AN 672: Transceiver Link Design Guidelines for High-Gbps Data Rate Transmission

2020.01.29

AN-672 Subscribe Send Feedback

As transceiver data rates increase and the unit interval time (UI) decrease, the end-to-end link design of a transceiver channel becomes increasingly critical to the overall performance of the link. Consider a Stratix® V FPGA with transceiver channels operating up to 28 Gbps. At this data rate, the UI is less than 36 ps. Any signal degradation of the channel can impact jitter margin and increase eye closure, resulting in increased bit error rates (BER). Two of the proposed standards for 100-Gpbs data transmission using fewer lanes are CEI-25G-LR and CEI-28G-VSR.

CEI-25G-LR is intended for 25 Gbps data transmission across long-reach backplane architectures. Likewise, CEI-28G-VSR specifies a 28-Gbps data rate for very short-reach chip-to-module and chip-to-chip applications. For these standards, the total insertion loss budget for the link at the Nyquist rate is approximately -25 dB for CEI-25G-LR and -10 dB for CEI-28G-VSR. Successful data transmission across these types of links requires the designer to minimize signal degradation caused by the channel to meet stringent loss requirements. Understanding the various factors that contribute to channel loss allows the designer to make appropriate design choices to mitigate adverse effects and achieve optimal link performance. This application note describes how to optimize a complete transceiver link for these very high-Gbps data transmission designs.

Figure 1: Typical System End-to-End Transceiver Links

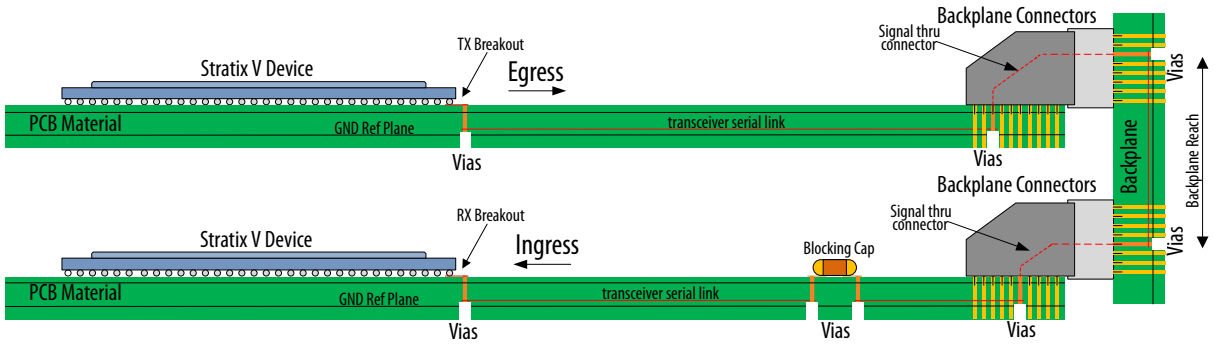

(a) Line Cards Over Backplane Link

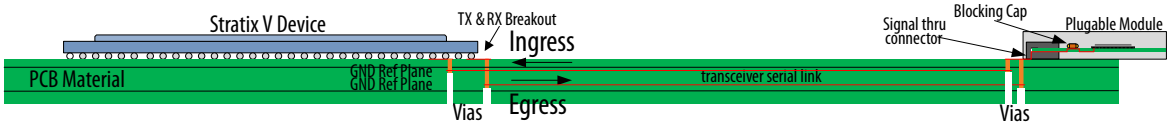

(b) Chip-to-module Link

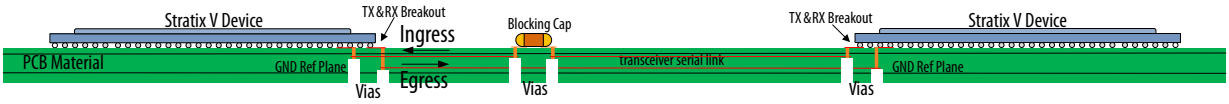

(c) Chip-to-chip Link

A typical end-to-end transceiver link can be separated into three main components:

- PCB material
- stack-up design
- channel design

Designers must give careful consideration to each of these components to avoid degraded link performance. Link optimization involves understanding and managing the first-order factors that impact signal performance for each of the three components. For high-speed signal transmission, these first-order factors are:

- Signal attenuation
- Impedance control and discontinuities
- Crosstalk

PCB Material Selection

The printed circuit board (PCB) material is the substrate upon which the channel design is constructed. It consists of the core, prepreg dielectric layers, and copper foils stacked and glued together to from the complete PCB stack-up. The dielectric and copper layers form the reference planes and routing layers for the channel design. The sources of signal attenuation resulting from the PCB material include dielectric loss, conductor loss, reflections due to mismatched impedance, and radiation loss. Losses due to radiation are usually very small and can be ignored. The remaining losses can be attributed to the various properties of the material choice. Material properties that directly affect the link performance include:

- Loss tangent (tan(δ)/Df)
- Dielectric constant (Er/Dk)
- Fiberglass weave composition
- Copper surface roughness

Loss Tangent and Dissipation Factor

Loss tangent (tan(δ)) (also referred to as dissipation factor (Df) by many PCB manufacturers) is a measure of signal attenuation as the signal propagates down the transmission line. This attenuation is the result of electromagnetic wave absorption in the dielectric material and is commonly known as dielectric loss. As frequency increases, the dielectric loss also increases proportionally. Common material choices for high-speed signal layers include Nelco 4000-13EPSI, Rogers 4350B, and Panasonic Megtron 6. For very high-Gbps transceiver designs, Intel has successfully used Panasonic's Megtron 6 material as the laminates for the core and prepreg layers surrounding the high-speed routing layers.

