

PCI Express* Board Design Guidelines

DRAFT

Intel Corporation
June 2003



DRAFT

*Other names and brands may be claimed as the property of others.

THIS DOCUMENT IS PROVIDED “AS IS” WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification.

Intel, the Intel logo, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

§Other names and brands may be claimed as the property of others.

Copyright © 2003, Intel Corporation. All rights reserved.

This document is accessible on the Web at: <http://www.express-lane.org>

The PCI Express Base Specification and PCI Express Electromechanical Specification can be found on PCI-SIG web site: <http://www.pcisig.com>

Contents

1.	Physical Interconnect Layout Design	5
1.1	Introduction	5
1.2	Topology and Interconnect Overview	5
1.2.1	Card Interoperability	7
1.2.2	Bowtie Topology Considerations	7
1.2.2.1	Lane Polarity Inversion.....	8
1.2.2.2	Lane Reversal and Width Negotiation	8
1.3	Physical Layout Design Constraints.....	11
1.3.1	PCB Stackup.....	11
1.3.1.1	Desktop System Board and Add-in Card (4-layer) Stackup	12
1.3.1.2	Server, Workstation and Mobile (6-layer, 8-layer and 10-layer) Stackups.....	15
1.3.1.3	Add-in Card and Mobile (6-layer) Stackup	16
1.3.2	PCB Trace and Other Element Considerations	17
1.3.2.1	Differential Pair Width and Spacing Impacts	20
1.3.2.2	Differential Pair Length Restrictions and Budgets	23
1.3.2.3	Length Matching.....	24
1.3.2.4	Reference Planes.....	25
1.3.2.5	Breakout Area Specific Routing Guidelines	27
1.3.2.6	Edge Finger Design: Add-in Card	29
1.3.2.7	Via Usage and Placement	30
1.3.2.8	Bends	32
1.3.2.9	Test Points and Probing	35
1.3.3	PCI Express Topologies	35
1.3.3.1	Interconnect Topologies for Two Components on the Baseboard	36
1.3.3.2	Interconnect Topologies for Baseboard with Add-in Card	37
1.3.4	Passive Components and Connectors	38
1.3.4.1	AC Coupling Capacitors	38
1.3.4.2	Connectors	40
1.4	Summary	41

List of Figures

Figure 1-1. Conceptual Example of a PCI Express Link	6
Figure 1-2. Polarity Inversion on a TX to RX Interconnect	8
Figure 1-3: Progressive Illustration of Lane/Polarity Reversal Benefits.....	10
Figure 1-4. Bowtie Routing to Alleviate Criss - crossing RX/TX Pairs	11
Figure 1-5. Example Stackup of 4-layer PCI Express Printed Circuit Board	12
Figure 1-6. Example of Soldermask Thickness Variations on Microstrip Traces.....	13
Figure 1-7. Detailed Cross-sections of Microstrip Traces.....	14
Figure 1-8. Example Stackup of 8-layer PCI Express Printed Circuit Board	15
Figure 1-9. Example Stackup of 10-layer PCI Express Printed Circuit Board	16
Figure 1-10. Reference Stackup for 6-layer Design	17
Figure 1-11. Top View of a PCB Illustrating Fiberglass Weave Patterns in the Dielectric	20
Figure 1-12. Trace Width and Spacing Recommendations for Microstrip	21
Figure 1-13. Trace Width and Spacing Recommendations for Stripline	22
Figure 1-14. Trace Width Variation Guidelines	22
Figure 1-15. Example of Symmetrical and Non-symmetrical Lateral Routing.....	23
Figure 1-16. Examples of Etch Located Within a Pad	24
Figure 1-17. Example of a PCB Showing the Differential Pair on the Top Layer and a GND Island on the Second Layer (PWR Plane) with its Potential Stitching Vias	26
Figure 1-18. Package Pinout/Breakout Possibilities.....	27
Figure 1-19. Breakout Area Traces	28
Figure 1-20. Example Scenario of a Trace Routing Over an Anti-pad Void	29
Figure 1-21. Example of PCI Express Reference Plane Voids for a x4 PCI Express Edge Finger Interface.....	30
Figure 1-22. Placing Vias as a Pair	31
Figure 1-23. Example Illustrations of Acceptable Bends vs. “Tight Bends”	32
Figure 1-24. Bend Illustration Diagram for Coupled Differential Pairs.....	33
Figure 1-25. Routing/Bend Diagram for Uncoupled Traces into Component Pins/Pads	34
Figure 1-26. Bend Illustration Diagram for Uncoupled Section of a Differential Pair.....	35
Figure 1-27. Baseboard Chip-to-chip Topology Example	36
Figure 1-28. Baseboard Chip-to-connector and Add-in Card Example	37
Figure 1-29. Symmetrical Routing into AC Caps	39

List of Tables

Table 1-1. Card Interoperability.....	7
Table 1-2. PCI Express Microstrip Trace Routing Guidelines.....	18
Table 1-3. PCI Express Stripline Trace Routing Guidelines.....	19
Table 1-4. Example Via Usage and Trace Length Tradeoffs	32
Table 1-5. Chip-to-chip Topology Segment Definitions.....	36
Table 1-6. Baseboard Chip-to-Connector and Add-in Card Segment Definitions	37
Table 1-7. AC Coupling Capacitor Guidelines	40

1. Physical Interconnect Layout Design

1.1 Introduction

This document provides practical, common guidelines for incorporating PCI Express interconnect layouts onto Printed Circuit Boards (PCB) ranging from 4-layer desktop baseboard designs to 10-layer or more server baseboard designs. Guidelines and constraints in this document are intended for use on both baseboard and add-in card PCB designs. This includes interconnects between PCI Express devices located on the same baseboard (chip-to-chip routing) and interconnects between a PCI Express device located “down” on the baseboard and a device located “up” on an add-in card attached through a connector.

This document is intended to cover all major components of the physical interconnect including design guidelines for the PCB traces, vias and AC coupling capacitors, as well as add-in card edge-finger and connector considerations. The intent of the guidelines and examples is to help ensure that good high-speed signal design practices are used and that the timing/jitter and loss/attenuation budgets can also be met from end-to-end across the PCI Express interconnect. However, while general physical guidelines and suggestions are given, they may not necessarily guarantee adequate performance of the interconnect for all layouts and implementations. Therefore, designers should consider modeling and simulation of the interconnect in order to ensure compliance to all applicable specifications.

The document is composed of two main sections. The first section provides an overview of general topology and interconnect guidelines. The second section concentrates on physical layout constraints where bulleted items at the beginning of a topic highlight important constraints, while the narrative that follows offers additional insight.

1.2 Topology and Interconnect Overview

A general summary and overview of the topology and interconnect considerations for a PCI Express interconnect is described below. Specific layout and routing constraint details for the various components of the interconnect are given in a separate section.

The basic PCI Express topology consists of a driver or transmitter (TX) located on one device connected through a differential pair interconnect, consisting of a D+ and a D- signal, to a receiver (RX) on a second device. The PCI Express Card Electromechanical Specification uses the designation of “PET” to signify TX originating signals and “PER” to signify RX destination signals with respect to the baseboard device. A “p” or “n” is also appended to the pin name in order to represent the D+ or D- signals, respectively (example PET2p). When referring to pinout descriptions for all components, this document follows the convention used in the PCI Express Card Electromechanical Specification for uniformity.

Because the PCI Express interface is defined as dual-simplex, each connection must consist of both a transmit pair and a receiver pair; meaning that data is sent on one differential pair while data is received on a separate differential pair. Daisy chaining PCI Express devices is not allowed, as all interconnects must be point-to-point only. The differential pairs connecting a TX to an RX, and vice versa, comprise what is called a “lane” in the PCI Express specification. The transmitters and receivers that can be grouped together as a signal interface on a given device are known as a “port.” A given port and its respective lanes are grouped together to form what is known as a “link.” PCI Express links may consist of 1, 2, 4, 8, 12, 16, or 32 lanes in parallel,

*Other names and brands may be claimed as the property of others.

referred to as x1, x2, x4, x8, x12, x16, and x32 links respectively. See Figure 1-1 for an example of a link with a width of x2 (read as “by two”). Note that the AC caps required in the design are not included in this illustration.

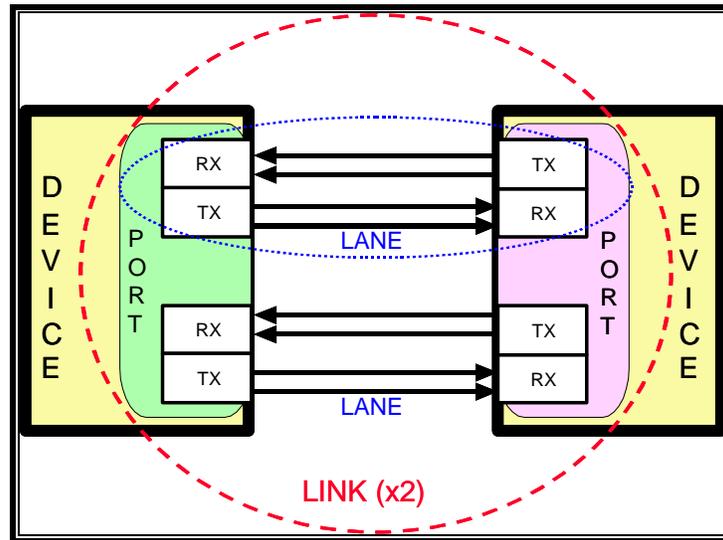


Figure 1-1. Conceptual Example of a PCI Express Link

The PCI Express specification dictates that each lane be AC coupled between its corresponding transmitter and receiver. The AC coupling is likely required to be located external to the transmitter or receiver component in the form of a discrete capacitor. Each transmitting device’s data sheet is required to inform the system designer whether or not an AC capacitor is required external to the TX component.

Each end of the link is terminated on-die into nominal 100Ω differential DC termination; therefore, no additional external termination is required for the differential pairs. Additionally, the spec requires each TX component to utilize on-die equalization by means of de-emphasis for all PCI Express signals. The de-emphasis is required to be a typical value of 3.5dB (+/-0.5dB) down with respect to the nominal output voltage. Therefore, no additional external equalization is required for the differential pairs on the PCB.

PCI Express links are formed when the TX and RX differential pairs of an “upstream” device connect to the RX and TX differential pairs of a “downstream” device. Figure 1-1 illustrates a x2 PCI Express Link connected as such. The term upstream device is used to refer to the PCI Express component that is on the end of the PCI Express Link that is hierarchically closer to the root of the PCI Express tree hierarchy. The term downstream device is used to refer to the component that is on the end of the Link that is hierarchically further from the root of the PCI Express tree hierarchy. Note that the port of the upstream device connecting to the downstream device is called the downstream port of the upstream device in the specification. Likewise the downstream device also contains an upstream port. An example of an upstream device might be a host memory controller hub chipset component, while an example of a downstream device might be a component on an add-in card.

It is important to note that each lane on a PCI Express link is assigned a sequential numerical value and is identified as such in the component’s pinout. For example, a link with a width of x2 has TX and RX pairs with respective labels that would be similar to the following: PETp0, PETn0, PERp0, PERn0, PETp1, PETn1, PERp1, PERn1.

It is also important to note that a TX located on one device must always connect to an RX on the other device. In general, this same style of connectivity must apply to each lane of a given link regardless of the link's size (in other words, with a link width x32, lane 0 and lane 31 of the transmitting device must connect to lane 0 and lane 31, respectively, of the receiving device).

1.2.1 Card Interoperability

For add-in cards, the PCI Express Card Electromechanical Specification defines an interoperability matrix between cards and slots with a variety of link widths. This is shown in Table 1-1.

Table 1-1. Card Interoperability.

Slot \ Card	x1	x4	x8	x16
x1	Required	Required	Required	Required
x4	No	Required	Allowed	Allowed
x8	No	No	Required	Allowed
x16	No	No	No	Required

The specification also defines several terms:

1. Down-plugging - plugging a larger link card into a smaller link slot (connector). This is not allowed and is physically prevented.
2. Up-plugging - plugging a smaller link card into a larger link slot (connector). This is allowed.
3. Down-shifting - plugging a card into a slot (connector) that is not fully routed for all lanes. This is not allowed with one exception. For a x8 slot (connector) on a system board, designers may choose to route only the first 4 lanes. To ensure interoperability, all x8 silicon and add-in card must support down-shifting to x4.

1.2.2 Bowtie Topology Considerations

It is quite possible that when interconnecting PCI Express devices, certain “bowtie” or signal crossing scenarios might occur when the link is physically routed on a printed circuit board. The three main types of bowtie scenarios that may exist are:

1. D+, D- crisscrossing within a pair. This scenario occurs when the D+ and D- signals of a differential pair from a transmitting device must crisscross in order to connect to the respective D+ and D- signals of the receiving device (support required – see below).
2. Crossing of transmit and receive pairs. This scenario occurs when the transmit and receive differential signal pairs must crisscross each other in order to properly connect from one device to another (not supported – see below).
3. Crossing of lanes. An example of this scenario occurs when lane 0 of the transmitting device on a x2 or greater link must crisscross the other lanes in order to connect with lane 0 on the receiving device (support optional – see below).

To help layout designers overcome scenarios #1 and #3 above, the PCI Express specification includes provisions for two features: Polarity Inversion and Lane Reversal. The specification does not include any provisions to address scenario #2, however, and any such instances of this occurrence will likely require a crisscross of the transmit and receive pairs during layout and routing (see Figure 1-4).

1.2.2.1 Lane Polarity Inversion

The PCI Express specification requires Polarity Inversion capability to be supported by each receiver on a link. The receiver accomplishes this by simply inverting the received data on the differential pair if it detects a polarity inversion during the initial training sequence of the link. In other words, a lane will still function correctly even if a positive (D+) signal from a transmitter is connected to the negative (D-) signal of the receiver. Note that the negative (D-) signal from the transmitter must now also connect to the positive (D+) signal of the receiver. Polarity Inversion can therefore become quite useful in eliminating bowties within a given differential pair. See Figure 1-2 for an example. AC caps, required in design for each signal, are not included in the illustration for clarity). It is important to note that Polarity Inversion does not imply support for “direction inversion” or “direction reversal”; that is, the TX diff pair from one device must still connect to the RX diff pair on the receiving device. In other words, a TX diff pair and an RX diff pair on a given device cannot switch functions and become an RX and TX pair, respectively. It is therefore likely that a bowtie will still need to be routed to solve the problem of crossing an RX and TX pair within a link.

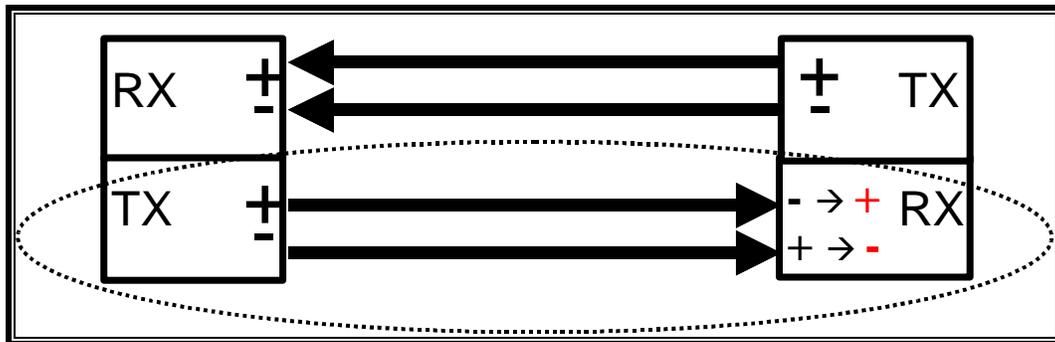


Figure 1-2. Polarity Inversion on a TX to RX Interconnect

1.2.2.2 Lane Reversal and Width Negotiation

The optionally supported feature of Lane Reversal can also help to alleviate certain bowtie scenarios that might present themselves on the link. Typically, care must be taken so that a link is properly connected between devices. For example, on a link with a width of x2, the basic rule is that lane 0 and lane 1 of the transmitting device must connect to the respective lane 0 and lane 1 signals of the receiving device. In general, this same style of connectivity must apply to each lane of a given link regardless of the link’s size. In other words, on a link with a width of x32, lane 0 and lane 31 of the transmitting device must connect to lane 0 and lane 31, respectively, of the receiving device.

