-up |
| STRAP_NT_UPSTRM_PORTSEL[4:0] | NT Port select | Internal pull-downs |

| Ball/Signal Name | Functions | Default |
|---|---|---|
| STRAP_STN0_PORTCFG[1:0] | Port configuration per station.<br><br>LL = x4x4x4x4<br>LH = x16<br>HL = x8x8<br>HH = x8x4x4 | Internal pull-downs |
| STRAP_STN1_PORTCFG[1:0] | | |
| STRAP_STN2_PORTCFG[1:0] | | |
| STRAP_STN3_PORTCFG[1:0] | | |
| STRAP_STN4_PORTCFG[1:0] | | |
| STRAP_STN5_PORTCFG[1:0] | | |
| STRAP_DEBUG_SEL[1:0] | ***Factory Test Only.*** | Internal pull-downs |
| STRAP_TESTMODE[4:0] | Test mode function select. Defines PEX_PORT_GOOD[23:0]# and GPIO[42:24] signal functionality. NOTE: 00h is not a valid setting. For normal operation, these balls should be pulled to 0Dh by external pull-ups. | Internal pull-downs |
| STRAP_UPSTRM_PORTSEL[4:0] | Upstream Port select | Internal pull-downs |
| STRAP_VS_MODE[2:0] | Determines the number of Virtual Switches that are enabled. If 000b, no Virtual Switches are enabled and the device is in Mode-1 and operates as a single switch. | Internal pull-downs |

# 10 GPIO Balls

The PEX 8696 has 43 GPIO balls – 19 are dedicated GPIO balls, and 24 share GPIO and PEX_PORT_GOOD[23:0]# functions in standard operation. Depending on the settings of the Test mode balls, STRAP_TESTMODE[4:0], the GPIO balls can be set as input or output. If Serial Hot Plug is implemented, using external I2C I/O Expanders, it is recommended that the 19 dedicated GPIO balls be configured as PERST# outputs (STRAP_TESTMODE setting 0Ch or 0Dh) and routed directly to the slots, bypassing the I/O Expander.

# 11 Power Supplies, Sequencing, and De-Coupling

The switch's maximum power consumption is approximately 20W. Special cooling requirements may exist, depending upon the system environment. (Refer to the *PEX 8696 Data Book* for details).

## 11.1 Power Supplies

The PEX 8696 has the following Power ball groups:

- VDD10 – Digital core logic supply
- VDD10A – SerDes analog supply
- VDD25 – Hot Plug, serial EEPROM, $I^2C$, JTAG, Port Status indicators, I/O buffers
- VDD25A – PEX_REFCLK PLL supply

At the board level, VDD10 and VDD10A can share a common 1.0V ±5% power plane, and VDD25 and VDD25A can share a common 2.5V power plane. The current demands for these supplies can be high, depending upon the device (approximately 80 mA per Lane, plus 32 mA); therefore, ensure that the power plane is sufficiently sized, to support the specified operating current. For best performance, the

1.0V ±5% plane should have an adjacent ground plane that provides an interplane capacitor to supply high-frequency transient currents. Provide a sufficient number of discrete capacitors for mid- and low-frequency de-coupling. The recommendation is that 0201-sized capacitors be used in close proximity to these power balls.

VDD10A has a lower noise tolerance than the digital supplies. Therefore, VDD10A might require additional filtering, depending upon the 1.0V ±5% power plane noise. The SerDes can tolerate ±5% variance on the supply rails, due to noise and IR drop. VDD25 power is used for the single-ended I/O buffers – Hot Plug, serial EEPROM, JTAG, $I^2C$, and the Port Status indicators. Although power consumption for this supply is relatively small, the output drivers have fast edge rates, and therefore, require that adequate power de-coupling be provided, to supply transient current to the drivers. It is preferred that VDD25 be implemented as a plane or partial plane, either on a signal layer or main power plane layer. Provide 0.1 and/or 0.01 µF ceramic capacitors, along with one or more 10-µF tantalum capacitors, to de-couple the VDD25 power balls. The number of capacitors required depends upon the number of 2.5V I/O balls utilized in the design, and the existence or absence of an interplane capacitance for the VDD25 rail.

## 11.2 Power Sequencing

There is no power sequencing requirement.

## 11.3 Board-Level De-Coupling

Board-level de-coupling requirements for high-speed digital designs are highly dependent upon several factors, including:

- Printed circuit board (PCB) layer stack-up
- Differential versus single-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 de-coupling solution that will work for all designs.

Board-level power supply de-coupling exists primarily in two forms:

- Parallel plane capacitance
- Use of discrete capacitors

Parallel plane capacitance exists between a PCB's DC power and ground planes. PCB reference planes have a small amount of 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 de-coupling. Figure 16 illustrates attenuation curves measured for a PCI Express test board. The plot illustrates the bare board power-to-ground impedance (indicated in black), compared with the impedance of various power planes after de-coupling capacitors are populated. Notice that as frequencies surpass 200 MHz, the impedance profile is affected only by the bare-board capacitance. Also note the impedance holes at 7 MHz. It is suggested that discrete capacitor values be adjusted to eliminate measured holes.