Figure 2: Comparison of Material Loss Tangent

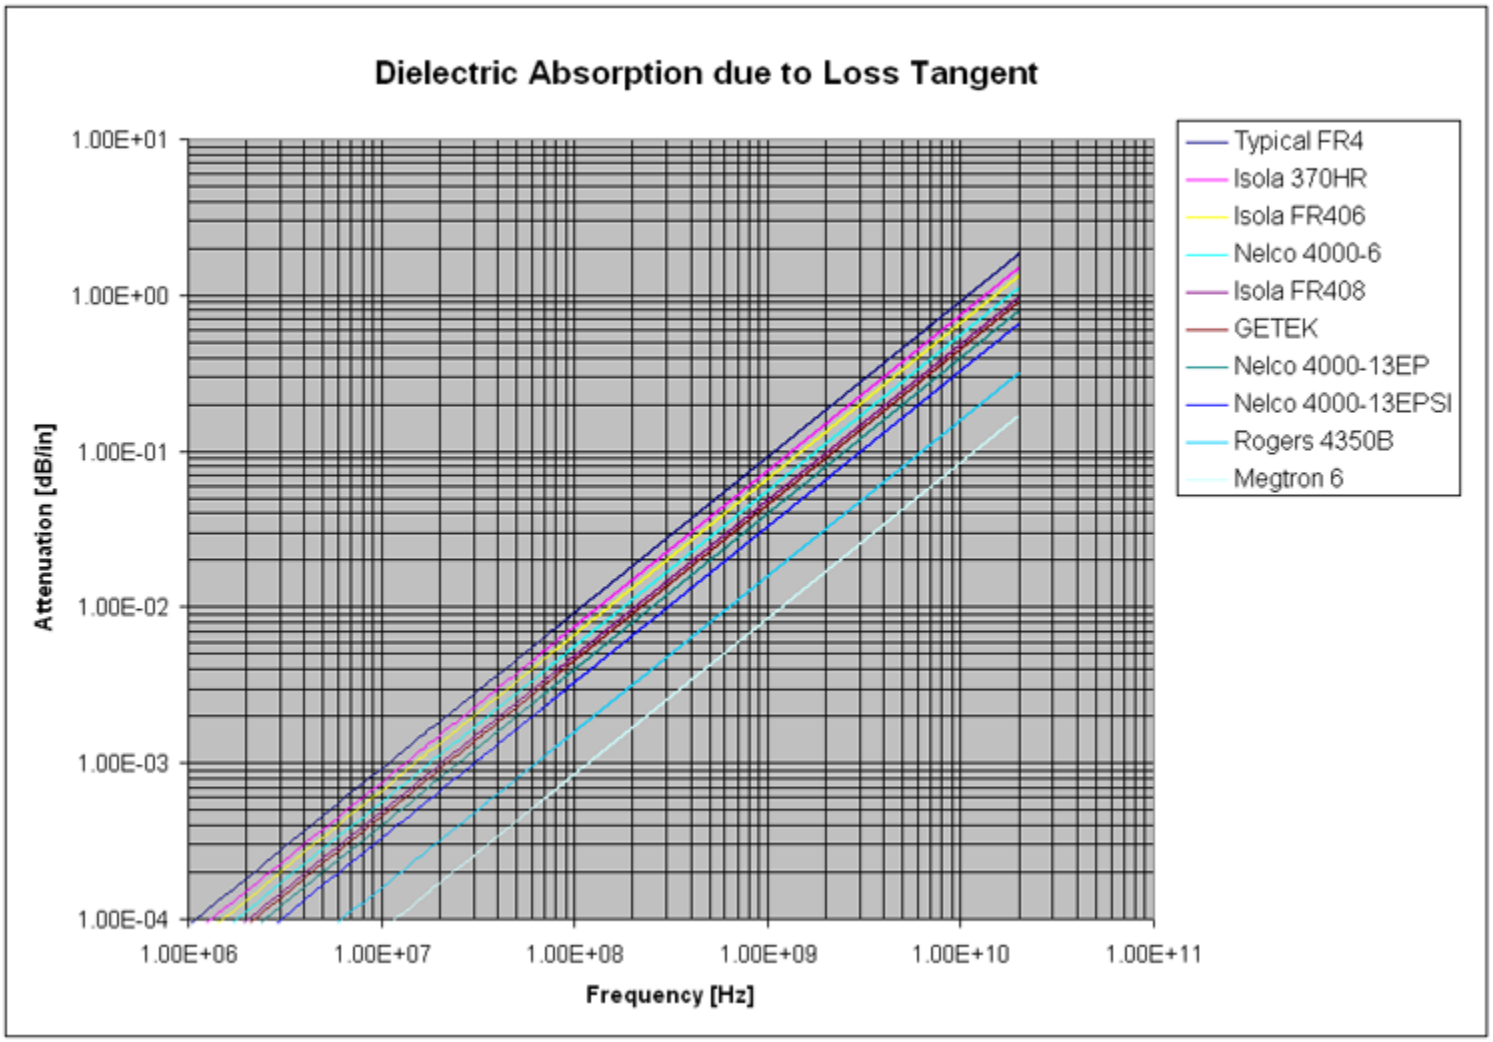

Note: Given the choice, select the material with the lower loss tangent to reduce signal attenuation from dielectric losses.

Dielectric Constant

The dielectric Constant (ε_r or Dk in many material data sheets) is a measure of the insulating properties of the material and affects the capacitance of the conductor embedded within it as well as the speed of signal propagation on the transmission line. Lower dielectric constant provides better insulation, faster signal propagation, higher trace impedance for a given trace geometry and smaller stray capacitance.