However, since the PCI Express spec optionally supports Lane Reversal, potential connectivity issues between two components may be alleviated if this option is indeed implemented. Lane Reversal may be also thought of as lane re-ordering. In essence, Lane Reversal allows for the re-ordering of lanes that connect transmitting and receiving devices. For example, on a link with a width of x8, lane 0 of the transmitting device may actually now be connected to lane 7 of the receiving device, and vice versa. Of course, this would only be allowed if such support for Lane Reversal were indeed implemented by one of the devices forming the link. The caveat is that

unless a system designer can guarantee that at least one of the components on a link supports Lane Reversal, it cannot be assumed that it would be supported. For example, when routing to a x8 connector, unless the device on the system board is known to support Lane Reversal, lane 0 should always connect only to lane 0 on the connector (as opposed to allowing it to connect to either lane 0 or lane 7).

Lane width negotiation is an element of link training that is also loosely associated with Lane Reversal. The significance of lane width negotiation is that it potentially allows for a given size port on an upstream device to connect to a different size port on a downstream device.

Consider the following two examples:

1. A system board chipset component may have a PCI Express port that supports link widths of x8, x4 and x1. A separate chipset component on the system board to which it needs to connect may only support widths of x8 and x1. If the two components were connected together by routing only 4 of the available TX and RX pairs from the first chipset to the second, then the two devices would in fact negotiate to a width of x1. The additional 3 TX and RX pairs connecting the two devices would then actually be deactivated. This is because the PCI Express Specification requires that all components must be at least capable of forming a x1 link as well as their maximum link size.
2. A system board chipset component may have a PCI Express port that supports link widths of x12, x8, x4 and x1. A separate chipset component on the system board to which it needs to connect may only support widths of x12, x4 and x1. If the two components were connected together by routing only 8 of the available TX and RX pairs from the first chipset to the second, then the two devices would in fact negotiate to a width of x4 since the second device doesn't support a width of x8 as does the first device. The additional 4 TX and RX pairs connecting the two devices would then actually be deactivated. However, if a new generation (but same pinout/footprint) of the second chipset was used on the system board and it supported a width of x8, the same system board would now operate at the x8 width!

It is important to note that for lane width negotiation to be beneficial, however, the system designer must determine—by means of each component's data sheet—what link width sizes are potentially supported and route the bus accordingly. Otherwise, available bandwidth may unintentionally go unused, especially if the interconnect defaults to a x1 link as illustrated in Example #1, above.

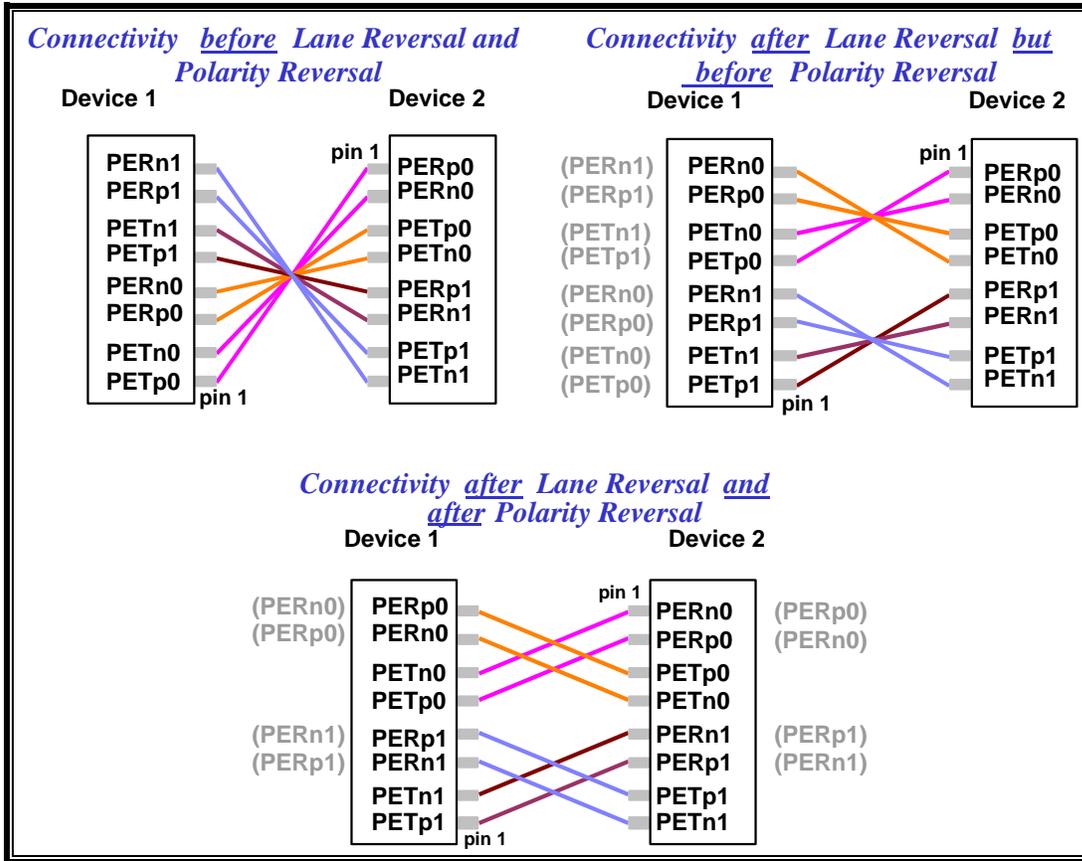


Figure 1-3: Progressive Illustration of Lane/Polarity Reversal Benefits

Important note: As with Polarity Inversion, Lane Reversal does not imply direction reversal. The TX diff pair from an upstream device must still connect to the RX diff pair on the downstream device.

A comprehensive illustration showing the progressive effects of both how Polarity Reversal and Lane Reversal can help alleviate bowtie scenarios is shown in Figure 1-3. In the figure, the previous pin name shown with parenthesis in light colored text highlights progressive changes. In this example, bowties still exist between the crisscrossed pairs of individual TX and RX differential signals. Figure 1-4 illustrates how this scenario might be accommodated through the use of vias and signal layer transitions. Although AC caps are required on all signals, they are not included in the Figure 1-3 and Figure 1-4 illustrations.

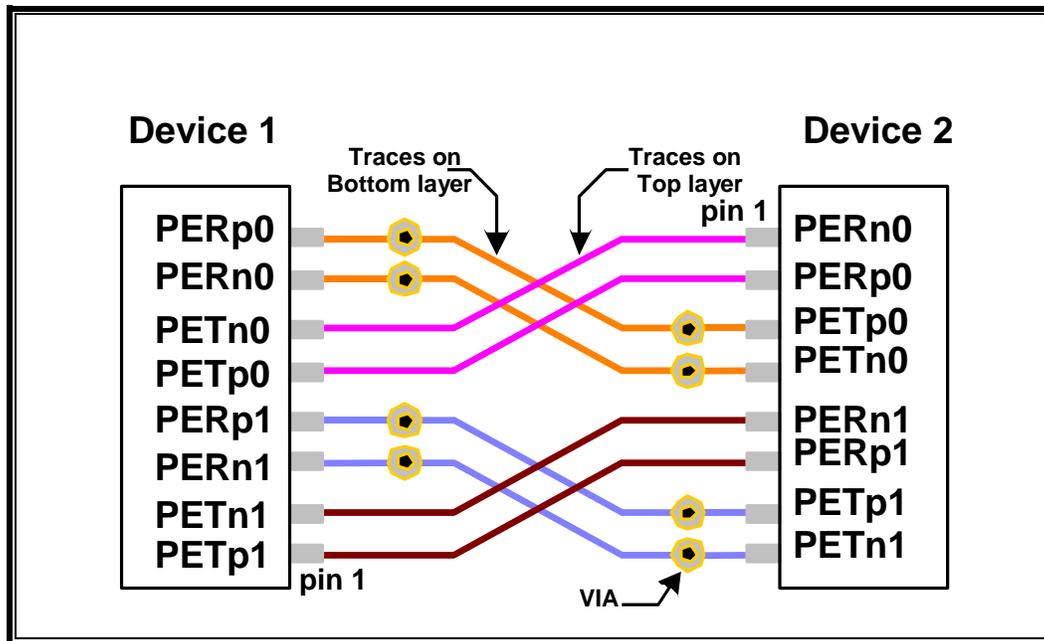


Figure 1-4. Bowtie Routing to Alleviate Criss - crossing RX/TX Pairs

1.3 Physical Layout Design Constraints

The physical layout of the PCI Express interconnect requires consideration of a number of different components that form the complete system design. This section focuses on the PCB stackup (1.3.1), PCB trace and other element considerations (1.3.2), PCI Express topologies (1.3.3), passive components and connectors (1.3.4).

1.3.1 PCB Stackup

This section presents the considerations associated with the printed circuit board and its stackup. Specific routing requirements are provided in each of the following sub-sections. Testpoint and probing considerations are not covered in this document. Below is an overview of some general stackup considerations.

- Along with other considerations, the particular stackup used for a given design generally depends upon cost, routing density, and signal referencing requirements
- Desktop system boards designed with 4-layer stackups due to cost constraints typically utilize microstrip trace routing (on the outer layers).
- Enterprise Server, Workstation and Mobile system boards commonly use 6-layer, 8-layer or 10-layer stackups utilizing microstrip and stripline trace routing.
- Add-in cards are commonly built as 4-layer microstrip designs, although other multi-layer stackups with stripline routing are also prevalent (e.g. 6 or more layer designs).
- Microstrip differential traces inherently display greater impedance variation than stripline traces.
- Thicker microstrip and stripline copper traces demonstrate less skin effect loss, while thinner traces demonstrate more skin effect loss.
- Wider traces demonstrate less loss.

- Plating of copper traces with non-copper materials (e.g. tin, nickel, etc.) will increase the relative loss vs. a pure copper trace of the same thickness.
- Thicker dielectrics demonstrate less loss.

1.3.1.1 Desktop System Board and Add-in Card (4-layer) Stackup

Though each system board and add-in card is different, many desktop boards and add-in cards often have a similar stackup or composition. An example stackup of a 4-layer printed circuit board is shown in Figure 1-5. This stackup assumes a nominal, overall board thickness of 0.062 inches, which is typical for most system boards and is in fact required for PCI Express add-in cards. Trends in manufacturing variances are also listed in the figure, though they are to be used as guidelines only. **All system board and add-in card layout recommendations presented in this document are based upon the example stackups shown in Figure 1-5, Figure 1-8, Figure 1-9 and Figure 1-10. Any major deviations from these stackups will impact and alter the routing guidelines presented later in this document.** Edge-finger layouts are located in section 1.3.2.6, and connector considerations are located in section 1.3.4.2 of this document.

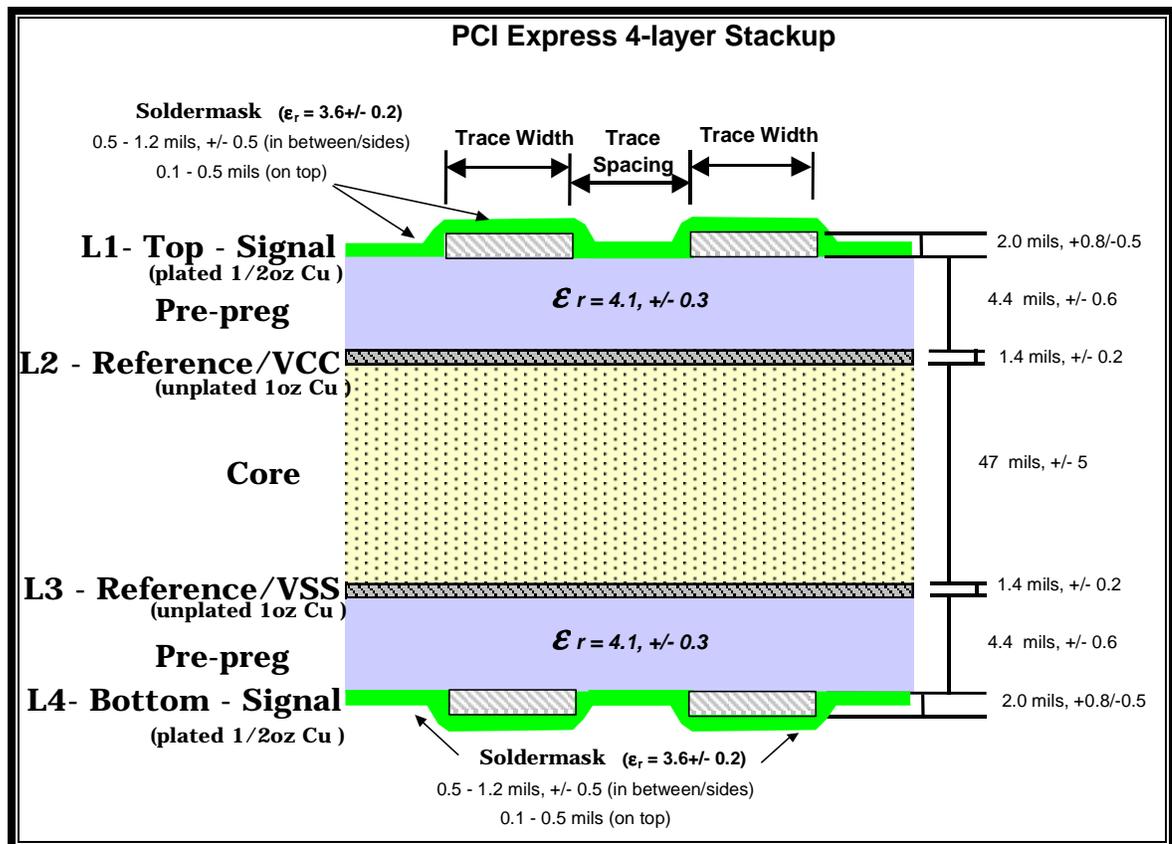


Figure 1-5. Example Stackup of 4-layer PCI Express Printed Circuit Board

Key aspects of the 4-layer example stackup of Figure 1-5 are as follows:

- Variations in the stackup of a system board such as changes in the dielectric height, trace widths, and spacing can impact the impedance, loss, and jitter characteristics of a PCI Express interconnect.
 - Note that geometry variations often occur due to PCB manufacturing process tolerances.
- Components (including AC coupling caps) are typically limited to single-sided placement on the top layer due to increased manufacturing costs associated with double-sided placement. This is especially true for High Volume Manufacturing (HVM) environments.
- Due to differences among various add-in card and system board manufacturers, the power and ground plane locations may be interchanged on a given board from what is shown in the figure above. Designers should follow the reference plane guidelines in Section 1.3.2.4, however, regardless of the plane locations.
- Signals are generally routed as microstrip traces on layers L1 and L4, and are subject to soldermask, plating, etching, and dielectric fiberglass bundle weave impacts:
 - Plating and etching have a greater impact on the impedance tolerance of differential pair signals than they do for signals consisting of a single trace.
 - Soldermask variation will also impact the impedance of microstrip traces, though to a lesser extent, as its application to the PCB is generally not uniform. For example, a thinner layer of soldermask usually exists immediately above the copper traces (e.g. ~0.5 mils), while a thicker layer is present next to and between the sides of the copper traces (e.g. almost equal to the thickness of the traces themselves). Please refer to Figure 1-6.

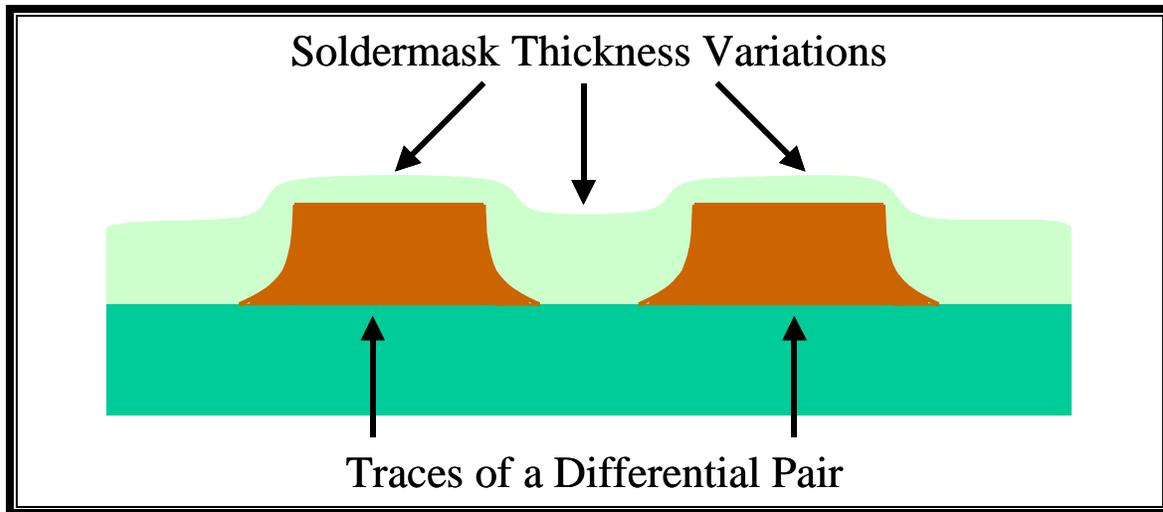


Figure 1-6. Example of Soldermask Thickness Variations on Microstrip Traces

- Variations in plating can negatively impact the differential impedance tolerance achievable. Please refer to situation #1, Figure 1-7.
- Over or under-etching of microstrip traces will alter the exact trace dimensions and subsequently impact the resulting differential impedance of coupled pairs. Please refer to situation #1, Figure 1-7.
- The varying alignment of the two traces of a differential pair over the fiberglass bundles and/or resin-rich areas (or resin troughs) in the PCB dielectric will

impact the resulting differential impedance and can contribute to an increase in the amount of common mode voltage found on the differential pair. Please refer to situation #2, Figure 1-7.

