IDT PCI Express Gen1 Hardware Design Guide

Notes

Introduction

This document provides general guidelines to help design systems based on IDT's PCI Express® 1.0a/1.1 base specification compliant devices. These devices offer 2.5 GT/S (Gen1) data speed per lane. IDT's 89HPES24N3 PCI Express 3-port switch (PES24N3) switch is used throughout this document as an example. However, the information is relevant to all of the IDT switches including PE3T3, PES4T4, PES5T5, PES6T5, PES8NT2, PES8T5, PES8T5A, PES12N3A, PES12NT3, PES16NT2, PES16T4, PES16T7, PES24N3A, PES24NT3, PES24T6, PES32T8, PES32H8, PES48H12 and PES64H16.

The document also describes the following switch interfaces (using PES24N3 as an example) and provides appropriate board design related recommendations:

- 2.5Gbps high-speed differential pairs
- Clocking
- Footprint example
- Power and decoupling
- Boot Configuration Vector (BCV)
- Reset scheme
- Hot-Plug
- SMbus Eprom (Optional)

The PES24N3 is a three port transparent switch that contains 24 PCI Express lanes. Each of the three ports is statically allocated 8 lanes with ports labeled as A, B, and C. Other switches may have a different number of lanes (between 3 and 64, depending on the device) and the port numbering may be different (Port 0, Port 1, Port 2, etc up to Port 16). Port A (or Port 0 in some switches) is always the upstream port while ports B and C are always downstream ports. Optional initialization from a serial EEPROM is selected via the Switch Mode (SWMODE[3:0]) inputs, part of the Boot Configuration Vector (BCV).

During link training, link width is automatically negotiated. Each PES24N3 port is capable of independently negotiating to a x8, x4, x2, or x1 width. Thus, the PES24N3 may be used in virtually any three port switch configuration (e.g., {x8, x8, x8}, {x4, x4, x4}, {x4, x2, x1}, etc.). The PES24N3 supports static lane reversal. Lane reversal for upstream port A may be configured by asserting the PCI Express Port A Lane Reverse (PEALREV) input signal or through serial EEPROM or SMBus initialization. Lane reversal for ports B and C may be enabled via a configuration space register, serial EEPROM, or the SMBus.

2.5Gbps Differential Pairs

- The PES24N3 includes 50 Ohm resistor on-die terminations on both transmit and receive pins so that no external termination is required per PCIe specifications.
- No extra vias should be added in addition to those needed for IC pads or a connector.
- The PES24N3 Transmit differential pairs TX could be routed on the top side of the board to take advantage of the switch pinout, to facilitate decoupling and to reduce the number of vias. 5 mil wide microstrip spaced 11 mils apart (center-to center) is an appropriate topology in order to yield a differential (odd mode) impedance of 100 Ohms. However, for a given trace impedance, the trace width may vary depending on which layer the trace resides or the overall board thickness.
- The PES24N3 Receive differential pairs RX could be routed on an internal layer near the bottom of the board to reduce board via stubs. 6 mil wide stripline line spaced 11 mils apart (center-to center) is an appropriate topology in order to yield a differential (odd mode) impedance of 100 Ohms. However, for a given trace impedance, the trace width may vary depending on which layer the trace resides or the overall board thickness.
- Individual traces within a given differential pair (positive and negative) must be matched in length to a tolerance of 5 mils.
- In addition, the spacing between different pairs of the 2.5G Serial Link differential pairs (transmit as well as receive) must be at least 20 mils edge-to-edge to prevent crosstalk.
- It should be noted that trace length matching among pairs is not required as the PCIe specifications allow up to 20ns of skew among differential pairs.
- AC coupling capacitors are associated with the Tx differential pairs and should be located symmetrically on the top or bottom layer between the PES24N3 and the PCIe connectors, as shown to Figure 1. Values between 75nF and 200nF are appropriate, per the PCI Express specification. For example, 100nF capacitors are used on the PES24N3 evaluation board.

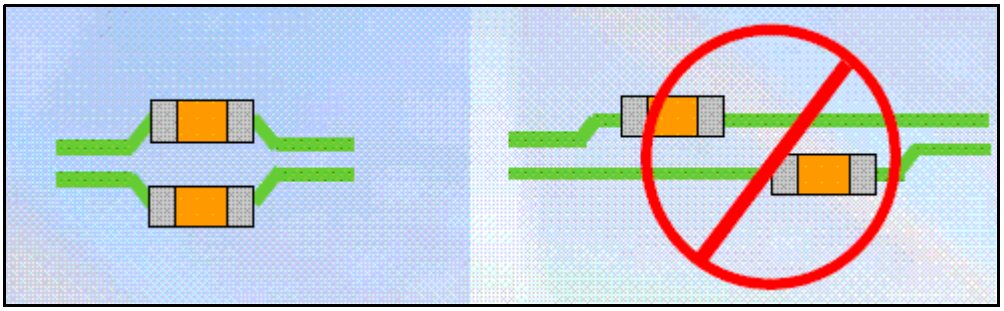

Figure 1 AC Decoupling Capacitor Layout

- Vias should be avoided as much as possible on the PCIe differential pairs (as described above) since they can result in up to a 0.25 dB loss.
- When a via is required, its pad size should be less than 25 mils, its hole size should be less than 14 mils and its anti-pads less or equal to 35 mils.
- Depending on the system topology and the maximum targeted trace length, regular FR4 material is appropriate dielectric material. In the case of a backplane type of application, higher quality, lower loss material, such as Nelco 4000-13, may be needed.
- Simulations are recommended and IDT provides an Hspice development kit.