Figure 3: Stripline Differential Trace

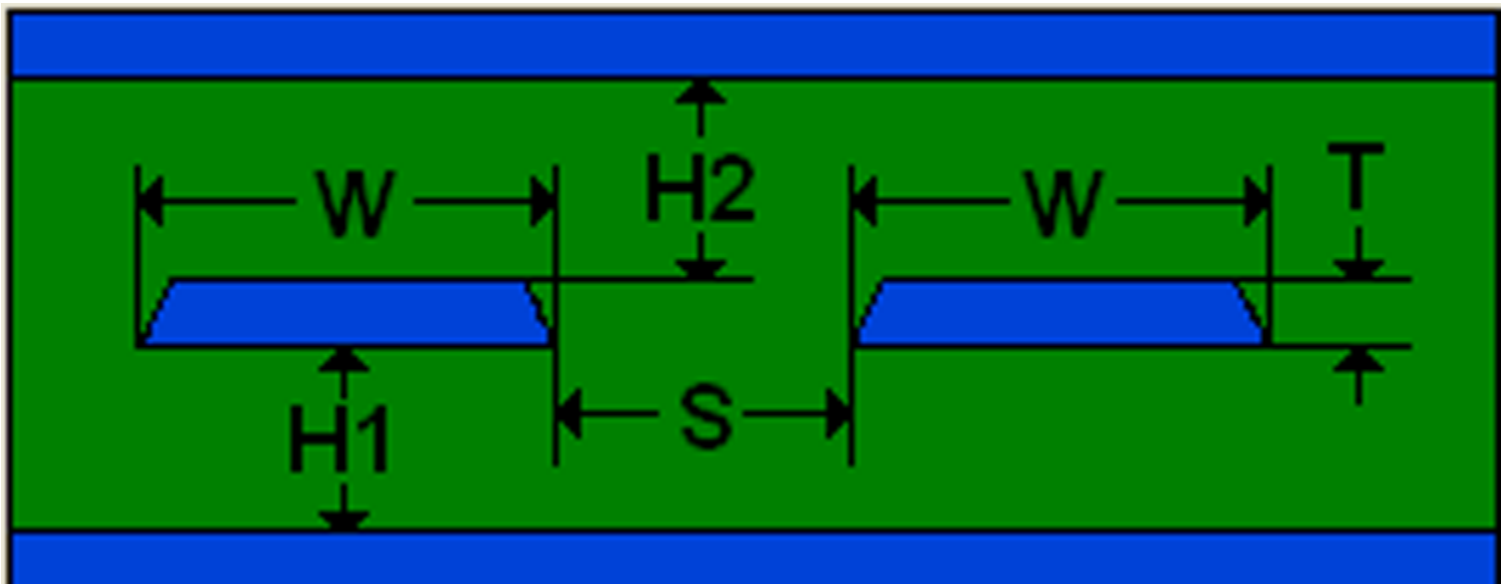

Table 1: 100-Ω Stripline FR4 vs. Megtron-6 Construction

Material	E_r	W (mils)	S (mils)	T (mils)	H1 (mils)	H2 (mils)	Zo (Ω)	Zdiff (Ω)	T_{pd} (ps/in)	C (pF/in)
FR4	4.6	6	15	0.7	7	11	52	100	182	3.5
Meg-6	3.4	6	15	0.7	5	8		100	156	3.1

Consider a 100-Ω differential stripline pair constructed with an optimal trace width (W) of 6 mils and pair separation (S) of 15 mils. For FR4, the total dielectric thickness is 18.7 mils (T+H1+H2). For the same trace construction dimensions using Megtron-6 the total thickness is 13.7 mils. This reduction of total thickness by 5 mils is significant because it quickly adds up for high density boards that require many signal layers. For instance, a typical FPGA board with 6 stripline routing layers can result in 30 mils of PCB thickness savings using the Megtron 6 vs. FR4. This can be the difference between boards that are easily manufactured using standard via drilling processes vs. ones that require more expensive laser drilling techniques because of the smaller board thickness to via aspect ratio. Furthermore, lower ε_r results in faster signal propagation (T_{pd}) and lower trace capacitance (C), which improves signal performance.

An additional consideration for ε_r is that it usually decreases as frequency increases. Decreasing ε_r affects the transmission line in two ways:

- The trace impedance increases with decreasing ε_r causing reflections that further degrade signal quality
- The signal velocity increases with decreasing ε_r causing dispersion of the different harmonics that comprise the digital signal. This results in increased phase jitter at the receiver

Note: Always choose lower ε_r with a flat frequency response for best signal performance and to reduce signal dispersion that adds phase jitter.

Fiberglass Weave

PCB boards are constructed of strands of fiberglass yarn woven into bundles with epoxy resin to form the fiberglass sheets that make up the core and prepreg laminates.

Figure 4: Commonly Available Styles of Fiberglass Weaves

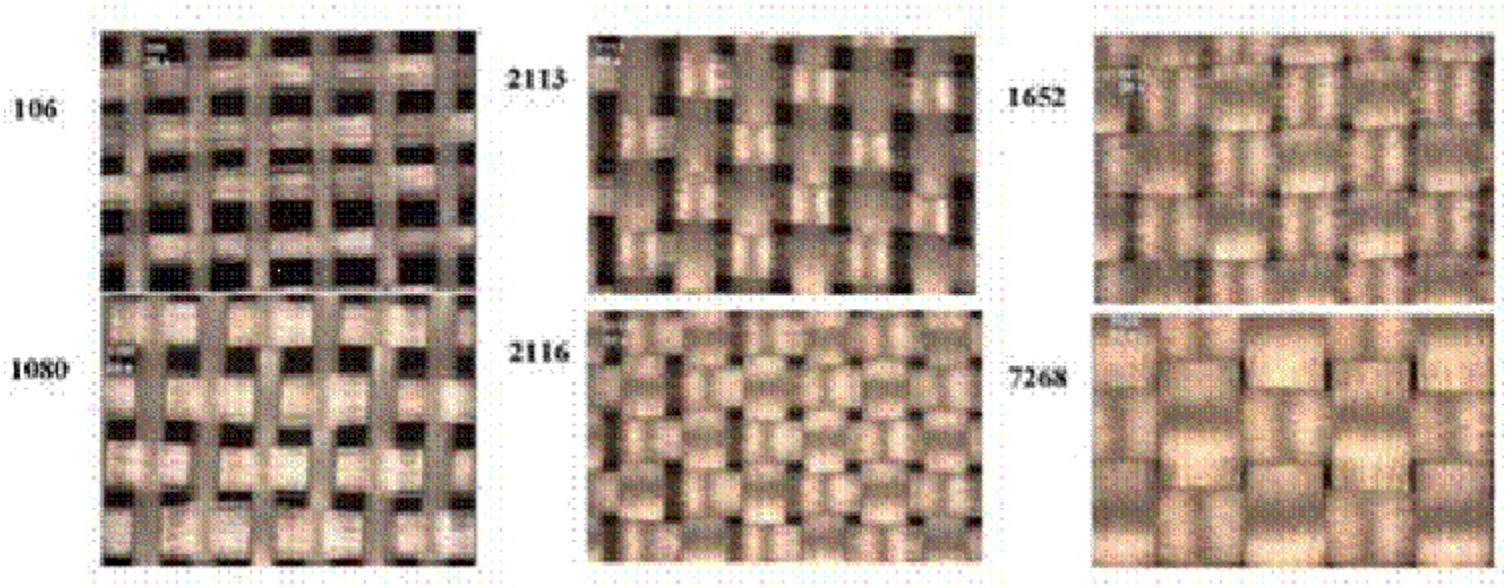

Since the dielectric constant of the fiber bundles and epoxy resin differ, the density of the fiber weaves affects the uniformity of the dielectric constant of the laminate. Sparsely woven cores and prepregs have less uniform dielectric constant and can result in impedance and signal velocity variations along the trace route. For example, consider a differential pair trace routed over 7268 style fiberglass vs. the same pair routed over 106 style fiberglass. The 7268 glass provides a more uniform dielectric constant since the copper traces always reside over the glass. However, traces routed over 106 style glass can result in having a differential pair with one leg routed over glass while the other alternates between glass and resin. The main concern is each leg having different impedances and skew in the differential signal. Although there are routing techniques such as using zig-zag, jogged routing, and placing the complete board design at a slight cant on the PCB panel to average out the on-weave/off-weave problem and mitigate the signal velocity variation, these techniques trade off valuable board real estate while not fully resolving the impedance variation and reflection problem. Additionally, these compensation techniques increases the trace lengths, causing additional loss.