- Stackups accommodating stripline routing (e.g. 6- or more-layer stackups, etc.) can produce a better differential impedance tolerance than microstrips due in part to the absence of soldermask and plating. Impacts due to etching variations are also somewhat reduced. However fiberglass bundle weave impacts may have more impact on striplines than microstrip traces.

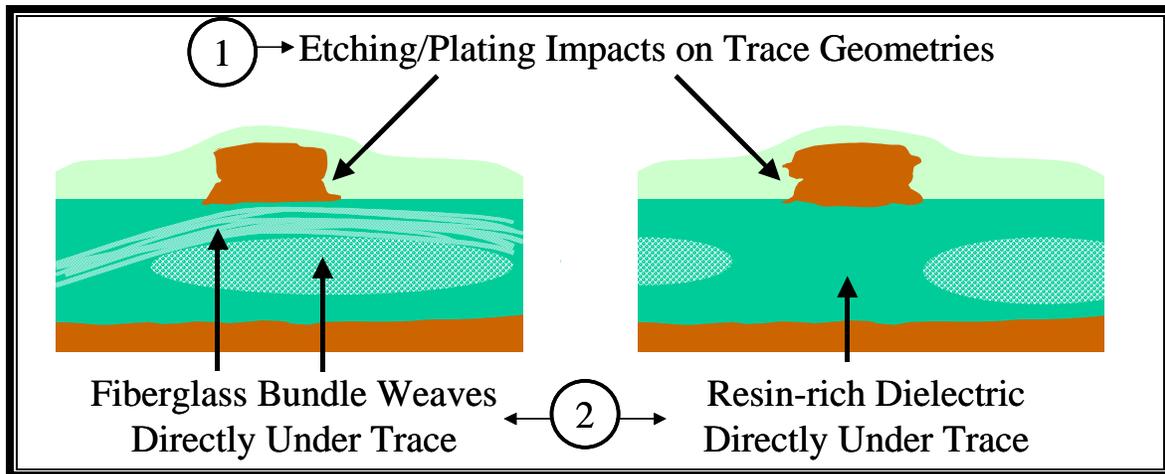


Figure 1-7. Detailed Cross-sections of Microstrip Traces

- Current high-volume manufacturing technology commonly restricts the minimum trace widths to 5 mils and the minimum spacing to a 5-mil air gap between traces.
 - Specific trace width and spacing targets for PCI Express interconnects can be found later in this document in the “PCB Trace and Other Element Considerations” (section 1.3.2); such targets are intended to reach a compromise between impedance targets, loss impacts, crosstalk immunity and routing flexibility.
- Common pre-preg or laminate material used between layers L1/L2 and layers L3/L4 is typically one sheet of “2116”.
 - Note that “2116” is the descriptor used by PCB manufactures that identifies a particular type of pre-preg/laminate with a given resin to glass ratio. This specific laminate will typically produce the nominal thickness and dielectric constant described in the stackup of Figure 1-5.
 - Note that thinner dielectrics (such as 1080, for example) typically have less-dense glass weave patterns, aggravating situation #2 illustrated in Figure 1-7.
- Several sheets of “7628” are commonly used to form the core.
 - Note that “7628” is the descriptor used by PCB manufactures that identifies a particular type and thickness of pre-preg/laminate with a given resin to glass ratio. This laminate is typically used for rigid or semi-rigid stackups.

1.3.1.2 Server, Workstation and Mobile (6-layer, 8-layer and 10-layer) Stackups

The stackups for Servers, Workstations and Mobile platforms may require more layers than the Desktop platforms. The Server and Workstation stackups typically, as with Desktop, are 0.062 inches overall thickness (some may also be targeted to 0.090 inches, but examples are not provided in this document) whereas the Mobile stackup may be 0.050 inches. 6-layer, 8-layer and 10-layer stackups will be more common for these platforms, but lower layer count stackups will work as long as all placement and electrical requirements are met. Examples of 8-layer and 10-layer stackups are shown in Figure 1-8 and Figure 1-9 respectively. A 6-layer stackup, also used for add-in cards, is shown in Figure 1-10. Routing will typically include both stripline and microstrip layers. Greater component placement and routing densities dictate more layers and smaller trace geometries (4-mil or 5-mil trace widths), but the additional layers and narrower traces may add cost. The narrower traces increase the loss per unit length, hence the total routable length decreases (if using 4-mil wide traces, for example). In addition, the denser layer stackup with thinner dielectric spacing may reduce the cost-effective impedance target (instead of 100 Ohms, the impedance target may be lower at 92 Ohms or 85 Ohms for example). The key is to design to a target impedance and small enough tolerance to minimize impact to loss and jitter budgets.

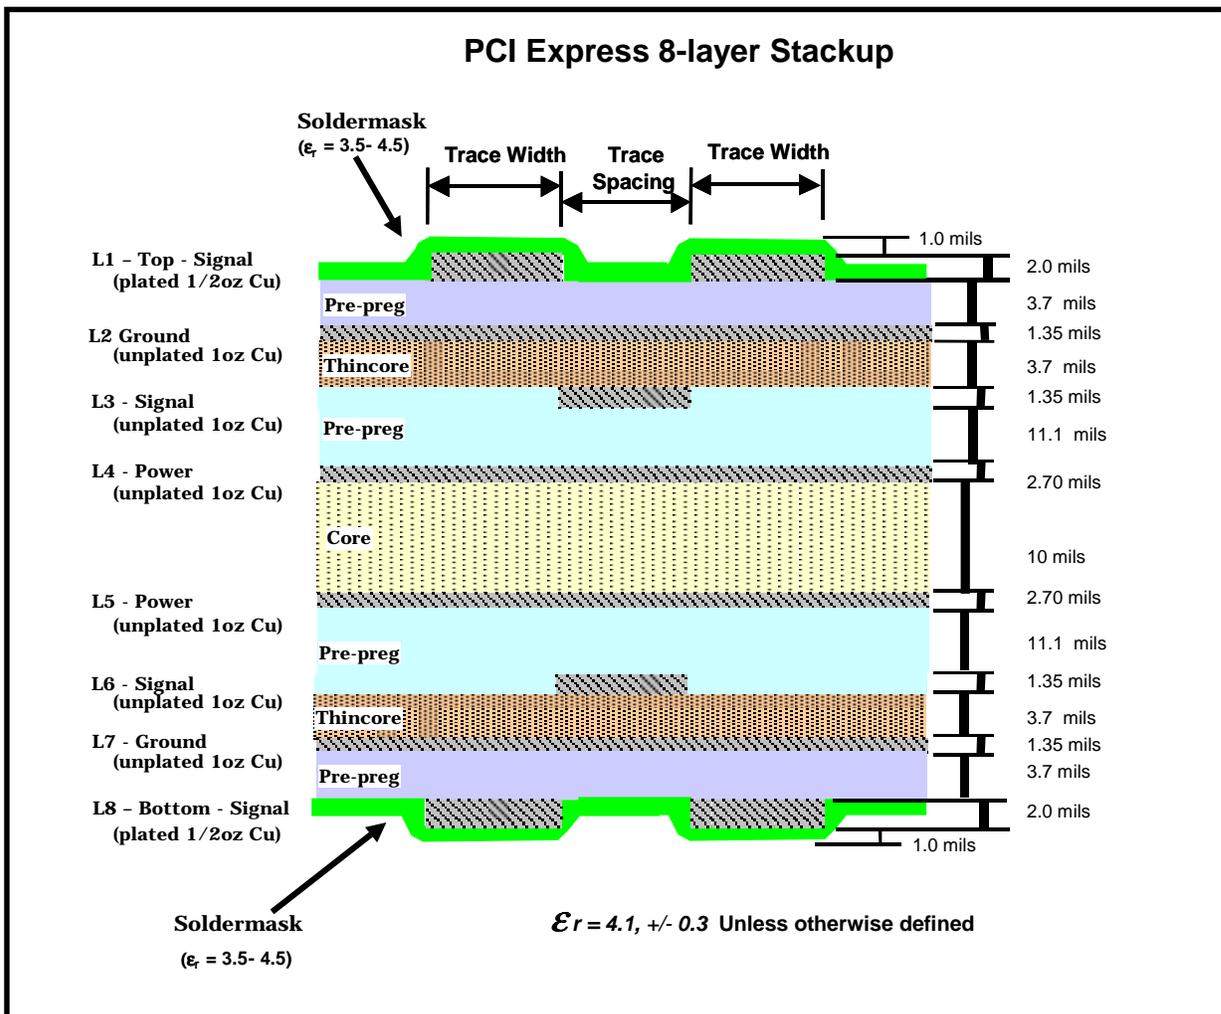


Figure 1-8. Example Stackup of 8-layer PCI Express Printed Circuit Board

*Other names and brands may be claimed as the property of others.

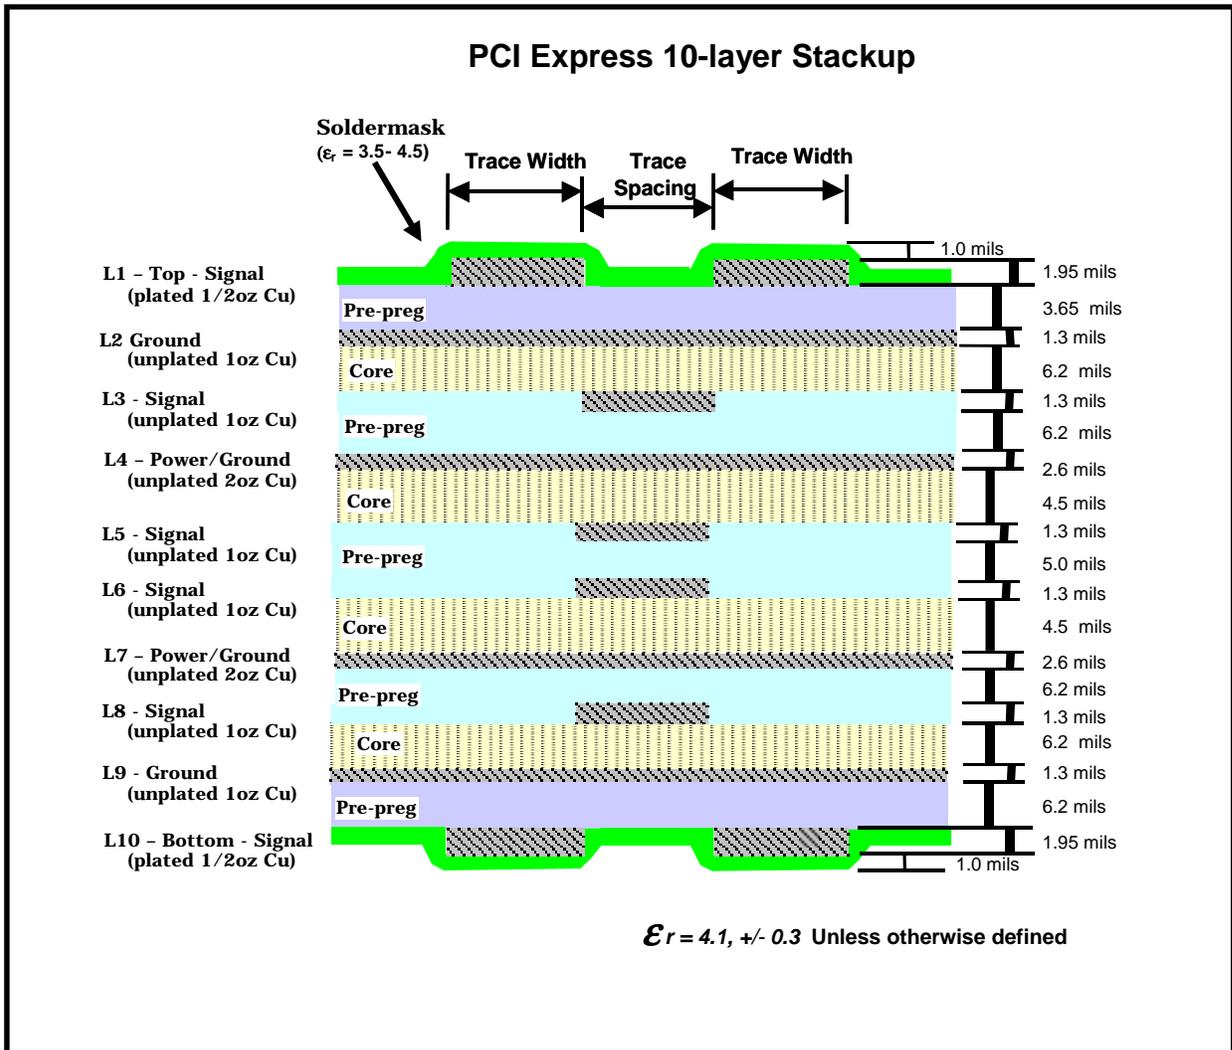


Figure 1-9. Example Stackup of 10-layer PCI Express Printed Circuit Board

1.3.1.3 Add-in Card and Mobile (6-layer) Stackup

Add-in cards and Mobile platforms (as well as other platforms) may use the 6-layer stackup in addition to other layer count stackups. Add-in cards that are plugged into PCI Express specification compliant edge card connectors must conform to an overall thickness of 0.062 inches, whereas Mobile platforms may have a thickness of 0.050 inches as well. The 6-layer stackup is shown in Figure 1-10 with microstrip trace geometries and tolerances matching that of the 4-layer stackup and including an additional two stripline layers.

All PCI Express traces for add-in cards are expected to be primarily routed as microstrip traces even with 6-layer stackups, but stripline routing may be used as well. Mobile and other platforms typically will use both layer types for routing.

Due to differences among various card manufacturers, the location of the power and ground planes may be interchanged, though PCI Express traces should still always be ground referenced, likely requiring ground floods on the VCC plane.