To facilitate the routing process, polarity inversion and lane reversal features can be used as followed:

- Polarity Inversion

  Each port of the PES24N3 supports automatic polarity inversion as required by the PCIe specification. Polarity inversion is a function of the receiver and not the transmitter. The

transmitter never inverts its data. Polarity inversion is a lane and not a link function. Therefore, it is possible for some lanes of link to be inverted and for others not to be inverted.

- Lane Reversal

  The PCIe specification describes an option lane reversal feature. The PES24N3 does not support the automatic lane reversal feature outlined in the PCIe specification. However, it does support static lane reversal on a per port basis.

  Associated with each PES24N3 switch port is a lane reversal signal. The lane reversal signal for port A is PEALREV, for port B is PEBLREV, and for port C is PECLREV. The status of the lane reversal signals sampled during a fundamental reset may be determined from the PALREV, PBLREV, and PCLREV fields in the PA_SWSTS register.

  The port lane reversal signals are sampled during a fundamental reset and used as the initial value of the PALREV, PBLREV, and PCLREV fields in the PA_SWCTL register. When these bits are set, then the lanes of the corresponding port(s) are reversed during link training.

  The operation of lane reversal is dependant on the maximum link width selected by the MAXLNKWDTH[1:0] pads. Lane reversal mapping for the various non-trivial maximum link width configurations is illustrated in the following figures.

  When link training occurs, the corresponding lane reversal bits in the PA_SWCTL register are examined. If a bit is set, then the lanes associated with that link are revered. This mechanism may be used to configure lane reversal via the serial EEPROM, slave SMBus, or root.