Figure 5: Zig-zag, Jogged Routing, and Canted PCB Compensation Technique

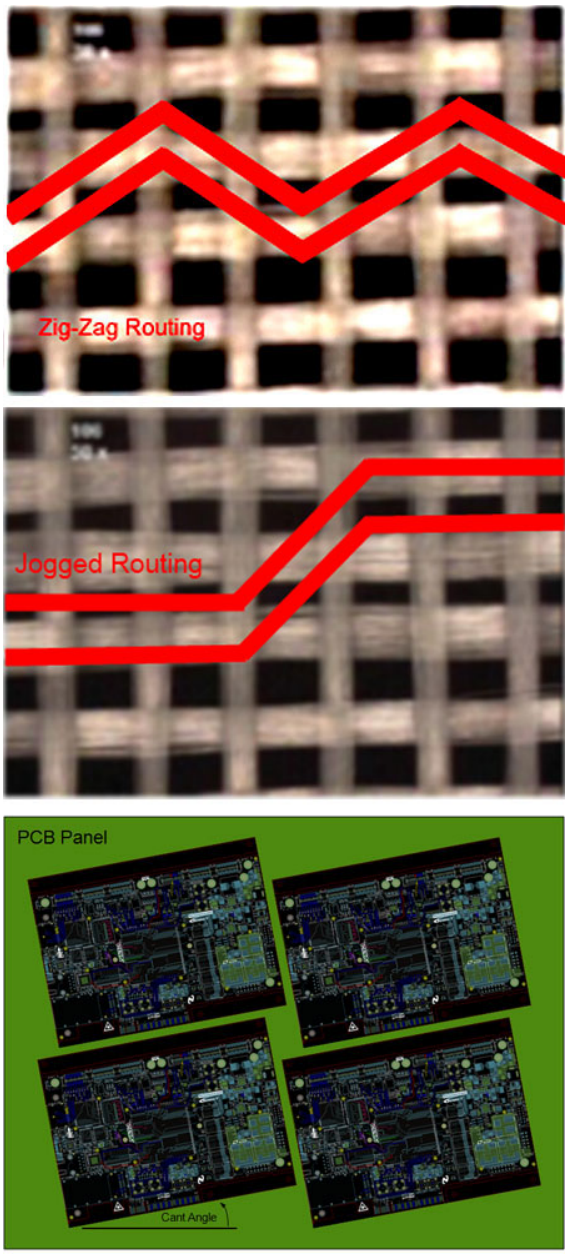

As a result, for very high-Gbps data rates, select higher density fiberglass laminates for the high speed routing layers while relegating sparsely dense fiberglass to lower speed and power layers to offset the PCB cost.

Note: Always choose more densely woven fiberglass style for the core and prepreg material surrounding the high speed signal layers for more uniform ε_r that minimizes impedance and signal velocity variations.

Note: Choose sparsely dense fiberglass styles for power layers and slower general purpose signal layers to reduce PCB cost.

Copper Surface Roughness

In addition to dielectric loss, material losses also result from conductor loss. Conductor loss is the resistive attenuation on the copper conductor. Resistive loss is usually mitigated by using appropriately wide traces. However, for very high frequency designs, the majority of the current distribution is pushed towards the outer surface of the copper conductor as a result of the skin effect. At frequencies where the skin depth approaches the average roughness of the copper, the current flow is further impeded by the copper surface roughness, causing increased resistance and conductor loss.