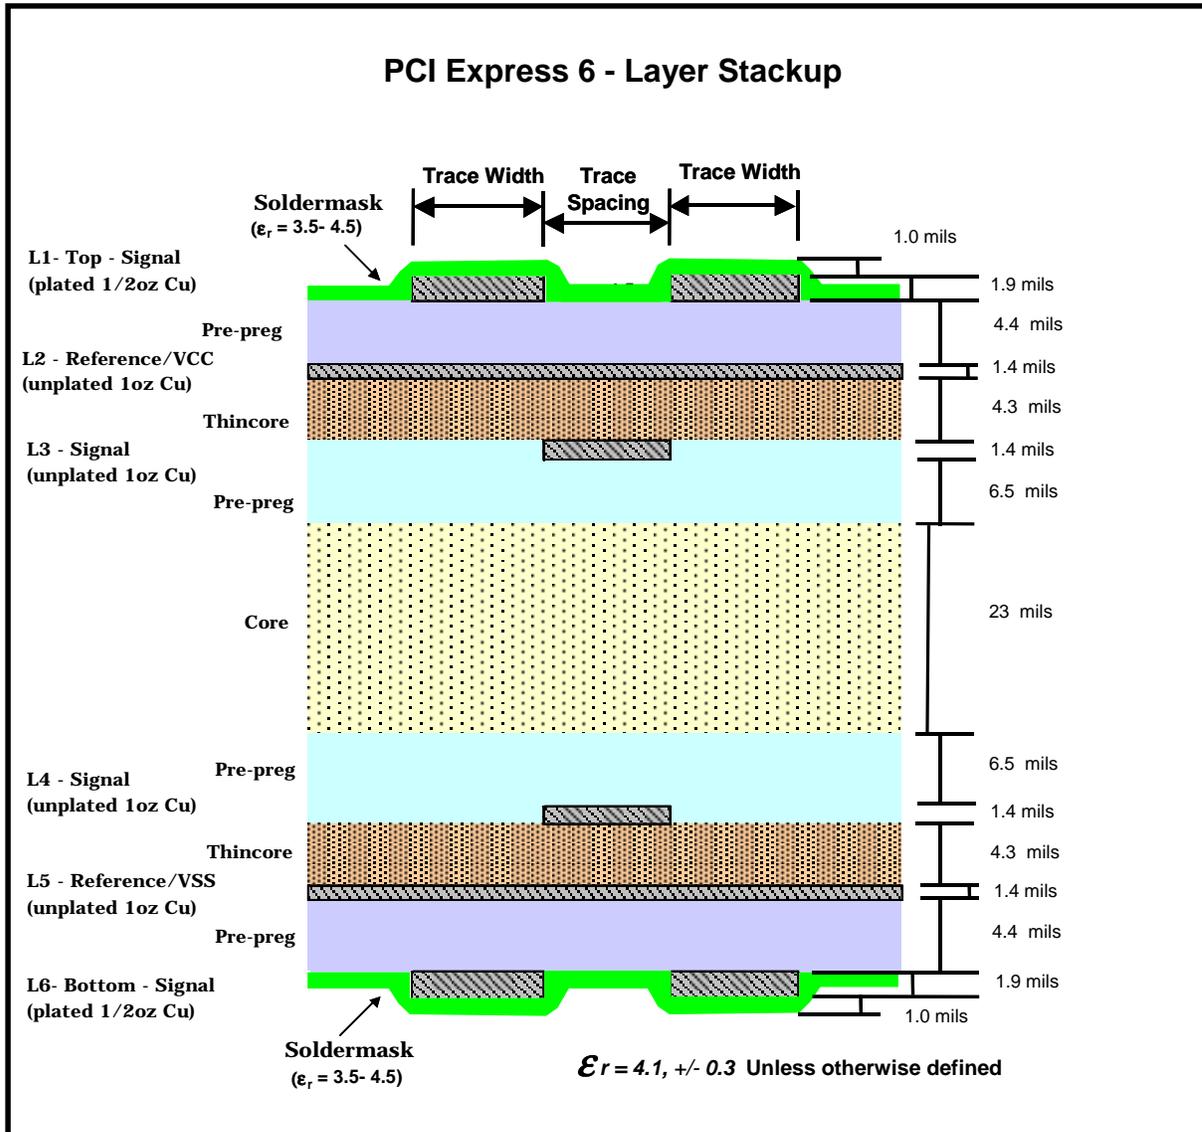


Figure 1-10. Reference Stackup for 6-layer Design

1.3.2 PCB Trace and Other Element Considerations

This section gives an overview of the requirements for the Printed Circuit Board trace-related aspects of the interconnect. These guidelines are derived from the stackups in the previous sections. Because these are general guidelines only, all interconnect paths must still be simulated to ensure proper performance and compliance with the loss and jitter budgets of the system.

While details are explained in subsequent sections, a summary of PCB-related routing criteria for PCI Express differential pairs is given in Table 1-2 and Table 1-3.

Table 1-2. PCI Express Microstrip Trace Routing Guidelines.

PARAMETER <i>(COUPLED TRACES)</i>	BASEBOARD – MAIN ROUTE	ADD-IN CARD – MAIN ROUTE	BREAKOUT AREA RULES/ ALLOWANCES
Differential Impedance	4, 6 layer: 100Ω +/- 20% ¹ 8, 10 layer: 85Ω +/- 20% ¹	SAME	
Single-ended Impedance	4, 6 layer: 60Ω +/- 15% ¹ 8, 10 layer: 55Ω +/- 15% ¹	SAME	
Trace Width	5 mils	SAME	7 mils if un-coupled length is >=100 mils
Intra-pair Spacing (within pair)	7 mils	SAME	5 mils gap is OK
Inter-pair Spacing (between pairs)	Greater of >= 20 mils or 4x dielectric height	SAME	>= 10 mils OK (TBD)
Chip to Chip Length Restrictions	15.5-inch max ² (TBD)	Not applicable	Breakout etch less than 150 mils ² (TBD) (part of total length requirement)
Chip to Card Length Restrictions	Mobile: 7-inch max ³ (TBD) All others: 12-inch max ² (TBD)	Mobile: 4-inch max ³ (TBD) All others: 4-inch max ² (TBD)	Breakout etch less than 150 mils ² (TBD) (part of total length requirement)
Length Matching Intra - pair	<ul style="list-style-type: none"> Max 5 mil delta Matching maintained segment to segment Match at point of discontinuity, but avoid “tight bends” 	SAME	Match within a 5 mil delta within Breakout Area
Length Matching Inter - pair	<ul style="list-style-type: none"> No strict electrical requirements Desirable to keep differences within a 3-inch delta to minimize latency 	<ul style="list-style-type: none"> No strict electrical requirements Want to keep within 3-inch delta to minimize latency 	
Reference Plane	<ul style="list-style-type: none"> GND Referenced preferred (stitching caps with PWR) >=20 mil trace edge to plane edge gap GND Stitching vias required by signal vias for layer changes 	SAME	
Splits/Voids	<ul style="list-style-type: none"> No routing over splits No routing over voids 	SAME	No more than ½ trace width routed over via antipad
Via Usage	<ul style="list-style-type: none"> 4 vias per TX trace 2 vias per RX trace (6 vias total, entire path) 	<ul style="list-style-type: none"> 4 vias per TX trace 2 vias per RX trace (6 vias total, entire path) 	Limit of one via OK in each breakout area
Via Size	<ul style="list-style-type: none"> Pad <=25 mils Finished hole <=14 mils 	SAME	
Bends	Match left/right turn bends where possible <ul style="list-style-type: none"> No 90-degree bends or “tight” bend structures 	SAME	Avoid “tight bends” when matching length in breakout areas

¹ Targets do not apply to Mobile with 0.050 inch thick stackup which may be closer to 92 Ohms.
² Preliminary from simulations and subject to change.
³ Mobile 4-mil trace applications. Preliminary from simulations and subject to change.
TBD – To Be Determined

Table 1-3. PCI Express Stripline Trace Routing Guidelines

PARAMETER (<i>COUPLED TRACES</i>)	BASEBOARD – MAIN ROUTE	ADD-IN CARD – MAIN ROUTE	BREAKOUT AREA RULES / ALLOWANCES
Differential Impedance	6 layer: 100Ω +/- 15% ¹ 8, 10 layer: 85Ω +/- 15% ¹	SAME	
Single-ended Impedance	6 layer: 60Ω +/- 15% ¹ 8, 10 layer: 55Ω +/- 15% ¹	SAME	
Trace Width	4 – 5 mils	SAME	7 mils if un-coupled length is >=100 mils
Intra-pair Spacing (within pair)	5 – 6 mils	SAME	5 mils gap is OK
Inter-pair Spacing (between pairs)	Greater of >= 15 mils or 3x dielectric height	SAME	>= 10 mils OK (TBD)
Chip to Chip Length Restrictions	4 mil traces: 11-inch max ² (TBD) 5 mil traces: 18-inch max ² (TBD)	Not applicable	Breakout etch less than 150 mils ² (TBD) (part of total length requirement)
Chip to Card Length Restrictions	<ul style="list-style-type: none"> Mobile: 4 mil traces: 7-inch max³ (TBD) All others: 4 mil traces: 5-inch max² (TBD) 5 mil traces: 12-inch max² (TBD) 	<ul style="list-style-type: none"> Mobile: 4-inch max Edge finger to Chip³ (TBD) All others: 4-inch max Edge finger to Chip² (TBD) 	Breakout etch less than 150 mils ² (TBD) (part of total length requirement)
Length Matching Intra-Pair	<ul style="list-style-type: none"> Max 5 mil delta Matching maintained segment to segment Match at point of discontinuity, but avoid “tight bends” 	SAME	Match within a 5 mil delta within Breakout area
Length Matching Inter-Pair	<ul style="list-style-type: none"> No strict electrical requirements Desirable to keep differences within a 3-inch delta to minimize latency 	<ul style="list-style-type: none"> No strict electrical requirements Want to keep within 3-inch delta to minimize latency 	
Reference Plane	<ul style="list-style-type: none"> GND Referenced preferred (stitching caps with PWR) >=20 mil trace edge to plane edge gap GND Stitching vias required by signal vias for layer changes 	SAME	
Splits/Voids	<ul style="list-style-type: none"> No routing over splits No routing over voids 	SAME	No more than ½ trace width routed over via antipad
Via Usage	<ul style="list-style-type: none"> 4 vias per TX trace 2 vias per RX trace (6 vias total, entire path) 	<ul style="list-style-type: none"> 4 vias per TX trace 2 vias per RX trace (6 vias total, entire path) 	Limit of one via OK in each breakout area
Via Size	<ul style="list-style-type: none"> Pad <=25 mils Finished hole <=14 mils 	SAME	
Bends	<ul style="list-style-type: none"> Match left/right turn bends where possible No 90-degree bends or “tight” bend structures 	SAME	Avoid “tight bends” when matching length in breakout areas

¹ Targets do not apply to Mobile with 0.050 inch thick stackup which may be closer to 92 Ohms.
² Preliminary from simulations and subject to change.
³ Mobile 4-mil trace application. Preliminary from simulations and subject to change.

1.3.2.1 Differential Pair Width and Spacing Impacts

PCI Express links are composed of differential pairs on the PCB. The following is a summary of general routing guidelines for the width and spacing of the differential pair traces. These guidelines help ensure that designs meet impedance and crosstalk targets.

- For microstrip traces, an impedance target of 100Ω with a max 20% impedance tolerance is desired for 4-layer and 6-layer printed circuit board stackups. For 8-layer and 10-layer stackups, an 85Ω with a max 20% impedance tolerance is recommended. It is recommended that the impedance be specified to the PCB fab vendor in terms of single ended impedance (e.g. $\sim 60\text{ Ohms } \pm 15\%$).
- For stripline traces, an impedance target of 100Ω with a max 15% impedance tolerance is desired for 6-layer printed circuit board stackups. For 8-layer and 10-layer stackups, an 85Ω with a max 15% impedance tolerance is recommended. It is recommended that the impedance be specified to the PCB fab vendor in terms of single ended impedance (e.g. $\sim 55\text{ Ohms } \pm 15\%$).
- Coupling (of the intra-pair differential signals), or relative close physical spacing within the differential pair, and increased spacing to neighboring differential pairs and other signals helps to minimize harmful crosstalk impacts and EMI effects.
- Where possible, trace routes of long distances should be routed at an off-angle to the X-Y axis of a PCB layer to distribute the effects of fiberglass bundle weaves and resin-rich areas of the dielectric (TBD). See Figure 1-11 for an illustration of diagonal routing.

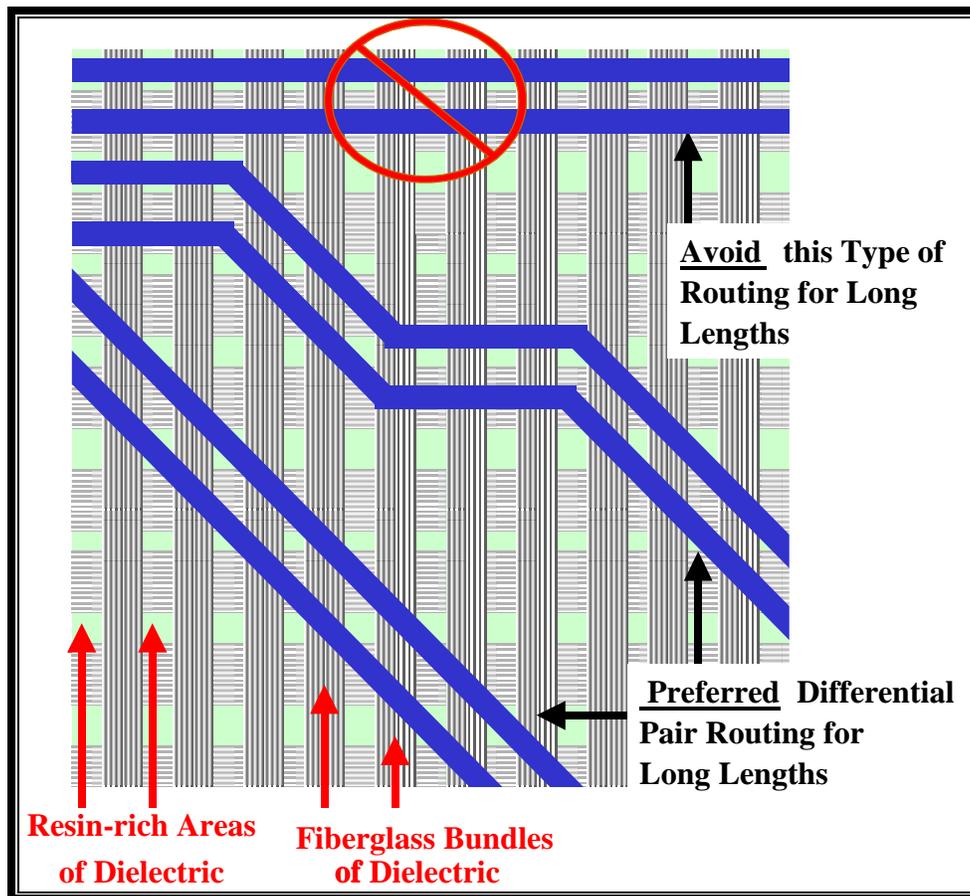


Figure 1-11. Top View of a PCB Illustrating Fiberglass Weave Patterns in the Dielectric

- All edge-coupled microstrip traces of a differential pair should be 5 mils wide with a 7-mil wide air gap spacing between the traces of the pair (also referred to as 5-on-7 or 5/7/5 routing). See Figure 1-12. Note that the 5/7/5 routing choice has been found to be the best compromise between impedance targets, loss impacts, crosstalk immunity and routing flexibility. This geometry actually produces a nominal impedance slightly below 100 Ohms.
- All edge-coupled stripline traces of a differential pair should be 5 mils wide with a 6-mil wide dielectric spacing between the traces of the pair. 4-mil width is acceptable, but limits the overall trace length due to the greater loss per unit length of routed trace.
- A 20-mil edge-to-edge gap (or 4x dielectric height, whichever is greater) should be kept between traces (inter-pair) of microstrip adjacent pairs. Please see Figure 1-12.
- A 15-mil edge-to-edge gap (or 3x dielectric height, whichever is greater) should be kept between traces (inter-pair) of stripline adjacent pairs. Please see Figure 1-13.
- Spacing to all non-PCI Express signals should be at least 20 mils; this should be increased to upwards of 30 mils in order to avoid harmful coupling issues if the non-PCI Express signals have significantly higher voltage levels and/or edge rates.
- While impedance control is important, the loss and jitter performance requirements are the primary goal – the impedance target needs to be selected to fit the printed circuit board stackup manufacturability and cost requirements while maintaining the loss and jitter requirements.