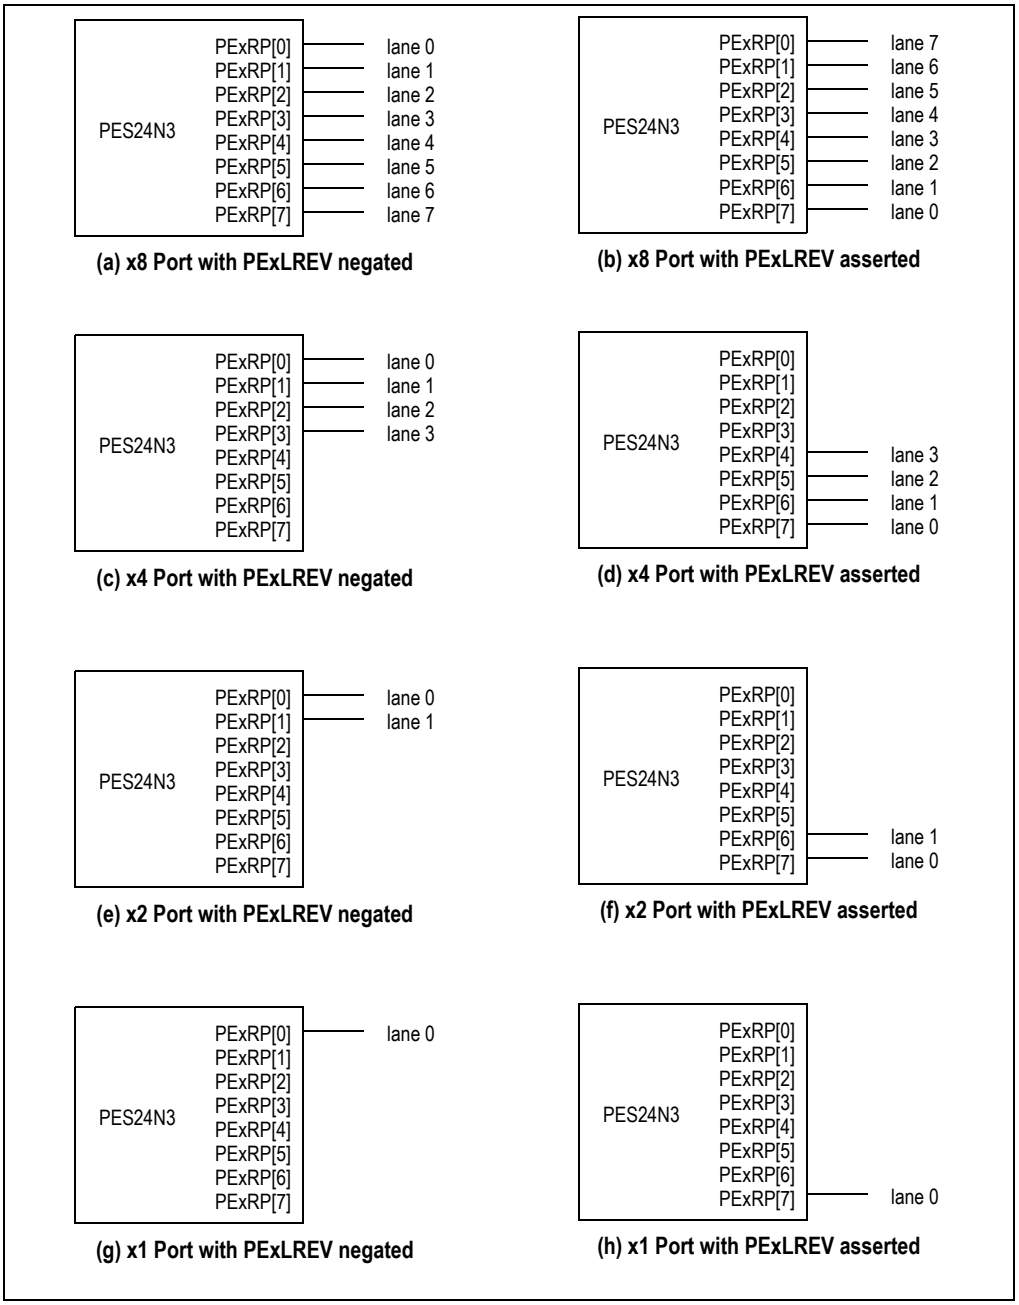

Figure 2 Lane Reversal for Maximum Link Width of x8 (MAXLNKWDTH [1:0]=0x3)

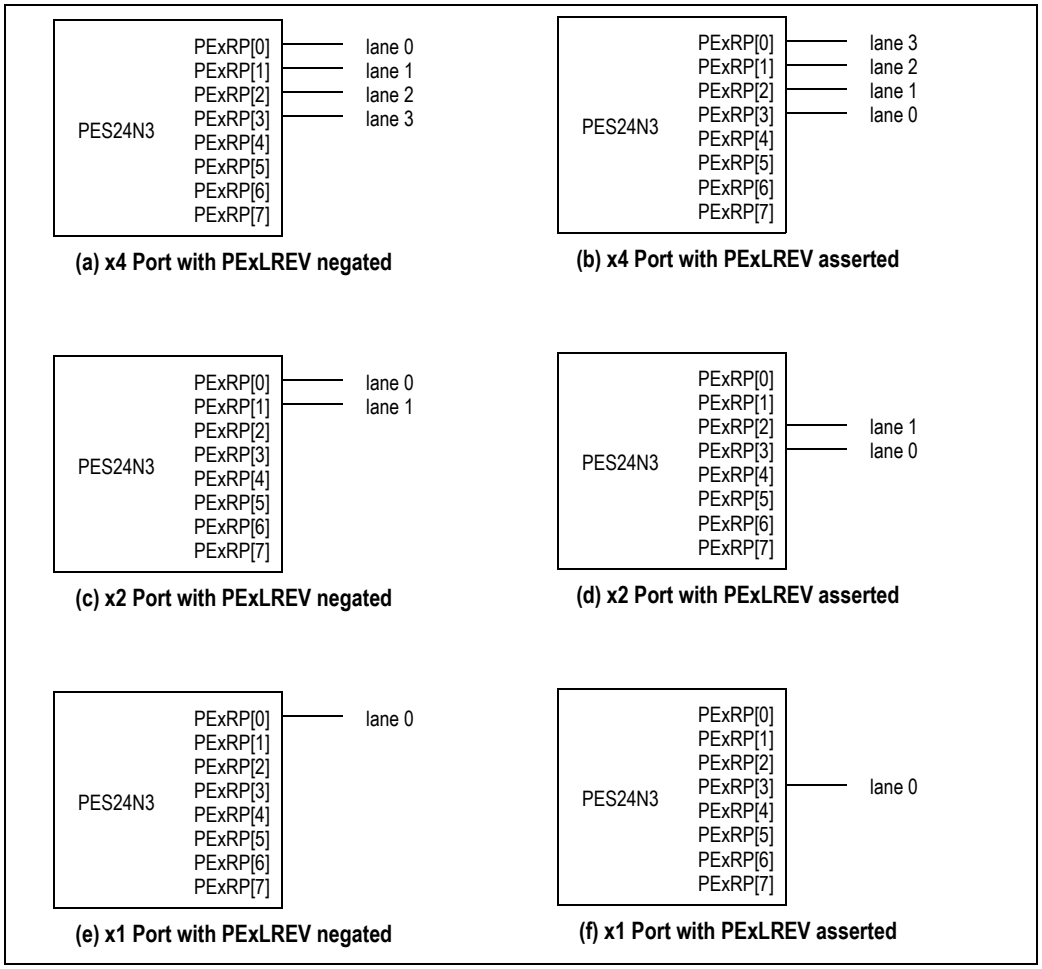

Figure 3 Lane Reversal for Maximum Link Width of x4 (MAXLNKWDTH [1:0]=0x2)