Figure 6: Skin Effect and Copper Surface Roughness

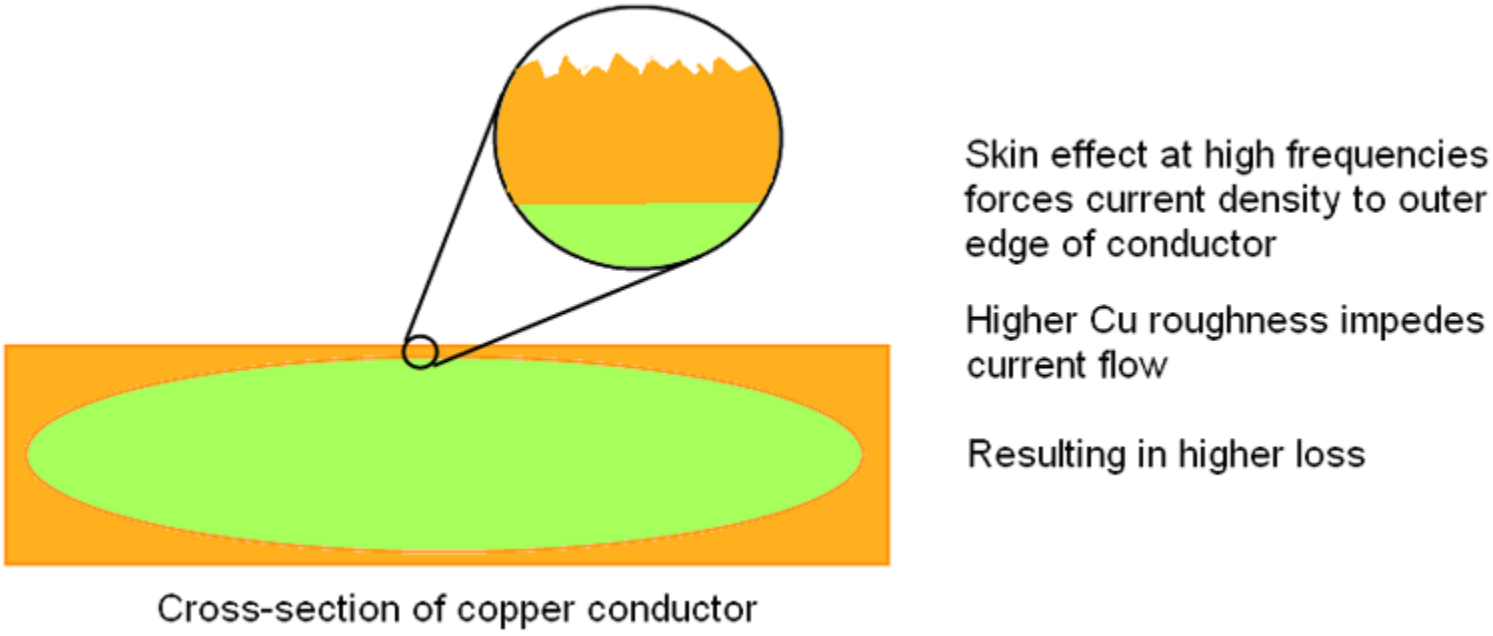

The roughness of the copper surface varies depending on the construction of the copper foils. These foils are either electrodeposited (commonly referred to as ED copper) or rolled and pressed to create a smoother copper foil surface. The copper roughness is specified as an average value (R_a) in micrometers (µm). As a comparison, typical ED copper roughness has R_a of 1 µm or more while rolled copper ranges from 0.3 to 0.4 µm. The effect of this copper roughness can be approximated by a correction factor (K_{SR}) to the attenuation of the conductor (α_{Cond}).

Figure 7: Copper Roughness Approximation Equations

$$\alpha'_{Cond} = \alpha_{Cond} \times K_{SR}$$

$$K_{SR} = 1 + \frac{2\tan^{-1}\left[1.4\left(\frac{R_a}{\delta}\right)\right]}{\pi}$$

$$\delta = \frac{1}{\sqrt{\sigma \cdot \pi \cdot \mu_0 \cdot \mu_r \cdot f}}$$

Where:

α_{Cond} = Conductor attenuation without copper surface roughness correction

α'_{Cond} = Conductor attenuation with copper surface roughness correction

K_{SR} = Surface Roughness correction factor

R_a = Average roughness usually specified in micrometers. For the K_{SR} equation above, R_a must be entered in meters.

δ = skin depth, in meters

σ = conductivity of copper conductor = 5.8×10^7 seimens/m

μ_0 = permeability of free space = $4\pi \times 10^{-7}$ H/m

μ_r = relative permeability of copper conductor = 1

f = frequency, in Hz

Figure 8: Attenuation Effects of Copper Surface Roughness of 6-inch Microstrip Test Trace

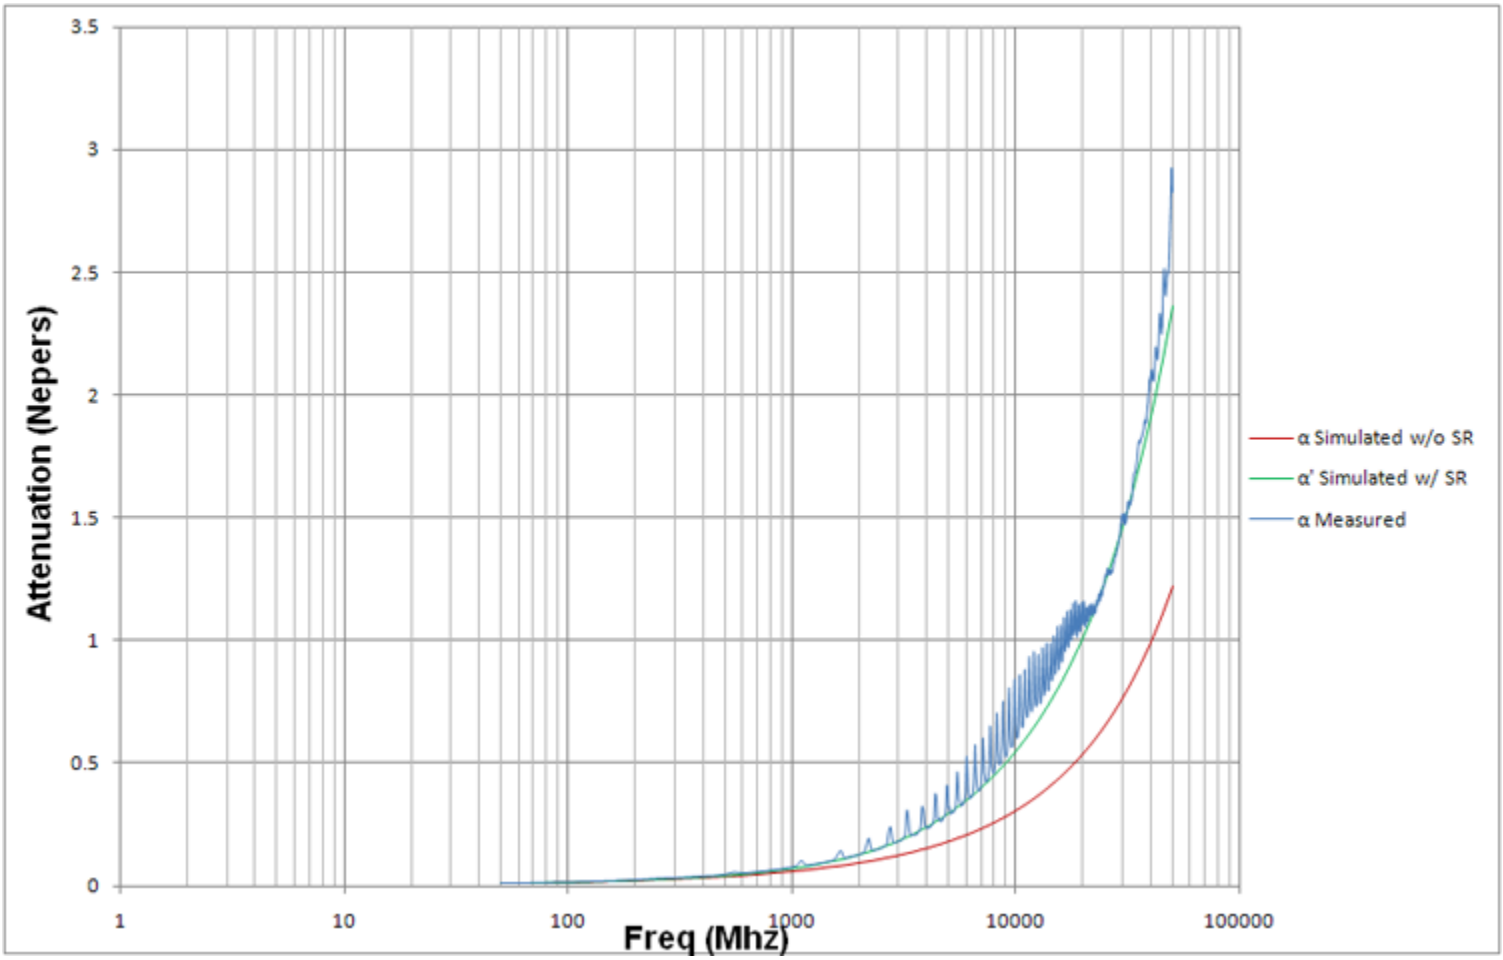

The red curve is the high frequency structural simulator (HFSS) simulated attenuation of the conductor (α_{Cond}) without the effects of surface roughness correction. The green curve is the same result with the surface roughness correction factor (K_{SR}) included. The blue curve is the actual vector network analyzer (VNA) measurement of the same test trace for correlation purposes.