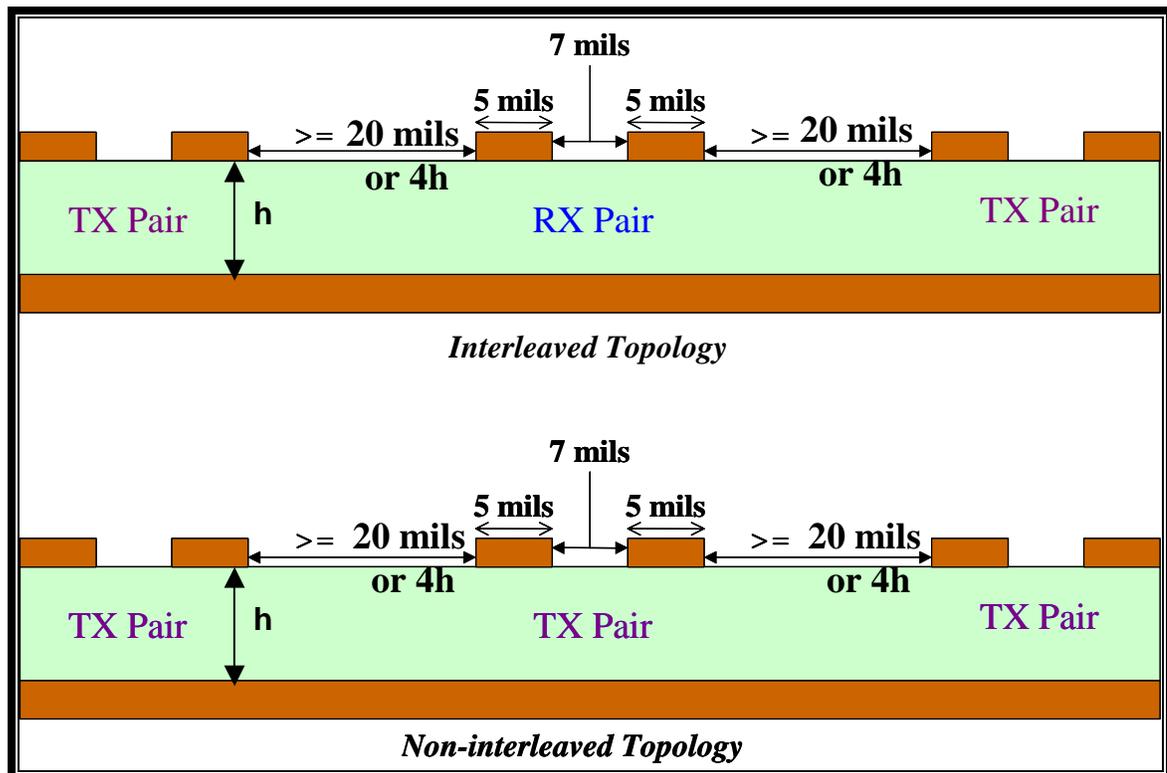


Figure 1-12. Trace Width and Spacing Recommendations for Microstrip

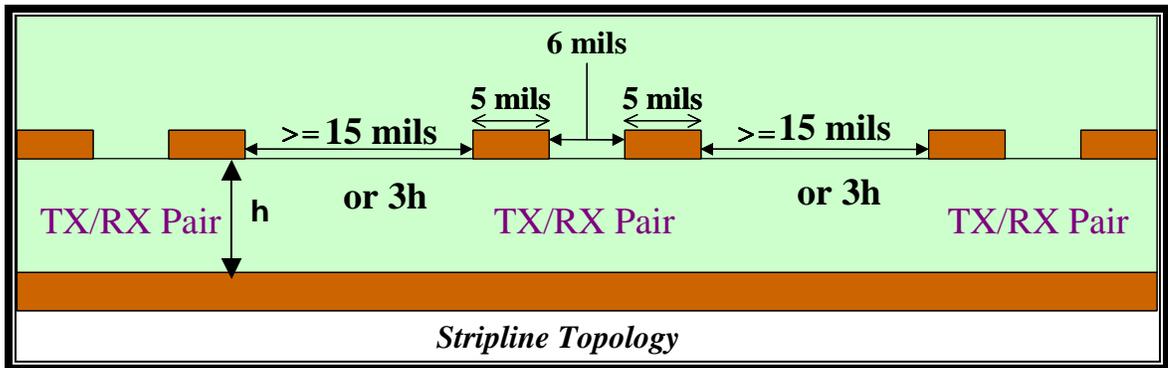


Figure 1-13. Trace Width and Spacing Recommendations for Stripline

- Any uncoupled sections (e.g. a trace within a differential pair that exceeds the intra pair spacing rule of 7 mils) may be routed as a 7-mil wide trace (vs. 5 mils) when the distance of the violation in air gap spacing is 100 mils or more. Please see Figure 1-14.

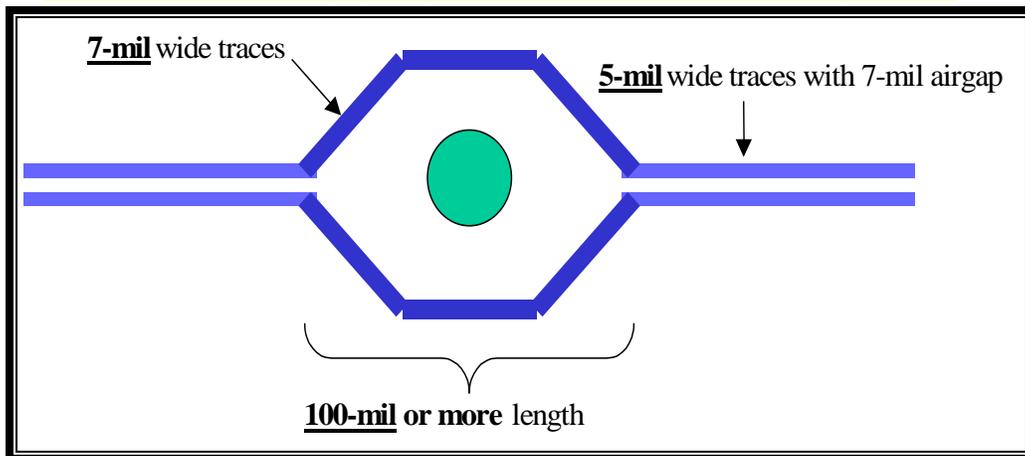


Figure 1-14. Trace Width Variation Guidelines

- Differential pairs should maintain lateral routing symmetry between the two signals of a differential pair whenever possible. Please see Figure 1-15.

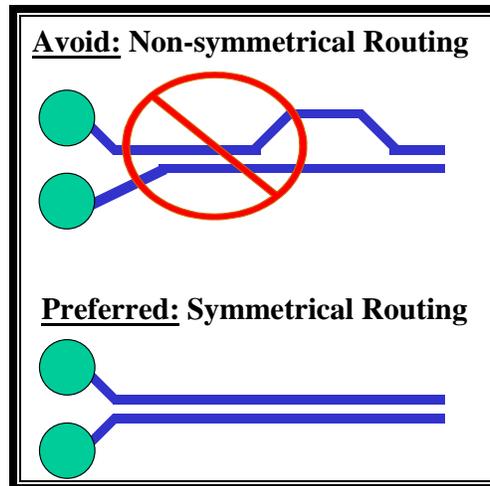


Figure 1-15. Example of Symmetrical and Non-symmetrical Lateral Routing

It's also important to note that even though we designate these as coupled pairs, there is still a large amount of coupling to the adjacent reference plane. Also, edge-coupled traces inherently force the current concentration to the extreme edge of the trace where it couples with its designated signal pair. Additionally, plating differences and etch variation from trace to trace can greatly effect the coupling characteristics of each pair. Thus, while a 5-on-5 differential pair may intuitively seem the best choice, in order to improve impedance continuity on the link, 5-on-7 routing actually will achieve an impedance closer to 100Ω .

1.3.2.2 Differential Pair Length Restrictions and Budgets

Trace length restrictions for the PCI Express interface depend on a variety of factors, many of which are described throughout this document. Trace lengths greatly impact the loss and jitter budgets of the interconnect. Though maximum trace routing lengths are given below, shorter lengths may be used, and subsequent tradeoffs can be made with the other restrictions placed on the interconnect, such as vias, spacing, and so forth. However, whenever tradeoffs are made, the new topology must be simulated to ensure compliance with the loss and jitter budgets of the system. The following are general trace length routing recommendations and restrictions based on simulations (please note that these are preliminary and may change):

- Keep trace routing lengths as short as possible.
- Trace routing lengths from chip to chip are limited to a maximum of 15.5 inches (TBD) for microstrip and 18 inches (TBD) for stripline. With simulation, specific stackup implementations may yield longer lengths.
 - Lengths are measured pin to pin. This includes any trace lengths in breakout regions as well. Package types at both ends of the interconnect path may impact this length. For this guideline, one wirebond and one flip-chip package are assumed.
- Trace routing lengths from a chipset to a connector (microstrip or stripline) are limited to 12 inches (TBD). With simulation, specific stackup implementations may yield longer lengths.
 - Lengths are measured pin to pin. This includes any trace lengths in breakout regions as well. Package types at both ends may impact this length.
- Trace routing lengths from the edge-finger pad edge to the chipset on an add-in card are limited to 4 inches (TBD).

- This includes any trace lengths in breakout regions as well (microstrip or stripline). Note that this length guideline assumes a Flip-chip package. Other package types may impact this length.
- Note that PCB traces may introduce as much as 1 ps or more of jitter per inch per differential pair, based on inherent physical properties of the PCB differential pairs.
 - Important note: other factors such as data patterns, edge rates, impedance discontinuities, crosstalk, etc. will also contribute to system jitter.
- Note that PCB traces may introduce as much as 0.25 to 0.35 dB of loss per inch per differential pair, based on inherent physical properties of the PCB differential pairs.

Depending upon whether routed as microstrip or stripline, the trace width, number of vias and other factors, the allowed routing length budget will vary. The guideline tables can be used directly, but since individual topologies will vary, these guidelines can also be used as a reference for setting initial expectations while determining specific topology lengths based on either the chipset platform design guide or simulation based on this document, or both.

1.3.2.3 Length Matching

Use the following length matching guidelines for PCI Express differential pairs. Exceptions to these guidelines may occur in the breakout regions of the PCB (please refer to Section 1.3.2.5):

- Match the length of both nets of the differential pair, allowing no more than a 5-mil delta between the lengths of the two signals (TBD).
 - Important note: The pad or pin edge-to-edge distance should be used vs. the total etch present when determining the overall length of a given net in a differential pair. This is due to the fact that although the amount of etch within a given pad is electrically part of the pad itself, most CAD tools count this length as part of the total etch. However, if the amount of etch routing into each and every pad or pin is identical, then the total etch measurement may be used. Please refer to Figure 1-16.

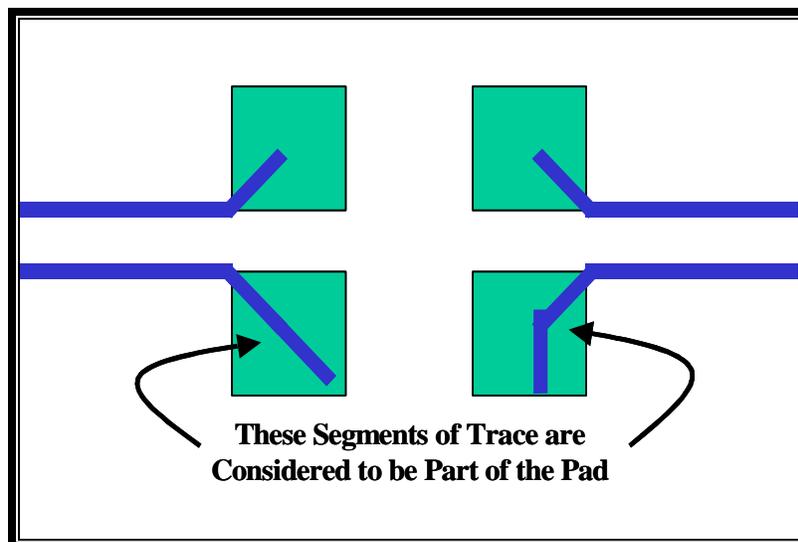


Figure 1-16. Examples of Etch Located Within a Pad

- Length matching should occur on a segment-by-segment basis vs. across the total distance of the overall route.

- Segments include the route between two via transitions, between the pin of the component and the AC capacitor pad, or between the capacitor pad and the connector pin or edge finger.
 - Whenever possible length matching should occur at the immediate point of mismatch or discontinuity.
- Important note: “Tight bends” should be avoided when trying to match lengths between nets of a differential pair. See Section 1.3.2.8 for more details.
- Length matching from one differential pair to another is not a strict requirement due to the large pair-to-pair skew allowed at a PCI Express receiving component; however, it is recommended that the following guidelines be used:
 - All differential pairs should be length matched to within a maximum of 7 inches of one another on the system board, though a smaller delta (~0 - 3 inches) is generally preferred.

Pair to pair length matching on the add-in card should be kept within a maximum delta of 3 inches.

1.3.2.4 Reference Planes

The concept of reference planes is important for PCI Express traces even though they are routed as coupled differential pairs. Since a large amount of coupling to the reference plane still occurs for microstrip differential pairs, the following guidelines are recommended:

- Traces should avoid discontinuities in the reference plane, such as splits and other voids.
 - Good reference planes help to minimize any AC Common mode voltage found on the differential pair as well as benefiting signal quality and EMI impacts.
- When routing near the edge of their reference plane, traces should adhere to the following:
 - Maintain at least a 20-mil air gap to the edge of the reference plane when the traces run parallel to the plane edge, excluding minor voids around pins or vias.
- Note that it is *required* for PCI Express traces on the add-in card to reference the ground (GND) plane per the PCI Express Card Electromechanical Specification.
- Note that it is *preferred* that PCI Express traces on the system board reference the main ground (GND) plane or be otherwise GND referenced wherever possible.
 - GND islands implemented on the power (PWR) reference plane should be tied to the main GND plane within 100 mils of the point of signal interface/transition and at varying intervals of no more than 0.5 inches apart along the edge GND island/plane. Points of signal launch may include the signal pins of the chipset, connector pins, AC cap pads, and edge fingers or via transitions. Please refer to Figure 1-17.
 - GND pins on a package can be used for this purpose as well.

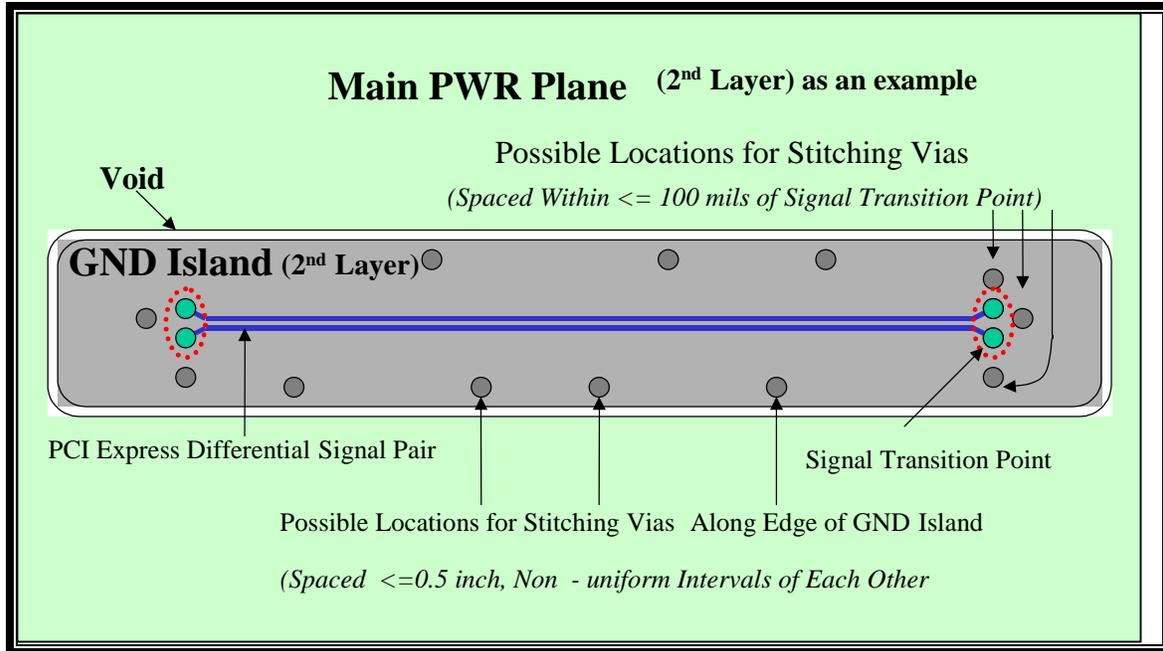


Figure 1-17. Example of a PCB Showing the Differential Pair on the Top Layer and a GND Island on the Second Layer (PWR Plane) with its Potential Stitching Vias

- If the differential pairs on the system board cannot be GND referenced, they are allowed to reference the PWR plane provided stitching caps are used.
 - Stitching caps should be placed within 100 mils (TBD) from the point where the differential pair transitions reference planes (e.g. switches from GND to PWR plane referencing or vice versa).
 - A minimum of one (1) stitching cap can be shared by up to four differential pairs as long as each pair is within the 100-mil (TBD) distance requirement stated above.
 - Stitching cap values of 0.1uF are sufficient.
 - Stitching caps are recommended at each instance of a reference plane change. This requires that a stitching cap also be placed at the connector on the system board, per the guidelines immediately above, for any non-GND referenced traces. This is due to the fact that the connector and add-in card are essentially GND referenced.
- When traces change layers and keep the same reference plane, stitching vias must be used.
 - Stitching vias should be placed as close as possible to the differential pair signal vias (≤ 100 mils TBD). Up to four (4) differential pairs may share 1 stitching via if they are within the 100-mil (TBD) distance requirement. Placement of additional stitching vias, where possible, is recommended.
 - Note that stitching vias serve only to connect the upper and lower reference planes together and thus help provide adequate signal return paths.
- A differential pair routed on a single layer should maintain a single reference plane along the entire routing path.

- Additionally, both traces of a differential pair should remain over the same reference plane (e.g. one net should not reference GND while the other references PWR, etc.).
- **Never route a trace straddling a plane split.** When routing parallel to a plane split, route over the plane and use the 4h rule (trace edge to plane split edge is at least 4 times the trace height above the plane or 20 mils, whichever is greater).

As mentioned above, PCI Express traces should not cross any splits or voids in the reference plane. However, although it should be avoided whenever possible, under certain circumstances it may be necessary for a trace to be partially routed over a via anti-pad in the chipset or connector escape area. This, along with other breakout area considerations is discussed in the following section.