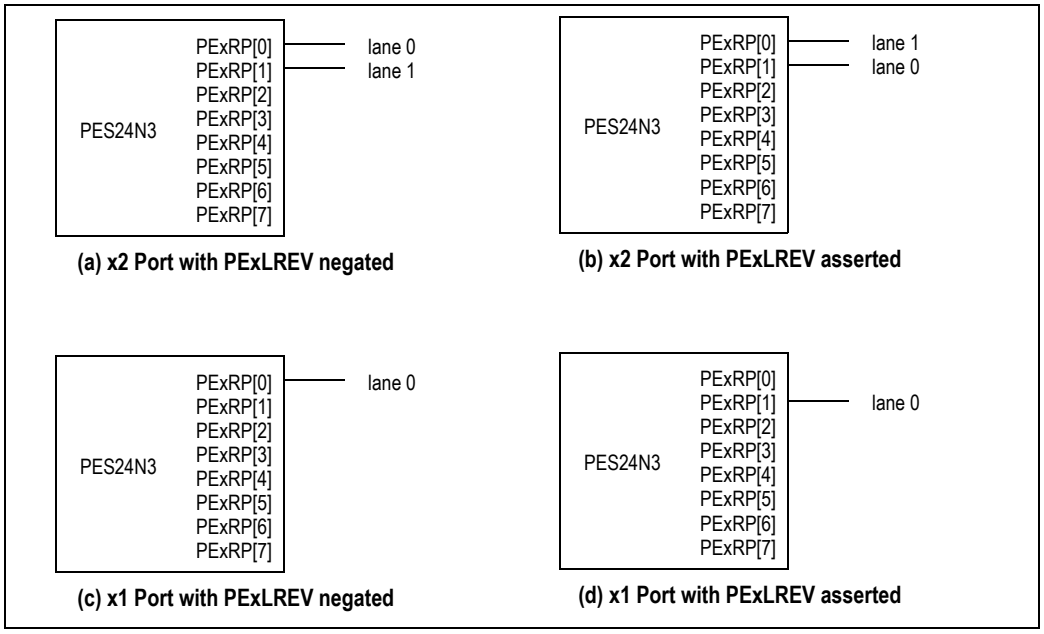

Figure 4 Lane Reversal for Maximum Link Width of x2 (MAXLNKWDTH [1:0]=0x1)

Clocking

- Clock lines must be kept at least 25 mils away from other signals.
- No extra via should be added in addition to those needed for IC pads.
- If an inner signal layer is used, the clock line should stay on this layer except when it connects to the IC pads.
- Individual traces within a given differential pair (positive and negative) must be matched in length to a tolerance of 5 mils.
- The clock lines are differential pairs and they should be routed according to the restrictions of this section as well as the 2.5Gbps Differential Pairs section.
- Any termination should be implemented so as to minimize the trace stubs.
- There is no skew/trace length constraint between the two pairs of clocks connected to the chip clock inputs.
- The PES24N3 input clock buffers support 100MHz and 125MHz clock frequencies as well as SSC as defined by PCI Express specifications Rev 1.0a.
- The PES24N3 input clock buffers support LVDS, LVPECL, CML, and HSTL clock signaling but require AC decoupling. A 0402 0.1µf capacitor is appropriate.
- For an LVDS clock source, a 2K common mode resistor must be included between the clk_p and clk_n signals in front of the decoupling capacitors. See Figure 5.

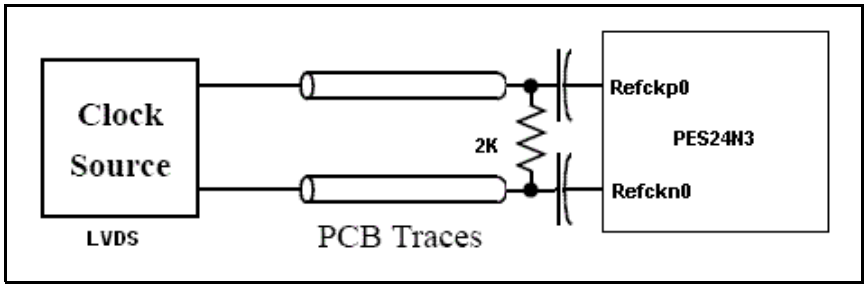

Figure 5 LVDS Clock Source Circuit Implementation

Footprint Example

In this paragraph, a top and a bottom footprint of the PES24N3 chip are suggested in Figures 6 and 7. There are many other possible implementations and these are only given as an example.

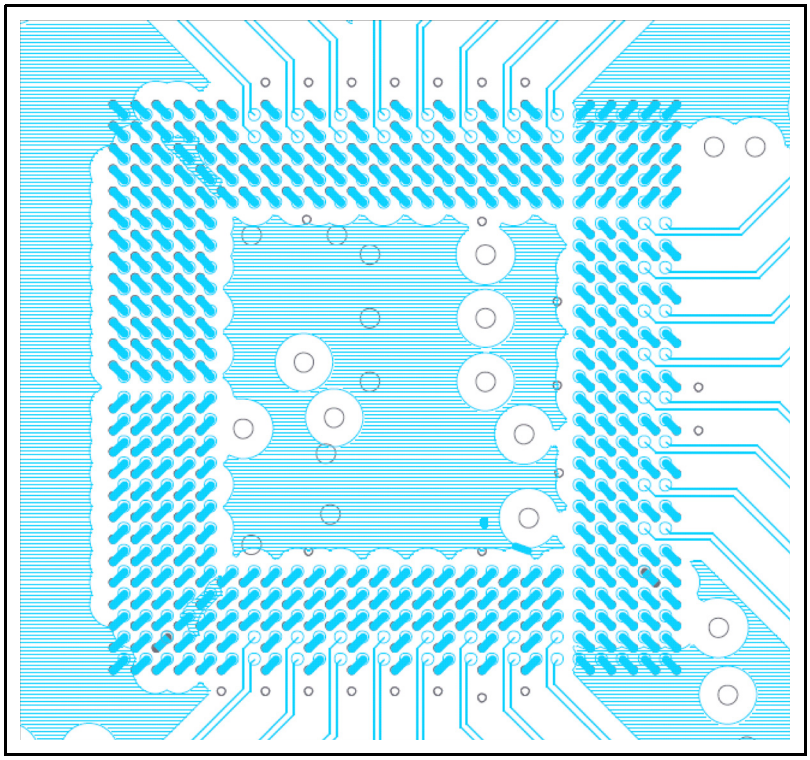

Figure 6 89PES24N3 Footprint Example (Top Layer)