Note: To mitigate conductor loss, use wider traces and choose rolled copper foils over traditional electro-deposited (ED) copper foils in the PCB construction.

Note: Certain simulation tools may not include loss effects from surface roughness. In these cases, the correction factor (K_{SR}) must be added to get realistic prediction of the actual loss.

Related Information

- **PCB Dielectric Material Selection and Fiber Weave Effect on High-Speed Channel Routing Application Note**
- **Modeling Copper Surface Roughness for Multi-gigabit Channel Designs**

Stackup Design

The assignment of critical high-speed routing layers within the PCB stack-up is a critical part of the design decision. The assignment of high-speed signal layers within the stackup directly affects the signal performance. Signals routed on external layers of the PCB board are referred to as microstrip, while those routed on internal layers are called stripline.

Figure 9: Differential Microstrip vs. Stripline Construction

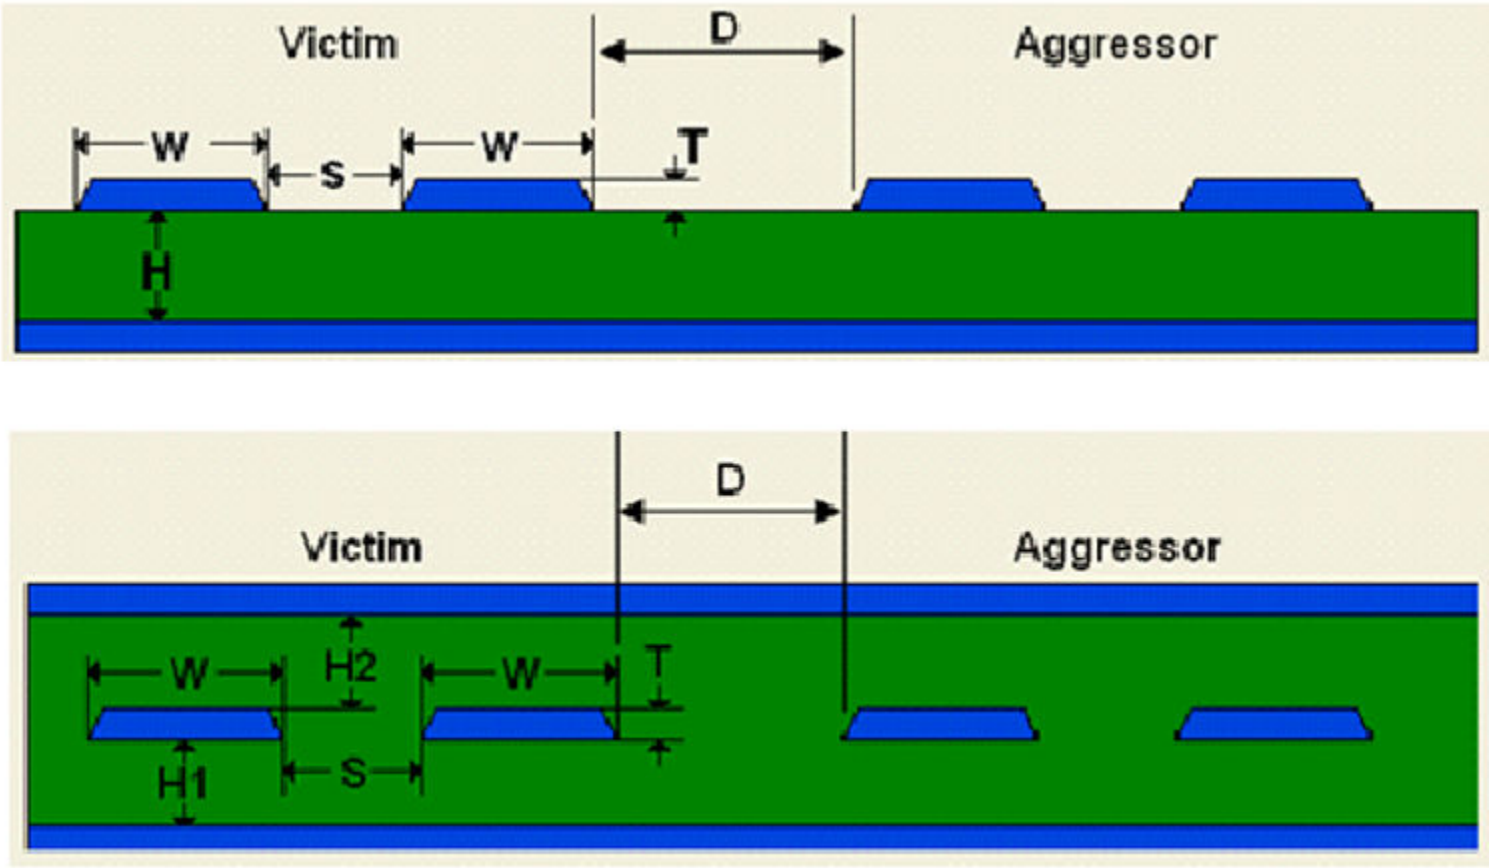

By manipulating the various parameters such as trace width (W), separation (S), and height from the reference plane (H for microstrip and H1, H2 for stripline), the trace impedance can be adjusted appropriately. Additionally, edge-coupled crosstalk from neighboring traces can be well-controlled by adjusting the pair separation (D). For more information on crosstalk, refer to **Crosstalk Control** on page 20.

Table 2: Microstrip vs. Stripline

Topology	Advantages	Disadvantages
Microstrip	• Thinner dielectric for 100-Ω traces • No via for top routing • No via stubs for bottom routing	• Only 2 routing layers possible • Higher EMI radiation • Both near-end crosstalk (NEXT) and far-end crosstalk (FEXT) concerns
Stripline	• Many routing layers possible • Inherent EMI shielding surrounding layers • No FEXT concerns	• Thicker dielectric required for 100 Ω • NEXT concerns • Via must be used • Via stubs require back-drilling

The decision to use one topology over the other examines the first-order factors that affect signal bandwidth. While impedance and crosstalk can be well-controlled in both routing topologies, stripline provides lower signal attenuation vs. microstrip for the same trace width and copper thickness.