1.3.2.5 Breakout Area Specific Routing Guidelines

The following guidelines should be used for breakout areas of a PCI Express interconnect, especially with regards to a chipset pin or ball field. Note that specific pinout recommendations for various interfaces (x1, x2, x16, etc.) are not included here.

- PCI Express signal pinouts are to be broken out, or grouped, as differential pairs.
 - Side-by-side breakout routing is preferred and recommended wherever feasible. Please see examples in Figure 1-18.
- As required, **length matching** in a chipset breakout area or connector pin field should occur as close as possible to the signal pins **without introducing any “tight bends”** as defined in Section 1.3.2.8. Otherwise length matching should occur within the **immediate segment** (See Section 1.3.2.3).

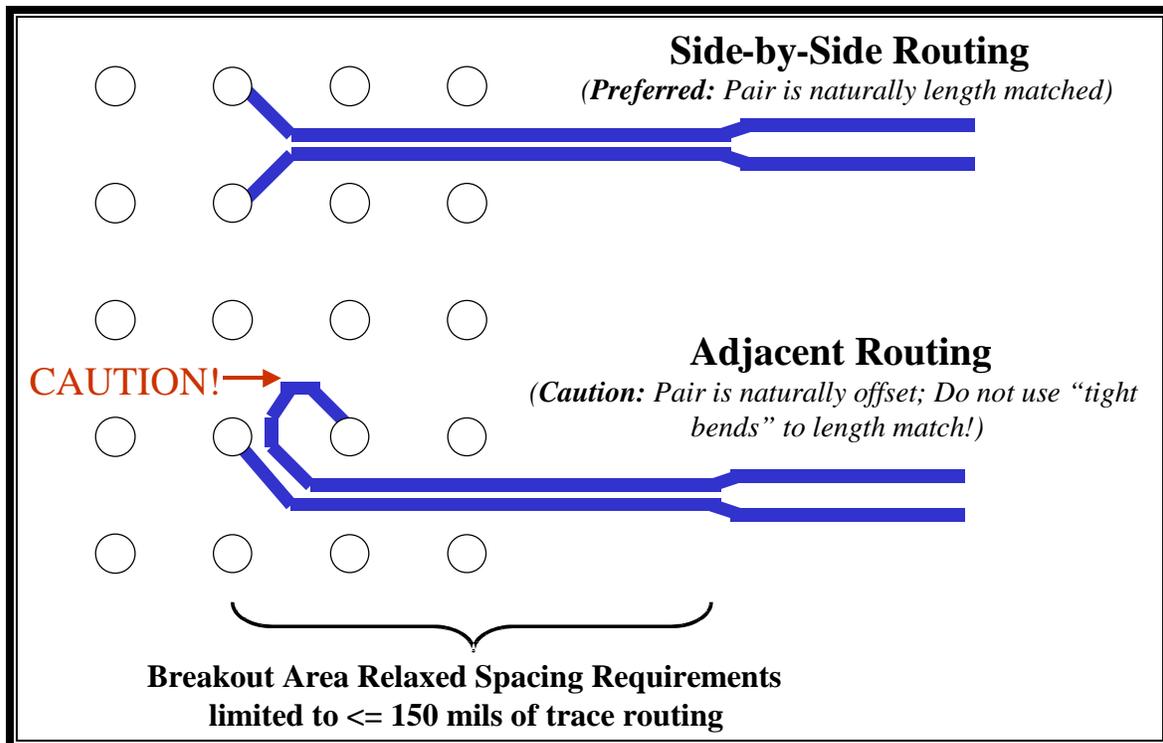


Figure 1-18. Package Pinout/Breakout Possibilities

- The specific breakout area, where exceptions to the general trace routing guidelines may occur, should be limited to no more than 150 mils.
 - Breakout areas are allowed at the chipset pins, the connector pins and/or at the add-in card edge fingers.
 - Tighter spacing is allowed at breakout areas (5 on 7 routing may become 5 on 5 routing) for the coupled differential pairs.
 - Closer pair to pair spacing of 10 mils is allowed in the breakout region.
 - While any deviations longer than the 150-mil breakout area guideline aren't specifically precluded, they should be simulated to ensure that the concessions allowed for the increases in length do not negatively impact the crosstalk margin or the jitter and loss budgets.
- Note that small sections of trace (in breakout of pins to a via, up to 50 mils is OK w/o reference) used to attach a chipset pin to a via do not necessarily require any special reference plane considerations. See Figure 1-19.

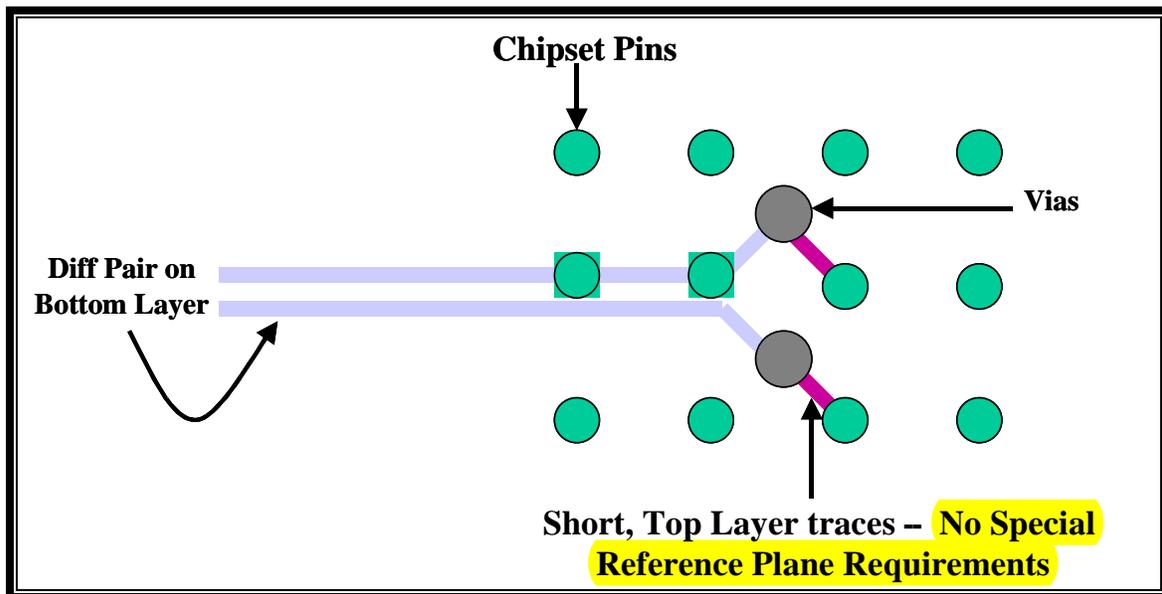


Figure 1-19. Breakout Area Traces

- When routing in a breakout area, via anti-pads may present discontinuities in the reference planes. The differential pair should avoid routing over these areas wherever possible. However, if they can't be avoided the following guideline should be used:
 - The amount of the trace that is over a void (e.g. via anti-pad) should be minimized in terms of the length and the percentage of the width of the trace; at worst, no more than half of the trace width should be over the via anti-pad at any given time. Please see Figure 1-20 for an example scenario. Note that this is only applicable to via antipads, not other voids like mounting holes.
- For other voids than antipads (mounting holes, for example) avoid routing over the void (specifically to avoid metal washers, etc.).

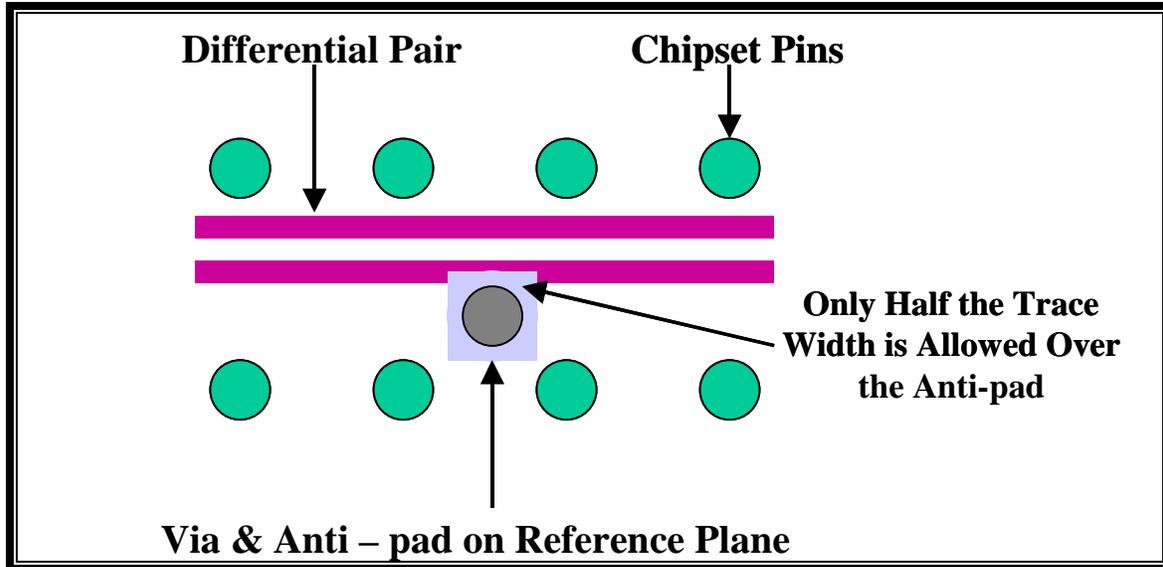


Figure 1-20. Example Scenario of a Trace Routing Over an Anti-pad Void

1.3.2.6 Edge Finger Design: Add-in Card

The edge finger of an add-in card is an important element of the interconnect path. The edge fingers themselves are designed to interact with the properties of the connector pins in order to produce an acceptable impedance target and ensure a robust signaling path. The following guidelines and considerations apply to the edge finger design:

- The reference planes on the inner layers should be removed immediately under the gold finger areas that exist on the outer layers of the PCB. The planes should be removed along the entire length of the edge finger component that contains PCI Express differential pair signal traces. See Figure 1-21 for a conceptual illustration of how this might be done.
 - This removal of the reference planes should be restricted to the actual area of the edge fingers only and not extend out to the trace routing area.
 - Note that the entire ground and power planes should be removed under the entire edge finger for PCI Express signals. This applies to both top and bottom layer edge fingers.
 - It is also important that any plating bars for the gold fingers be removed during the PCB manufacturing process.
- Actual mechanical dimension requirements for the edge fingers and add-in card are specified in the PCI Express Card Electromechanical Specification.

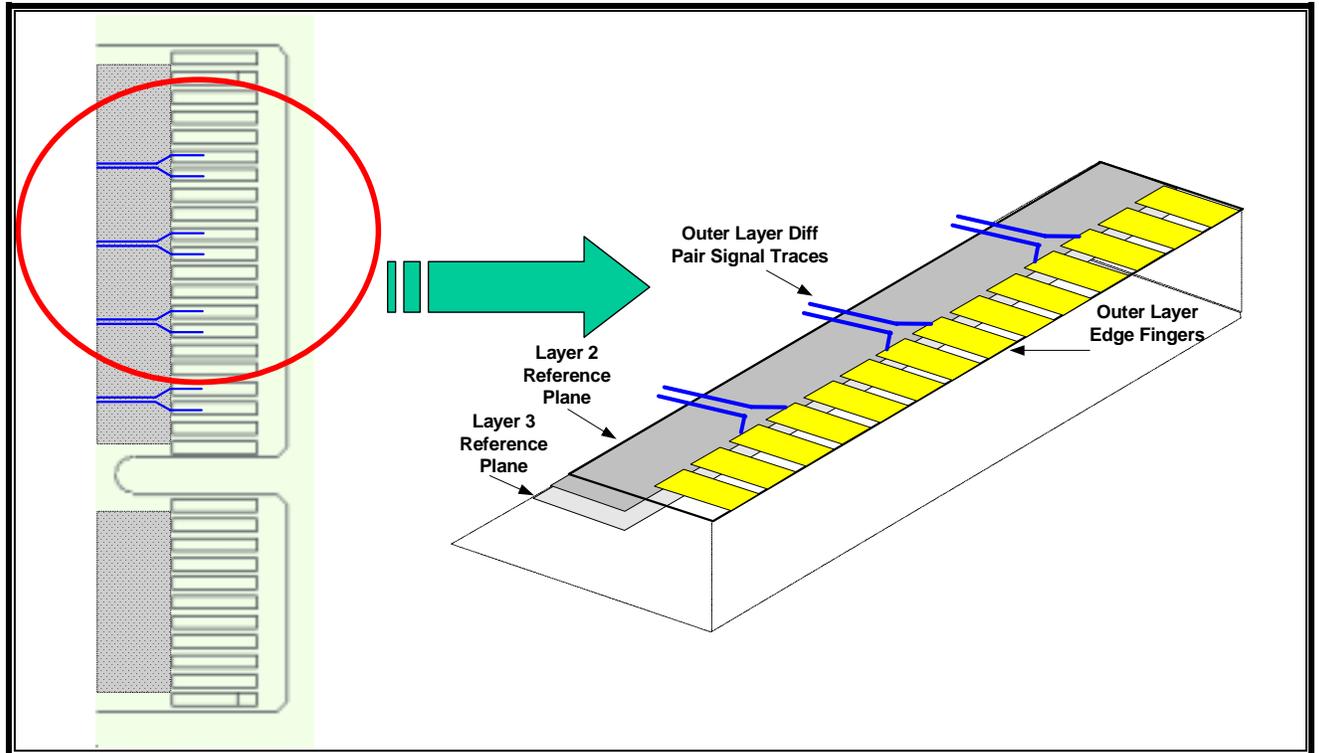


Figure 1-21. Example of PCI Express Reference Plane Voids for a x4 PCI Express Edge Finger Interface

1.3.2.7 Via Usage and Placement

Vias are allowed on a PCI Express differential interconnect; however, their use should be somewhat limited. Guidelines concerning via usage are as follows:

- Signal vias have the potential to adversely affect the overall loss and jitter budgets and may limit the maximum routing length achievable; therefore, they should be carefully used.
- It is recommended that vias have a pad size of 25 mils or less and a finished hole size of 14 mils or less.
- Signal vias should have pads on unused internal layers removed.
- Vias should always be placed as a pair - they should be at the same relative location and placed in a symmetrical fashion along the differential pair (See Figure 1-22).
- When a differential pair transitions reference planes through the use of vias, the guidelines in Section 1.3.2.4 should be followed.

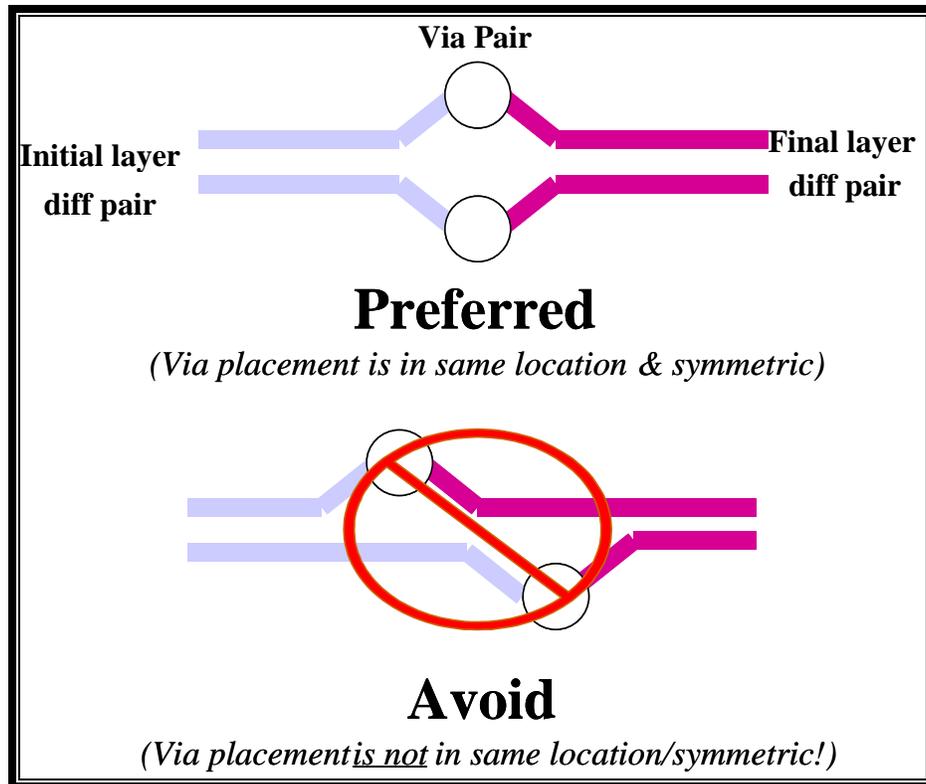


Figure 1-22. Placing Vias as a Pair

- A maximum of four (4) via pairs may be used on each TX differential pair of a system board/add-in card topology:
 - One via pair may be located at the breakout section of a chipset to transition layers and/or adjacent to the edge finger of the add-in card PCB. The plated through-holes (PTHs) for the connector pins should also be considered as a via pair.
 - Two pairs of vias are allowed to transition layers and may be used in order to accommodate placement of the AC coupling capacitors.
- A maximum of two (2) via pairs may be used on each RX differential pair of a system board/add-in card topology.
 - One via pair may be located at the breakout section of a chipset to transition layers and/or adjacent to the edge finger of the add-in card PCB. The plated through-holes (PTHs) for the connector pins should also be considered as a via pair
- Note that the above guidelines provide for a maximum of six (6) via pairs for the entire interconnect path. A maximum of six (6) via pairs may also be used for chip-to-chip topologies.