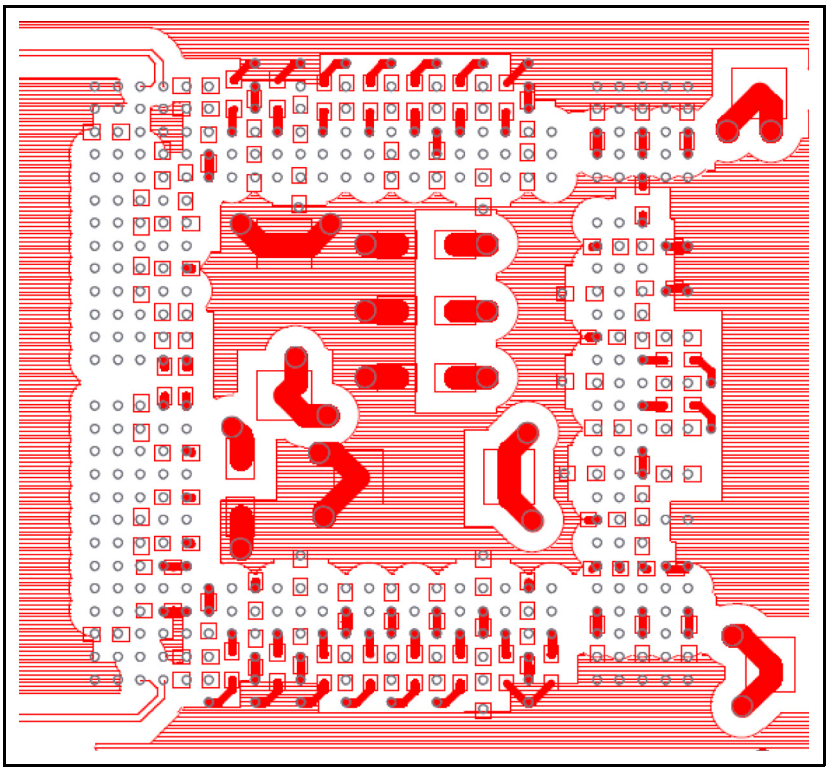

Figure 7 89PES24N3 Footprint Example (Bottom Layer)

Power and Decoupling Scheme

The PES24N3 has five different power supplies pins:

- V_{DD}Core powers the digital core of the switch.
- V_{DD}PE powers the SERDES core and the receive clock recovery logic.
- V_{DD}APE powers analog circuits such as PLL and bias generator.
- V_{TT}PE is the termination voltage used on the SERDES TX lines. Vtt can be adjusted to modify the TX common mode voltage as well as the voltage swing.
- V_{DD}IO powers the low speed IOs of the switch.

The respective standard and maximum power consumption are listed in the 89PES24N3 Data Sheet. Figure 8 describes a possible bypassing scheme for high frequency noise (0.1µF and 0.01µF). A bigger capacitor should be used to filter out low frequency noise. Because capacitors can be located around the chip, they do not represent any specific design challenge.

Note the following:

- 0402 package ceramic capacitors are recommended for 0.1µF and 0.01µF capacitors
- It should be noted that some of the vias have to be shared in order to create space for placing a capacitor next to a pin.
- Capacitors are placed based on space availability
- Larger 1µF and 47µF capacitors should be added around the part. Two bigger capacitors per voltage supply are appropriate. One option is spreading out the big capacitors at four corners, top and bottom layers of the chips.
- Prioritize the bypass capacitors in the following order for each supply:
  1. V_{DD}Core
  2. V_{DD}PE
  3. V_{DD}APE
  4. V_{TT}PE
  5. V_{DD}IO

- V_{DD}Core, V_{DD}PE, and V_{DD}APE can be generated from the same voltage source provided that additional filtering is implemented for V_{DD}APE and that the voltage source current rating meets the chip maximum requirement without exceeding its own thermal limitations.
- V_{TT}PE will require its own voltage source, preferably adjustable since its value may be subject to change if bigger voltage swings are required. Refer to 89PES24N3 Data Sheet for more details on the termination voltage requirements.
- Depending on whether switching or linear regulators are used, a total of two or three voltage sources will be required: one for V_{DD}Core due to its current need, one for V_{DD}PE and V_{DD}APE, and one for V_{TT}PE.