Figure 10: Stripline vs. Microstrip Insertion Loss (Sdd21)

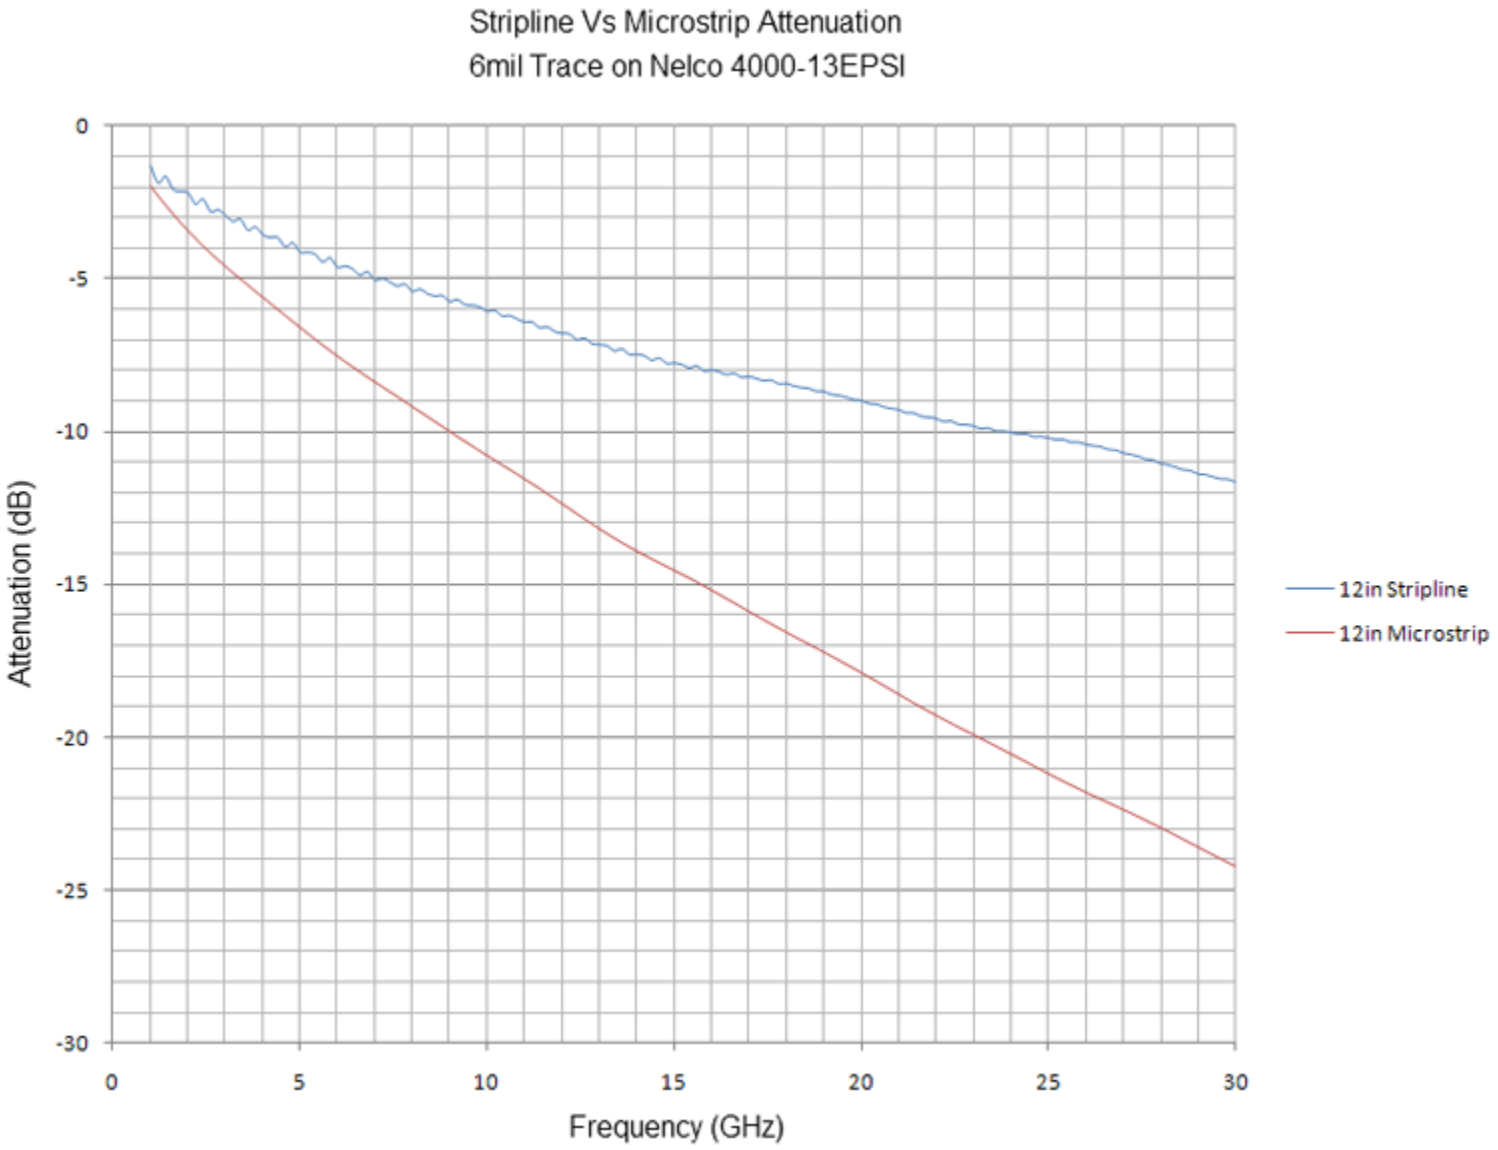

Figure 11: Stackup Construction

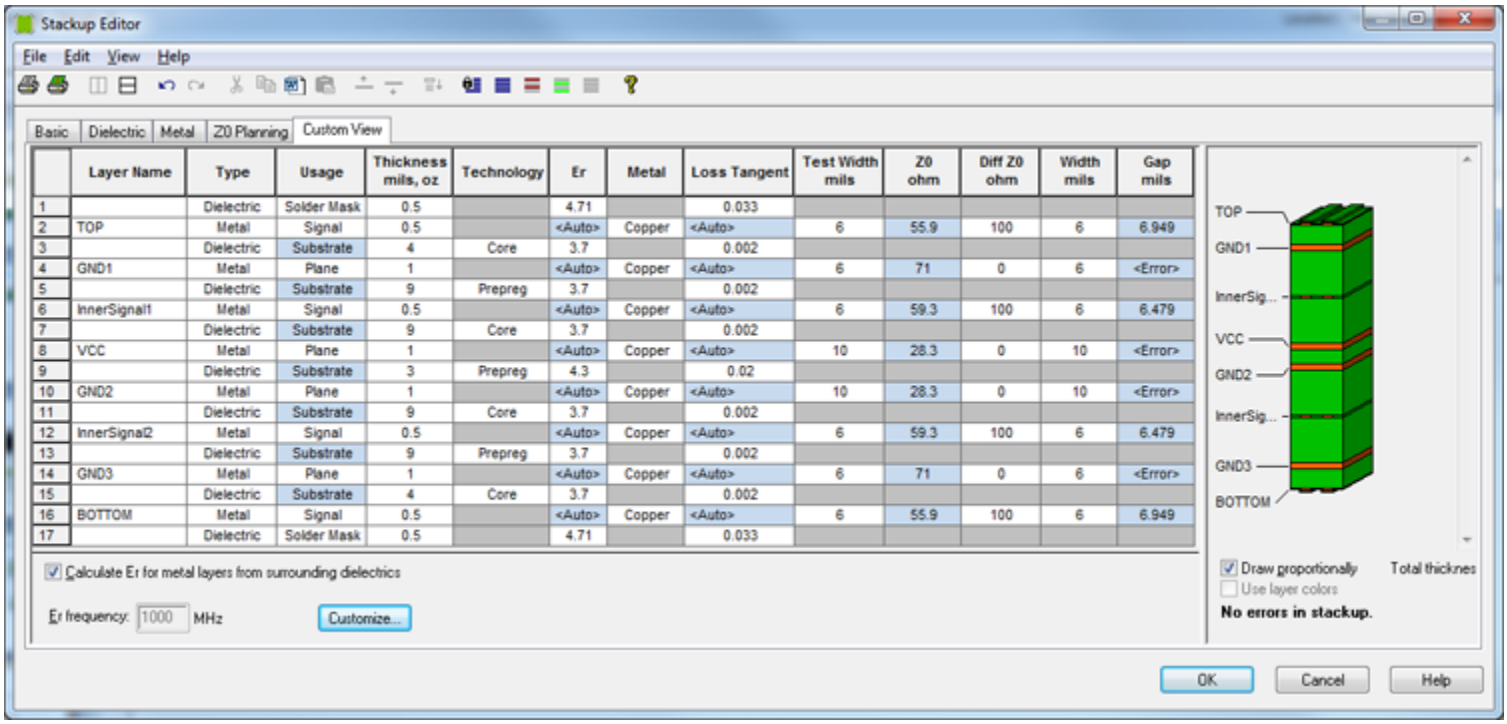

Note: For the same trace width and copper thickness considerations, stripline results in less signal attenuation compared with microstrip.

Related Information

PCB Stack-up Design Considerations for Altera FPGAs Application Note

Channel Design

The channel is the physical medium for sending data from a transmitting device to a receiving device. The five common components of a typical transceiver channel construction include:

- BGA channel breakout of the sending and receiving device
- Trace route design
- Signal via