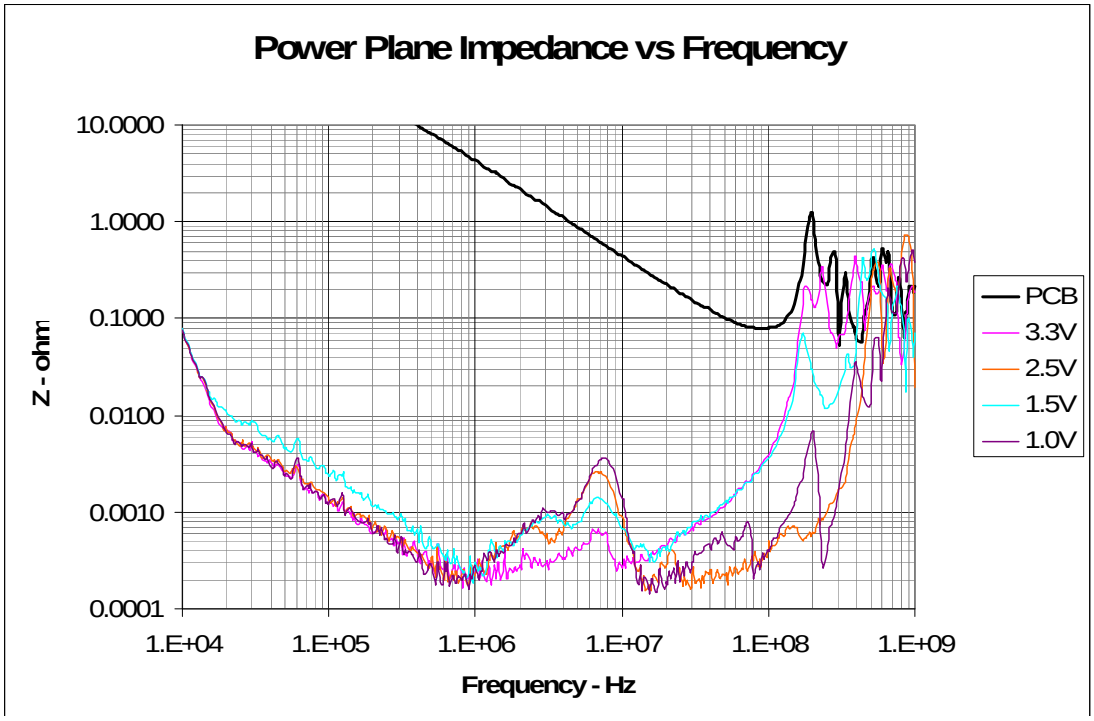


**Figure 16. Power Plane Impedance versus Frequency**

A power and ground plane separation of 0.254 mm (0.010 in.) results in approximately 100 pF/in$^2$, while a separation of 0.102 mm (0.004 in.) provides approximately 200 pF/in$^2$. 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.

As for discrete capacitors, the footprint and physical size of discrete capacitors have a significant effect on the frequencies in which the capacitors provide effective de-coupling. To minimize series inductance, use smaller-packaged ceramic capacitors (*such as* 0402 or 0201) for mid-ranged frequency de-coupling (20 to 250 MHz). Use a mixed selection of capacitor values, *such as* 0.1 and 0.01 µF, to lower the impedance across a wide frequency range.

Capacitor footprint layout is important in determining the frequencies at which they are effective. Avoid adding trace segments from the capacitor pads to the vias. These segments add more series inductance, thereby lowering the discrete capacitor LC resonant frequency. 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, the placement of small discrete capacitors is not critical. Place the capacitors on the solder side of the board, under 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), place the capacitors as close to the balls as possible. If a PCB layer stackup is such that plane capacitors are not possible, add power or ground fill areas on the signal layers, as follows:

- If a signal layer is referencing a DC ground plane, fill with power
- 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 the component balls.

Multi-layer ceramic chip capacitors (*such as* 10 to 22 µF) can be used for bulk de-coupling of lower-frequency components. The proximity of these capacitors is not critical; therefore, they can be placed outside the BGA matrix.

It is strongly recommended to measure the attenuation-versus-frequency profile of each power rail, on a completed board that is loaded only with bypass capacitors. This serves to confirm that there are no attenuation holes in the power-de-coupling design.

Figure 17 illustrates examples of how various footprints for 0603-size capacitors can change series inductance.

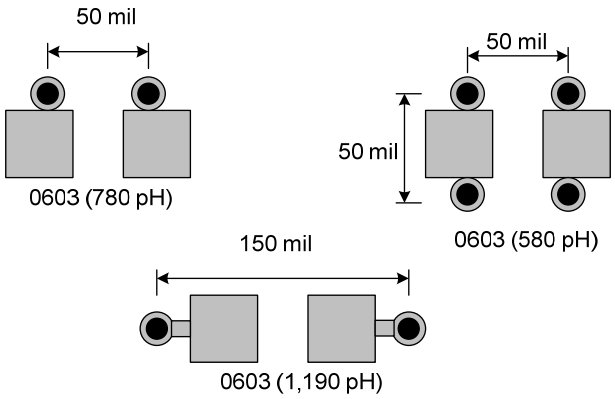


**Figure 17, Capacitor Footprint Effects on Series Inductance**

# 12 References

The following is a list of documentation to provide further details.

- PLX Technology, Inc.
  870 W Maude Avenue, Sunnyvale, CA 94085 USA
  Tel: 800 759-3735 (domestic only) or 408 774-9060, Fax: 408 774-2169, www.plxtech.com
  - *PEX 8696AA Data Book, Version 1.0* or higher
- PCI Special Interest Group (PCI-SIG)
  3855 SW 153rd Drive, Beaverton, OR 97006 USA
  Tel: 503 619-0569, Fax: 503 644-6708, www.pcisig.com
  - *PCI Local Bus Specification, Revision 3.0*
  - *PCI Bus Power Management Interface Specification, Revision 1.1*
  - *PCI to PCI Bridge Architecture Specification, Revision 1.1*
  - *PCI Express Base Specification, Revision 2.0*
  - *PCI Express Card Electromechanical (CEM) Specification, Revisions 1.0a and 1.1*
- *Right the First Time: A Practical Handbook on High Speed PCB and System Design*, by Lee Ritchie