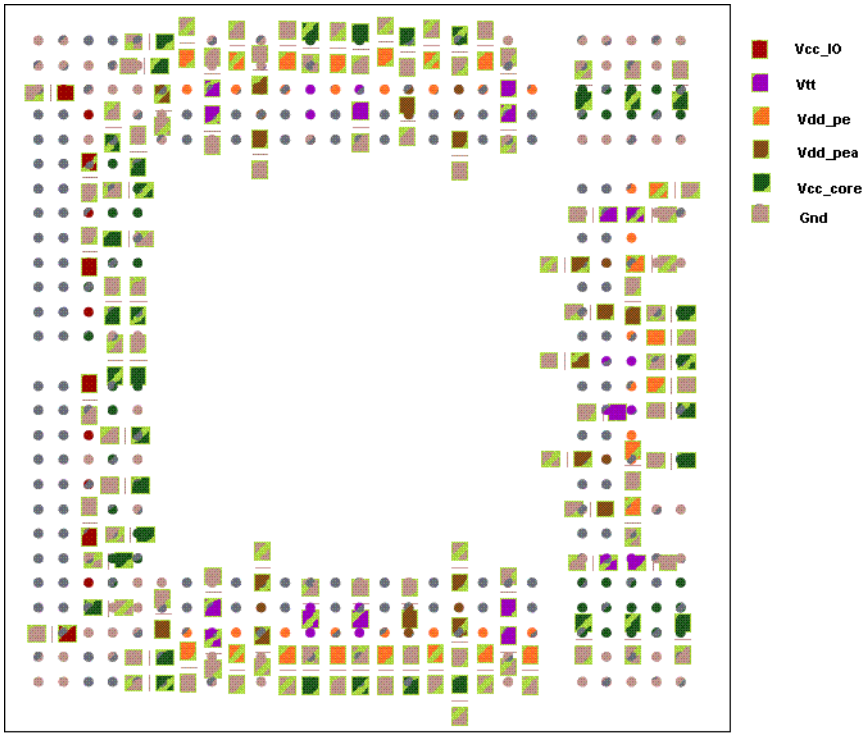

Figure 8 Bypassing Scheme Example Based 0.1nF 0402 Caps

Boot Configuration Vector (BCV)

A boot configuration vector consisting of the signals listed in Table 1 is sampled by the PES24N3 during a fundamental reset when PERSTN is negated. The boot configuration vector defines essential parameters for switch operation.

Signal	Type	Name/Description
CCLKDS	I	**Common Clock Downstream.** The assertion of this pin indicates that all downstream ports are using the same clock source as that provided to downstream devices.This bit is used as the initial value of the Slot Clock Configuration bit in all of the Link Status Registers for downstream ports. The value may be override by modifying the SCLK bit in the PB_PCIELSTS or PC_PCIELSTS register.
CCLKUS	I	**Common Clock Upstream.** The assertion of this pin indicates that the upstream port is using the same clock source as the upstream device. This bit is used as the initial value of the Slot Clock Configuration bit in the Link Status Register for the upstream port. The value may be overridden by modifying the SCLK bit in the PA_PCIELSTS register.
MSMBSMODE	I	**Master SMBus Slow Mode.** The assertion of this pin indicates that the master SMBus should operate at 100 KHz instead of 400 KHz. This value may not be overridden.
PERSTN	I	**Fundamental Reset.** Assertion of this signal resets all logic inside the PES24N3 and initiates a PCI Express fundamental reset.
RSTHALT	I	**Reset Halt.** When this signal is asserted during a PCI Express fundamental reset, the PES24N3 executes the reset procedure and remains in a reset state with the Master and Slave SMBuses active. This allows software to read and write registers internal to the device before normal device operation begins. The device exits the reset state when the RSTHALT bit is cleared in the PA_SWCTL register by an SMBus master.
SWMODE[3:0]	I	**Switch Mode.** These configuration pins determine the PES24N3 switch operating mode. *0x0 - Transparent mode* *0x1 -Transparent mode with serial EEPROM initialization* *0x2 - Reserved* *0x3 - Reserved* *0x4 through 0x7 - Reserved* *0x8 - 10-bit loopback test mode* *0x9 - Reserved* *0xA - Internal pseudo random bit stream self-test test mode* *0xB - External pseudo random bit stream self-test test mode* *0xC - Reserved* *0xD - SerDes broadcast test mode* *0xE - 0xF Reserved*

Table 1 Boot Configuration Vector Signals

Pull-up and pull-down resistors directly connected to these signals are appropriate. For debugging purposes, dip-switches provide greater flexibility.

Reset Scheme

PCI Express defines two reset categories: fundamental reset and hot reset. This section will focus on the former one, which is implemented through side band signals. For more information on the PES24N3 reset characteristics, please refer to the 89HPES24N3 PCI Express Switch User Manual.

There are two sub-categories of fundamental reset: cold reset and warm reset. A cold reset occurs following a device being powered on and assertion of PERSTN. A warm reset is a fundamental reset that occurs without removal of power.

The PES24N3 implements a reset input pin:

- PERSTN: this is the Fundamental input Reset pin. Assertion of this signal reset all logic inside the PES24N3 and initiates a PCI Express fundamental reset.

Figure 9 shows one possible reset implementation where the two downstream endpoints can independently have a fundamental reset. This reset scheme should be used if Hot-Plug support on either one of the two downstream ports is required.

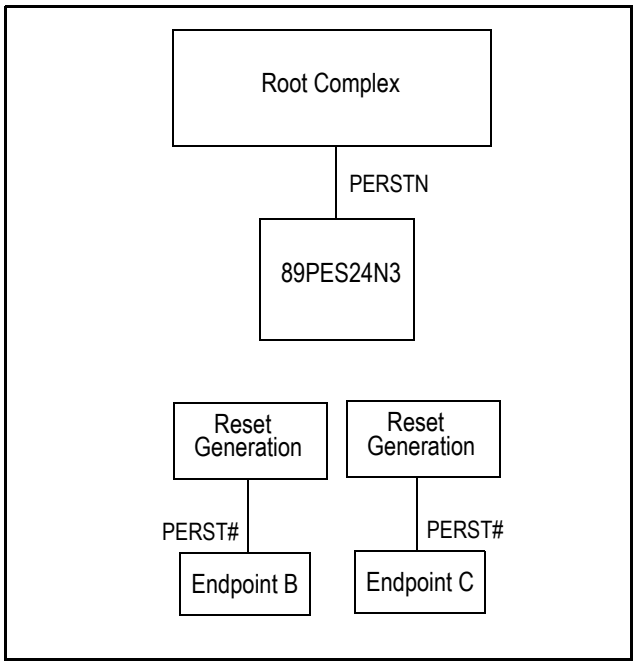

Figure 9 Reset Scheme for Hot Plug Support

If Hot Plug support is not required, the reset circuit can be simplified and the EPROM becomes optional in most applications. Figure 10 displays an appropriate reset circuit.