Though a restriction on the number of via pairs has been recommended, tradeoffs in the design may be allowed as long as the loss and jitter budgets are still met for the interconnect. Such tradeoffs will likely require further simulations of the interconnect. For example, a designer may choose to sacrifice overall trace length in order to allow an extra via pair to be present in the interconnect path. It has been found that, in some corner cases, each via pair could potentially contribute upwards of 0.25 dB of loss. Also recall that each inch of trace is anticipated to contribute approximately 0.25-0.35 dB to the overall loss budget. Therefore, by subtracting one

or two inches of trace from the previously recommended maximum trace lengths (See Section 1.3.2.2) enough margin may exist for an additional via pair to be present in the system. See Table 1-4 for an example of this scenario. Keep in mind, however, that this is a simplistic approach and that the location of the via pair among other things will play a role in exactly how much of an impact the vias will have on the loss and jitter budgets. For this reason, even if tradeoffs are made with trace the trace lengths that are routed, topologies that utilize vias beyond the boundaries of the guidelines described in the bulleted items above should be simulated to determine their true impacts on the system.

Table 1-4. Example Via Usage and Trace Length Tradeoffs

Via Tradeoffs in Terms of Loss/Attenuation	
<i>(EXAMPLE EQUIVALENTS ONLY!)</i>	
One or Two Via Pairs	= 1 or 2 inches of PCB Diff Pair Routing
<i>Hypothetical Example: If a given interconnect with 10" of trace was known to consume the entire interconnect budget, reducing the interconnect length by 4" to a total length of 6" may allow 2 to 4 more via pairs to be added to the interconnect path.</i>	

1.3.2.8 Bends

While serpentine routing (adding additional bends to arbitrarily increase trace length) isn't required to match lengths from one differential pair to another, bends (or turns) are still likely to be encountered while routing a PCI Express interconnect. For example, a coupled differential pair cannot always route in a straight line and often times must change direction, thus requiring some type of bend or turn. Also, a single trace of the differential pair may require some type of bend or turn as it routes into a component's pin or pad. However, "tight bends" as defined below, should be avoided as they can substantially impact the loss and jitter budgets. Please refer to Figure 1-23 for example scenarios. This is applicable to both coupled differential pairs and uncoupled sections of trace.

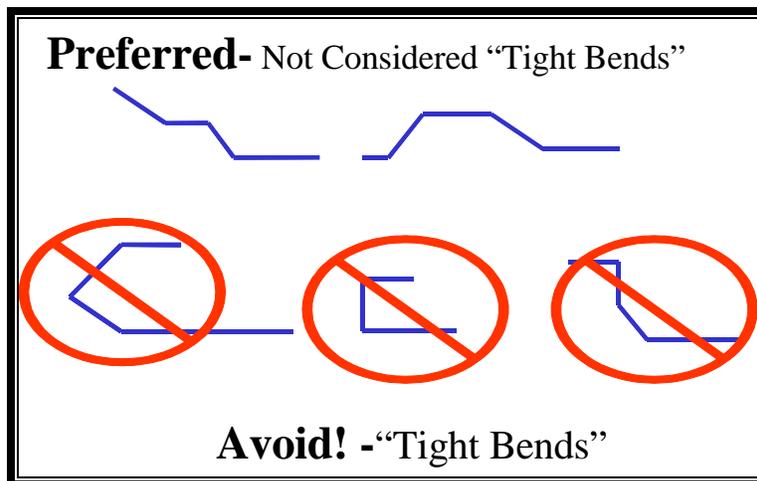


Figure 1-23. Example Illustrations of Acceptable Bends vs. "Tight Bends"

Bend guidelines for coupled sections of a differential pair (Please refer to Figure 1-24):

- Keep bends to a minimum for coupled differential pairs. The bends can introduce common mode noise into the system, which can affect the signal integrity/EMI of the differential pair.
- If bends are used, the following prioritized guidelines should be followed to avoid “tight bends”(see Figure 1-24):
 1. All angles (α) between traces should be ≥ 135 degrees (i.e. traces at 45-degrees); there should be no 90-degree bends or turns.
 2. The “inner air gap” (A) of a bend on a coupled differential pair should be \geq the min pair to pair spacing of 20 mils (see section 1.3.2.1).
 3. Segments (B) and (C) should have a length of $\geq 1.5x$ the width of the trace (the length being measured with respect to the middle of the trace/bend vertex).

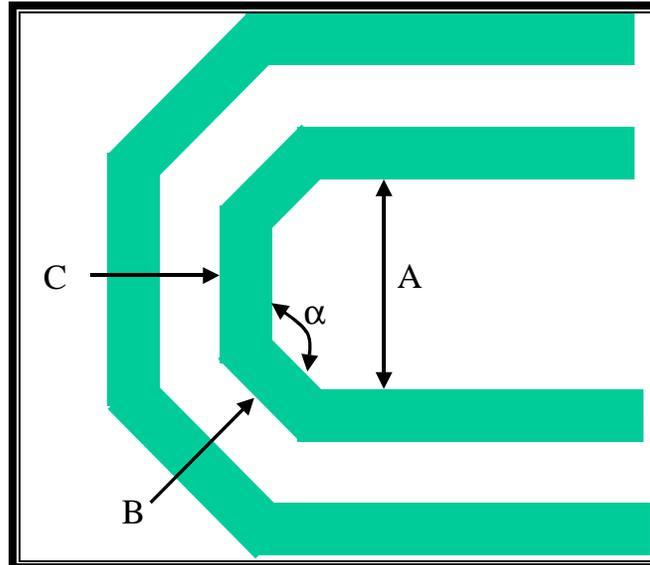


Figure 1-24. Bend Illustration Diagram for Coupled Differential Pairs

- Where possible, the number of left and right bends should be matched as closely as possible.
 - Alternating left and right turns will help to minimize skew due to length differences between each signal of the differential pair.
- Whenever it is not possible to match the number of left and right bends, the length required to match the two traces as a result of the difference in the number of bends should be determined. Half of this length should be added to each end of the shorter trace of the pair.
 - For example, if the D+ of a tightly coupled differential pair signal was 10 mils shorter than the D- signal due to bends, add 5 mils to the D+ signal line before the bends occur and add 5 mils after the bends occur to achieve an overall length match between the D+ and D- signals within the pair. This solves skew mismatch, but does not guarantee against the possibility of introducing any common mode noise into the system.

Bend guidelines for uncoupled sections of a diff. pair which route into a component's pin or pad:

- The uncoupled section of trace routing into a pin or pad should be ≤ 45 mils when using multiple bends (e.g. when length matching the differential pair). Refer to Figure 1-25.
 - The 45-mil length measurement is made with respect to the middle of the trace/bend vertex from the start of the uncoupled section to the edge of the pin or pad.
 - When trace lengths are closely matched, the preferred method of routing is a symmetrical route without multiple bends. The uncoupled length in this case should be minimized as best as possible.

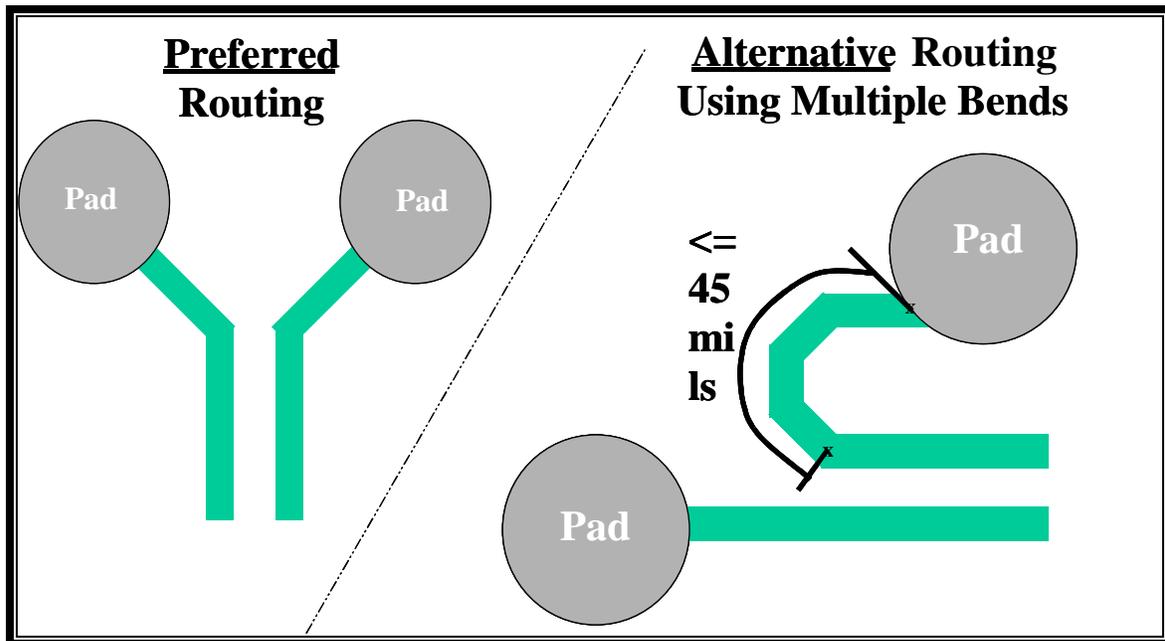


Figure 1-25. Routing/Bend Diagram for Uncoupled Traces into Component Pins/Pads

- When using multiple bends, follow these prioritized guidelines should be followed to avoid “tight bends”(see Figure 1-26):
 1. All angles (α) between traces should be ≥ 135 degrees (i.e. traces at 45-degrees); there should be no 90-degree bends or turns.
 2. The minimum “air gap” (A) of a bend should be $\geq 3x$ the width of the trace.

Note that the direct concerns addressed by rules 3 and 4 are generally alleviated when following rules 1 and 2 and the 45 mil max length guideline. Therefore, rules 2 and 3 don't necessarily need to be verified during CAD for each and every net, but are instead provided for completeness.

3. Segments (B) and (C) should have a length of $\geq 1.5x$ the width of the trace (the length being measured with respect to the middle of the trace/bend vertex).
4. Segment (D) should be minimized as much as possible.

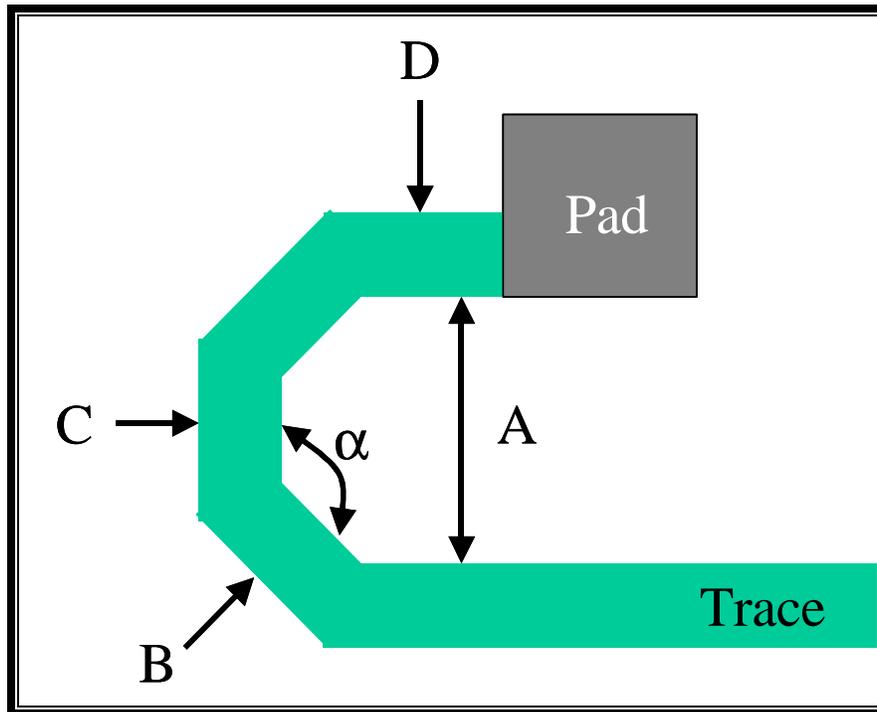


Figure 1-26. Bend Illustration Diagram for Uncoupled Section of a Differential Pair

1.3.2.9 Test Points and Probing

The inclusion of test points and probing structures has the ability to impact the loss and jitter budgets of a PCI Express interconnect. This isn't to say that they can't be tolerated, however. In general, test points and probe structures should not introduce stubs on the differential pairs or cause them to wildly deviate from the recommendations given throughout this document. Existing vias, pads or pins should be used wherever possible to accommodate such structures. Careful consideration must be taken whenever additional probing structures are used.

1.3.3 PCI Express Topologies

The PCI Express link is a point-to-point topology for each signal of the transmit and receive differential pairs. However, the interconnect path needs to include DC blocking capacitors, and may need LAI probe pads and allowances for package breakouts, vias and connectors. On the following pages are some topology examples that illustrate some specific cases to be used for design reference.

The topologies presented are generic in the sense that any of the stackups may use these topologies (with the exception of the 4-layer topology which is restricted to microstrip topologies). The segment lengths for these topologies are dependent upon the stackup layer to trace spacing, trace width, thickness and trace-trace spacings. For example a 4-mil wide trace stripline topology with the same trace to plane spacing, trace to trace spacing can not be routed as long as a 5-mil wide trace that otherwise has the same stackup and spacing. These topologies are provided as a reference to be used to derive specific platform or add-in card topology rules.

1.3.3.1 Interconnect Topologies for Two Components on the Baseboard

An example topology that has both transmitting and receiving components located on the baseboard is shown in Figure 1-27. The example shown is a 10-layer stackup, but this topology illustration also can apply to a 4-layer, 6-layer or 8-layer stackup. In this example, the routing segments illustrated are defined in Table 1-5. Both microstrip and stripline layers are shown in this topology. Where the figure shows more than one layer trace illustrated as connected between adjacent vias, only one of the layers is actually selected for the specific differential pair routing. For example, for the L3 segment layer 3 (stripline), layer 8 (stripline) and layer 10 (microstrip) are shown, but only one of these layers is actually to be used for the two traces of a differential pair. Four vias max are recommended for the baseboard topology, so it is expected that not all of the vias shown will actually be used. This example shows the AC capacitor located near the transmitter, but this is not required. Also illustrated is a probe pad (connected without a stub in the topology by segments L4 and L5).