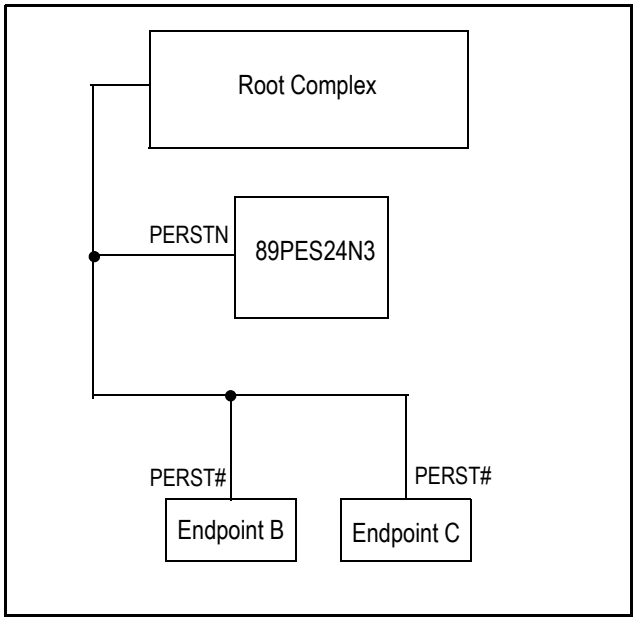

Figure 10 Simplified Reset Scheme

Hot-Plug

The PES24N3 switch may be used in one of three hot-plug configurations:

- An application in which the two downstream ports are connected to slots into which add-in cards may be hot-plugged.
- An add-in card application. Here the two downstream ports are hardwired to devices on the add-in card and the upstream port serves as the add-in card's PCIe interface. In this application, the upstream port may be hot-plugged into a slot on the main system.
- A carrier card application. In this application, the two downstream ports are connected to slots which may be hot-plugged and the entire assembly may be hot-plugged into a slot on the main system. Since this application requires nothing more than the functionalities described above, it will not be discussed further.

The first application requires an external IO expander to be connected to the switch SMBus master interface, along with a power controller in charge of providing power supply selectively to either one or the other downstream slot.

Since this document focuses on system level implementation, it will not describe the switch internal Hot-Plug features any further. Additional information can be found in Chapter 6 of the PES24N3 User Manual referenced earlier. Figure 11 displays a possible Hot-Plug implementation.

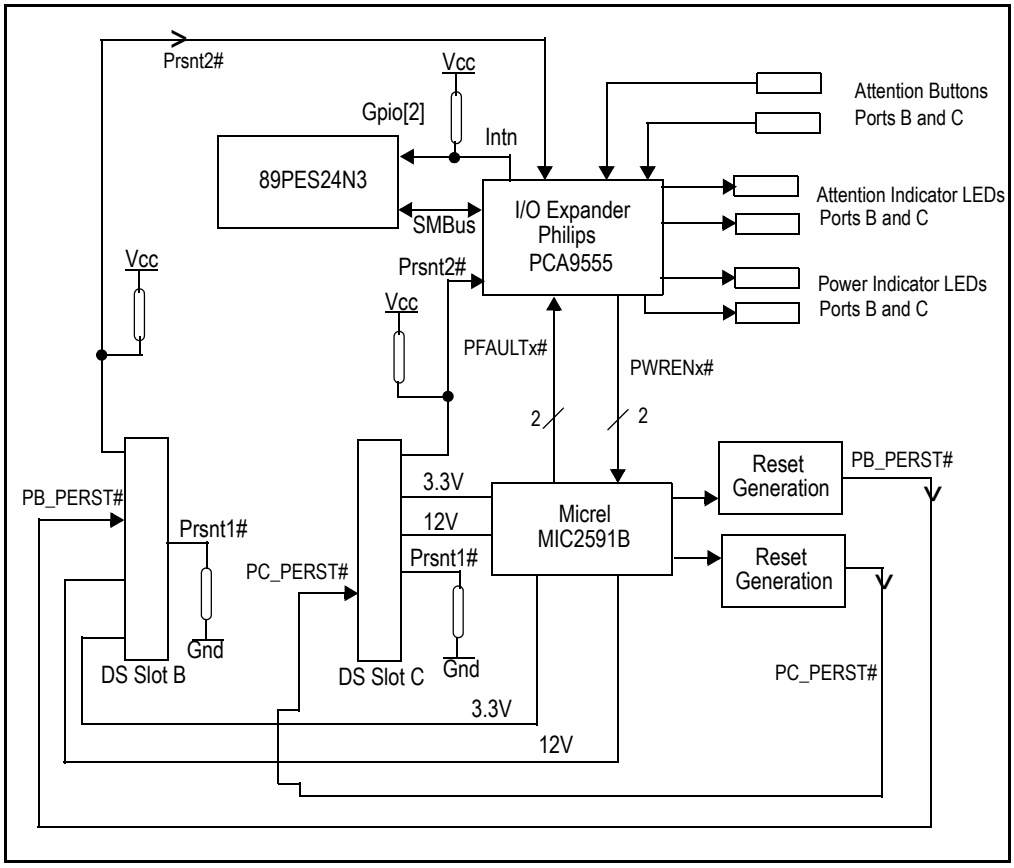

Figure 11 Hot-Plug Circuit Implementation

SMBus

The PES24N3 contains two SMBus interfaces:

- The slave SMBus interface provides full access to all software visible registers in the PES24N3, allowing every register in the device to be read or written by an external SMBus master. The slave SMBus may also be used to initialize the serial EEPROM used for initialization.
- The Master SMBus interface provides connection for an optional external serial EEPROM used for initialization and an optional I/O expander used for hot-plug signals.

Figure 12 provides a possible implementation involving all the SMBus devices that can connect to the PES24N3 Master and Slave SMBus interfaces.

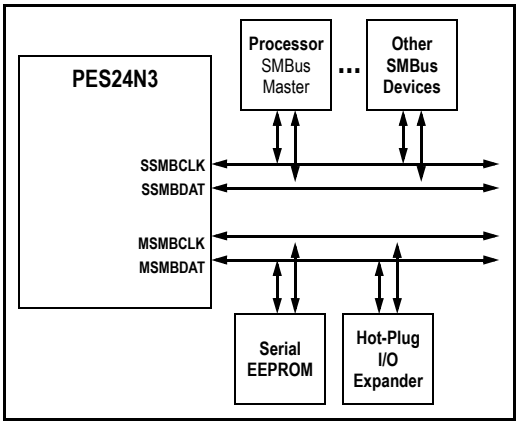

Figure 12 SMBus Split Configuration

The SMBus interface consists of two signals, clock and data, along with several address pins which define the interface SMBus address. 2K pull-up resistors are required on the clock and data lines.

Master SMBus interface/EPROM

During a fundamental reset, an optional serial EEPROM may be used to initialize any software visible register in the device.

Serial EEPROM loading occurs if the Switch Mode (SWMODE [3:0]) field selects an operating mode that performs serial EEPROM initialization (e.g., transparent mode with serial EEPROM initialization).

The address used by the SMBus interface to access the serial EEPROM is specified by the MSMBADDR [4:1] signals, as shown in Table 2.

Address Bit	Address Bit Value
1	MSMBADDR[1]
2	MSMBADDR[2]
3	MSMBADDR[3]
4	MSMBADDR[4]
5	1
6	0
7	1

Table 2 Serial EEPROM SMBus Address

Any serial EEPROM compatible with those listed in Table 3 may be used to store PES24N3 initialization values. Because some of these devices are larger than the total available PCI configuration space that can be initialized in the PES24N3, EEPROM space may not be fully utilized.

Serial EEPROM	Size
24C32	4 KB
24C64	8 KB
24C128	16 KB
24C256	32 KB
24C512	64 KB

Table 3 PES24N3 Compatible Serial EEPROMs

Slave SMBus interface

The slave SMBus interface provides the PES24N3 with a configuration, management, and debug interface. Using the slave SMBus interface, an external master can read or write any software visible register in the device.

When a static slave SMBus address is selected during a fundamental reset, the address is specified by the SSMBADDR [5,3:1] signals, as shown in Table 4. Refer to the PES24N3 user manual for more details on the command sequence to access the PES24N3 registers through the slave SMBus interface.

Address Bit	Address Bit Value
1	SSMBADDR[1]
2	SSMBADDR[2]
3	SSMBADDR[3]
4	0
5	SSMBADDR[5]
6	1
7	1

Table 4 Static Slave SMBus Address

Revision History

January 30, 2006: Initial publication.

March 6, 2006: Removed reference to SSMBARP and to alternate functions for GPIO [0:1] pins. Revised Figures 9 and 11.

September 20, 2007: Added reference to the entire family of PCIe 1.0a/1.1 compliant family of switches.

IMPORTANT NOTICE AND DISCLAIMER

RENESAS ELECTRONICS CORPORATION AND ITS SUBSIDIARIES (“RENESAS”) PROVIDES TECHNICAL SPECIFICATIONS AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT OF THIRD-PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for developers who are designing with Renesas products. You are solely responsible for (1) selecting the appropriate products for your application, (2) designing, validating, and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. Renesas grants you permission to use these resources only to develop an application that uses Renesas products. Other reproduction or use of these resources is strictly prohibited. No license is granted to any other Renesas intellectual property or to any third-party intellectual property. Renesas disclaims responsibility for, and you will fully indemnify Renesas and its representatives against, any claims, damages, costs, losses, or liabilities arising from your use of these resources. Renesas' products are provided only subject to Renesas' Terms and Conditions of Sale or other applicable terms agreed to in writing. No use of any Renesas resources expands or otherwise alters any applicable warranties or warranty disclaimers for these products.

(Disclaimer Rev.1.01)

Corporate Headquarters

TOYOSU FORESIA, 3-2-24 Toyosu,
Koto-ku, Tokyo 135-0061, Japan
www.renesas.com

Contact Information

For further information on a product, technology, the most up-to-date version of a document, or your nearest sales office, please visit www.renesas.com/contact-us/.

Trademarks

Renesas and the Renesas logo are trademarks of Renesas Electronics Corporation. All trademarks and registered trademarks are the property of their respective owners.