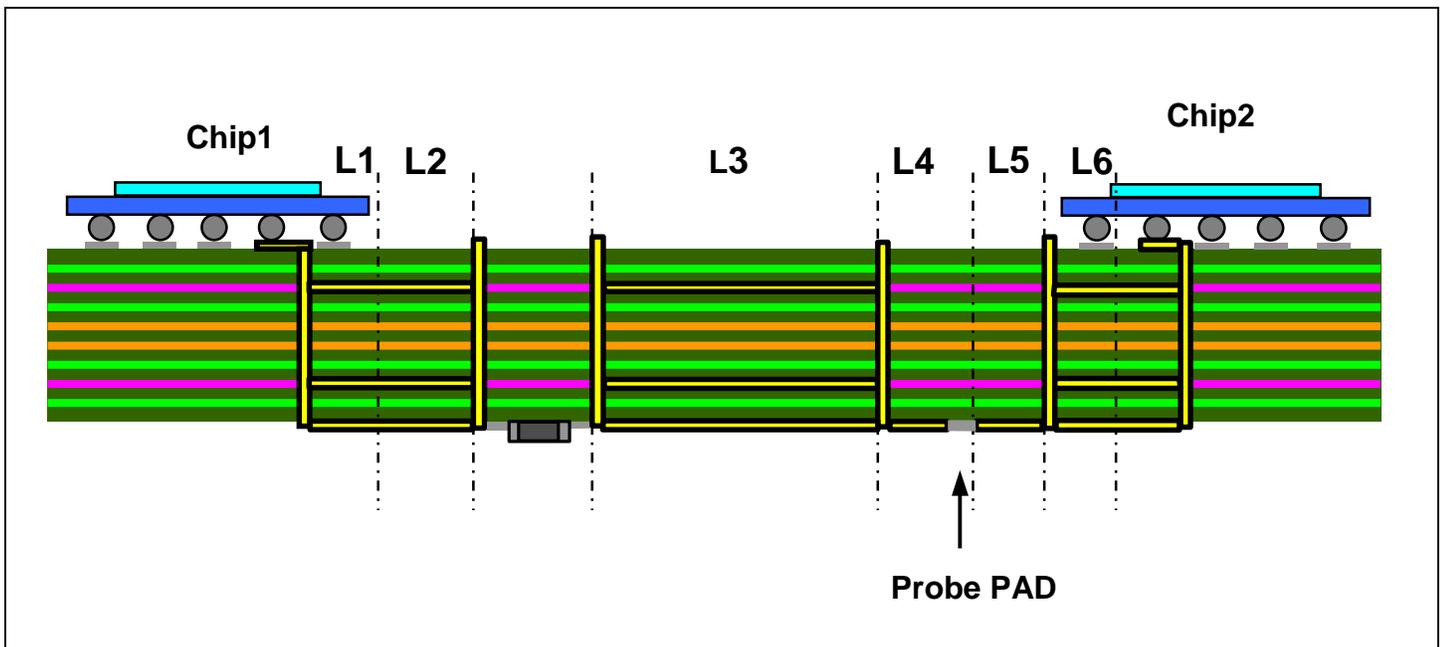


Figure 1-27. Baseboard Chip-to-chip Topology Example

Table 1-5. Chip-to-chip Topology Segment Definitions

Segment	Definition	Notes
L1	Breakout region	Less than 0.150"
L2	Breakout region to AC coupling cap	No stubs allowed.
L3	Main trunk	
L4, L5	Probe pad breakout traces	No stubs allowed.
L6	Receiver routing region	

1.3.3.2 Interconnect Topologies for Baseboard with Add-in Card

An example topology that illustrates a baseboard or add-in card component to connector case is shown in Figure 1-28. The example shown is a 10-layer stackup, but this topology illustration also can apply to a 4-layer, 6-layer or 8-layer stackup. In this example, the routing segments illustrated are defined in Table 1-6. Both microstrip and stripline layers are shown in this topology. Where there is more than one layer trace illustrated as connected between adjacent vias, only one of the layers is actually used for the specific differential pair routing. Four vias max are recommended for the baseboard topology (if TX trace) and only two for an add-in card (if RX trace), so it is expected that not all of the vias shown will be used. This example shows the AC capacitor, which is present for the transmitter on this baseboard or add-in card. If the component port is a receiver, there will be no AC capacitor in this topology. It will be on the substrate on the other side of the connector. Also illustrated is a probe pad (connected without a stub in the topology by segments L4 and L5).

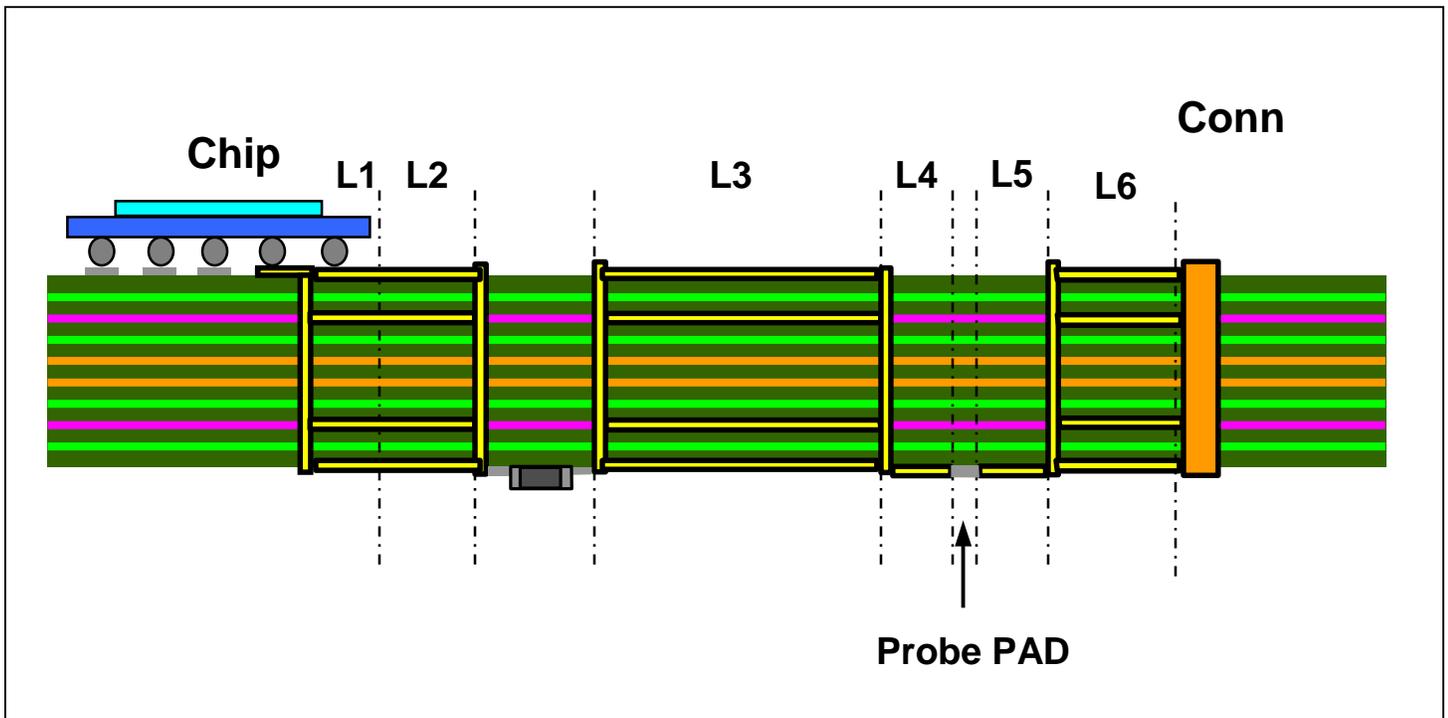


Figure 1-28. Baseboard Chip-to-connector and Add-in Card Example

Table 1-6. Baseboard Chip-to-Connector and Add-in Card Segment Definitions

Segment	Definition	Notes
L1	Breakout region	Less than 0.150"
L2	Breakout region to AC coupling cap	No stubs allowed.
L3	Main trunk	
L4, L5	Probe pad breakout traces	No stubs allowed.
L6	Connector segment routing region	

1.3.4 Passive Components and Connectors

1.3.4.1 AC Coupling Capacitors

The PCI Express specification requires that each lane of a PCI Express link be AC coupled between the transmitter and receiver. The PCI Express Base Specification allows for the AC coupling components to be located either on or off the die. It is expected that in most cases, the AC coupling will be separate from the die and in the form of discrete capacitors on the printed circuit board itself. AC coupling capacitors are associated with the transmitter for each signal so capacitor placement depends upon transmitter port location. Use the following guidelines for AC coupling capacitors:

- The spec requires that a value between 75nF and 200nF be used for each of the capacitors.
- Specific tolerance values of the capacitors are not required as long as the overall range of capacitance falls within the spec minimum/maximum values.
 - Note that temperature or voltage effects on capacitor tolerance values should also be considered (Y5V, for example, may be rated at +/- 20%, but may exhibit closer to +/- 80% with temperature changes).
- Dielectric properties are not a major consideration for PCI Express AC coupling capacitors.
 - It is anticipated that any type of capacitors including X7R, NPO, COG, Y5V, etc. are acceptable as long as they meet all other requirements.
- **C-packs are NOT allowed for PCI Express AC coupling capacitor purposes.**
- The exact same package size and value of capacitor should be used for each signal in a differential pair.
 - While it is anticipated that size 603 capacitors will be acceptable, size 402 capacitors are strongly encouraged.
 - The smaller the package size, the less ESL (Equivalent Series Inductance) introduced.
- Routing the signals on one layer (bottom layer) and then transitioning through vias to place the capacitor is allowed, but affects the loss budget and is not preferred.
- Locate capacitors for coupled traces at the same location along the differential traces (they should not be staggered from one trace to the other).
 - **Place them as close to each other as possible as allowed by DFM rules to avoid creating large uncoupled sections within the differential pair traces.**
- The relative location of AC capacitors from one differential pair to another is not important.
- Pad sizes for each of the capacitors are to be the minimum allowed for DFM to minimize parasitic impacts.
- For baseboard chip-to-chip interfaces, the AC coupling capacitor can be placed flexibly between the transmitter and receiver, but the location is recommended to be in the first or last 1/3 of the trace.
- For add-in cards, the AC coupling capacitors are located on the card itself for each of the TX pairs originating from the add-in card PCI Express device.
- For baseboard add-in card topologies, the AC coupling capacitors must be located on the baseboard for the TX pairs originating from the baseboard PCI Express device.
- For 4-layer baseboards that only allow single sided (top) placement, the TX to RX pairs that come from the baseboard and go to the add-in card, respectively, and are likely to all be routed on the top layer.

The “breakout” into and out of the capacitors should be symmetrical for both signal lines in a differential pair. Minimize trace separation for routing to pads in order to maximize the amount of tight coupling between the signal pairs. See Figure 1-29 for a conceptual example of this.

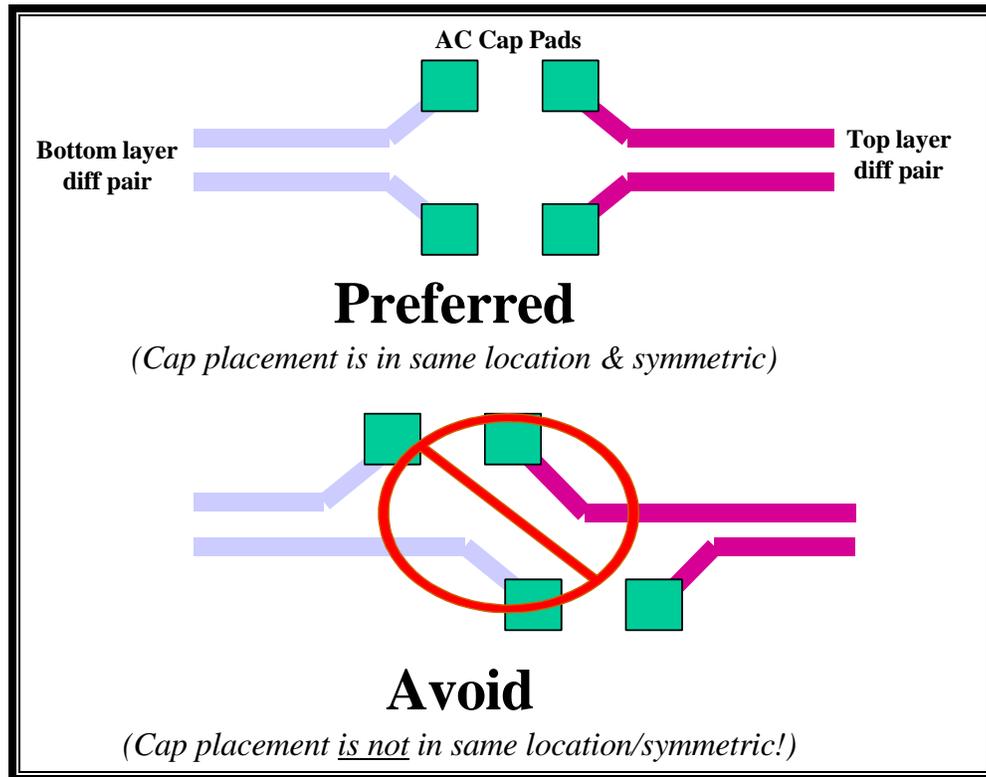


Figure 1-29. Symmetrical Routing into AC Caps

Though the capacitors may theoretically be placed anywhere along the interconnect, and tend to reside near the connector, it is recommended that they are no more than 400 mils (TBD) away from the connector pins when routing TX pairs on the system board to the connector. Likewise they should be placed no more than 250 mils (TBD) from the edge finger for add-in card TX pairs. The pin closest to the edge of the connector should be used as the reference point if the D+ and D- pin locations are staggered, as is the case with a through-hole style of connector. When routing chip-to-chip, the caps should be off-centered. It is actually desired that they should be placed closer to the RX inputs on the receiving chip vs. the TX outputs on the transmitting chip. It is recommended, however, that the considered location of the capacitors actually be simulated as their exact placement may introduce potential resonance issues within the system.

Table 1-7. AC Coupling Capacitor Guidelines

PARAMETER <i>(FOR AC COUPLING CAPACITORS)</i>	REQUIREMENT
Cap size	<ul style="list-style-type: none"> • 402 preferred • 603 acceptable, impacts budget more than 402 • Must use same size on both pair signals • C-Packs NOT allowed.
Cap Value	<ul style="list-style-type: none"> • 75 nF min and 200 nF max • Must use same value on both pair signals
Cap Value Tolerance	Not applicable as long as specified min/max range is met when the tolerance is considered.
Cap Placement — Within a Differential Pair	<ul style="list-style-type: none"> • Must be placed at same exact location within the two differential pair signal traces • Symmetric routing into the caps strongly recommended. Matched line lengths on either side of the caps for each line of the diff pair.
Cap Location —Chip to Connector Routing	<ul style="list-style-type: none"> • Within of 400 mils (TBD) of connector pins • Within of 250 mils (TBD) of add-in card edge finger
Cap Location —Chip to Chip Routing	<ul style="list-style-type: none"> • Minimum of 100 mils from component pins • Should be closer to Tx or Rx component, not in center of interconnect (not middle 1/3 region)

1.3.4.2 Connectors

The PCI Express Electromechanical Specification defines the connector to use in conjunction with edge finger add-in cards. However, several items worth mentioning with respect to the connector are:

- For through-hole connectors, the pins of a differential pair are offset from each other. This delta of mismatch between the pins should be directly accounted for by the PCB trace on the baseboard. Refer to the Length Matching section of this document for more details on length matching requirements.
- The two traces of a differential pair should both route into a connector pin field from the same layer.
- Note that the pinout names of the connector are defined with respect to the baseboard:
 - PETp(x) and PETn(x) pins (the transmitter differential pair of the connector) shall be connected to the PCI Express transmitter differential pair on the baseboard, and to the PCI Express receiver differential pair on the add-in card.
 - PERp(x) and PERn(x) pins (the receiver differential pair of the connector) shall be connected to the PCI Express receiver differential pair on the baseboard, and to the PCI Express transmitter differential pair on the add-in card.
- The connectors and the add-in cards are keyed such that smaller cards can be put in larger connectors. For example, an x1 card can be inserted into the x4, x8, and x16 connectors. This is referred to as up-plugging. Up-plugging layout adjustments need to be made to accommodate up-plugging. However, down-plugging is not allowed and is physically prevented.

- Adjacent differential pairs on the connector are separated by two ground pins to manage the connector crosstalk. Ground pins should directly tie to the ground plane on the baseboard.

It is presently assumed that either one connector (add-in card attaches directly to baseboard) or two connector (baseboard with riser card attached between add-in card and baseboard) will be present in any given PCI Express interconnect path. Special considerations must be taken to include the riser card in addition to the add-in card. The riser card loss and jitter budget is considered part of the baseboard budget, not the add-in card budget. Simulations should be run to validate that all appropriate budgets are still met when more than one connector is used in the interconnect path.

1.4 Summary

This document provides practical layout guidelines to design a PCI Express interconnect on a PCB with typical FR4 materials. It is intended to supplement the PCI Express Base Specification and PCI Express Card Electromechanical Specification to help designers implement a PCI Express interconnect. The PCI Express interconnect is a point-to-point layout of the serial differential signal trace pairs. It is generally easier to implement compared to the multidrop, parallel layout topology of the conventional PCI interconnect. For PCI Express, board designers are not constrained by tight trace length matching requirements between a wide data bus and its associated clock to meet timing skews. Instead, careful PCB layout techniques are used to minimize loss and to maintain symmetry for each differential trace pair while a generous allowance for pair-to-pair timing skew remove the length matching requirements across pairs (between lanes and also between transmit and receive pairs of each lane). These layout guidelines serve as a starting point for board designers to implement a PCI Express interconnect. Further design optimization can be achieved by performing interconnect simulations and signaling validations.