- DC blocking capacitors
- Interface connectors.

Any channel usually involves two or more of these components. To optimize the channel for best signal performance requires careful consideration of each of these components to minimize their first-order impact on the channel.

Figure 12: Channel Design Considerations

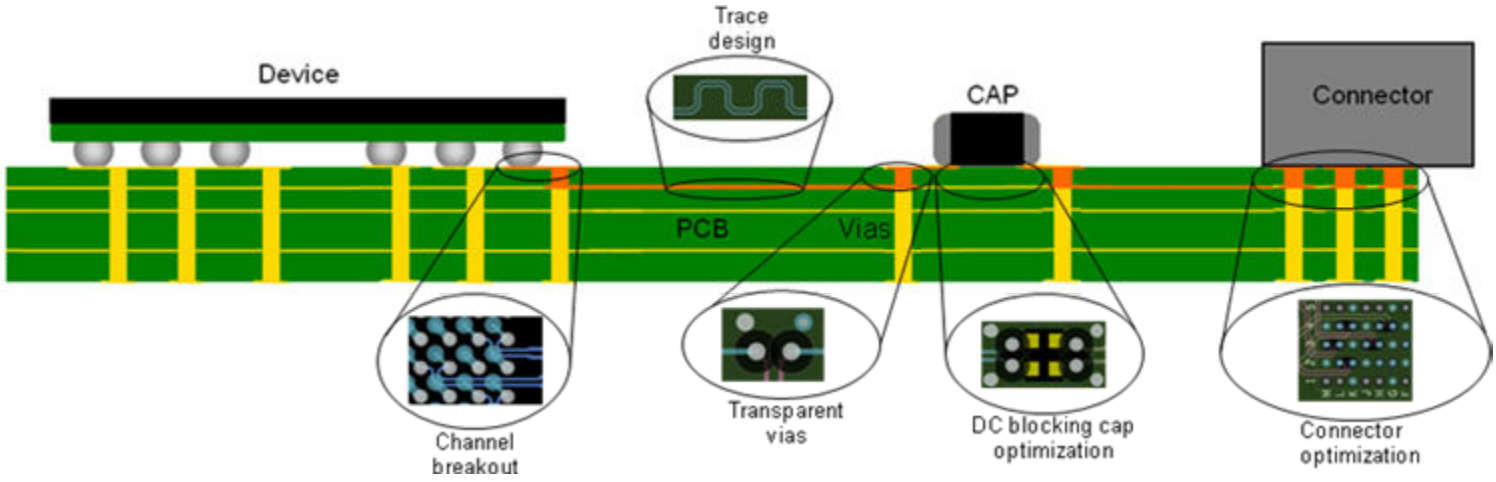

BGA Channel Breakout

As the number of very high-speed transceiver pairs on FPGAs continue to increase, the complexity of the channel breakout design increases as well. Trace breakout can typically use either a single trace or dual trace topology between the BGA via grid.

Figure 13: Single Trace Breakout

Routing layers are separated by a GND plane.

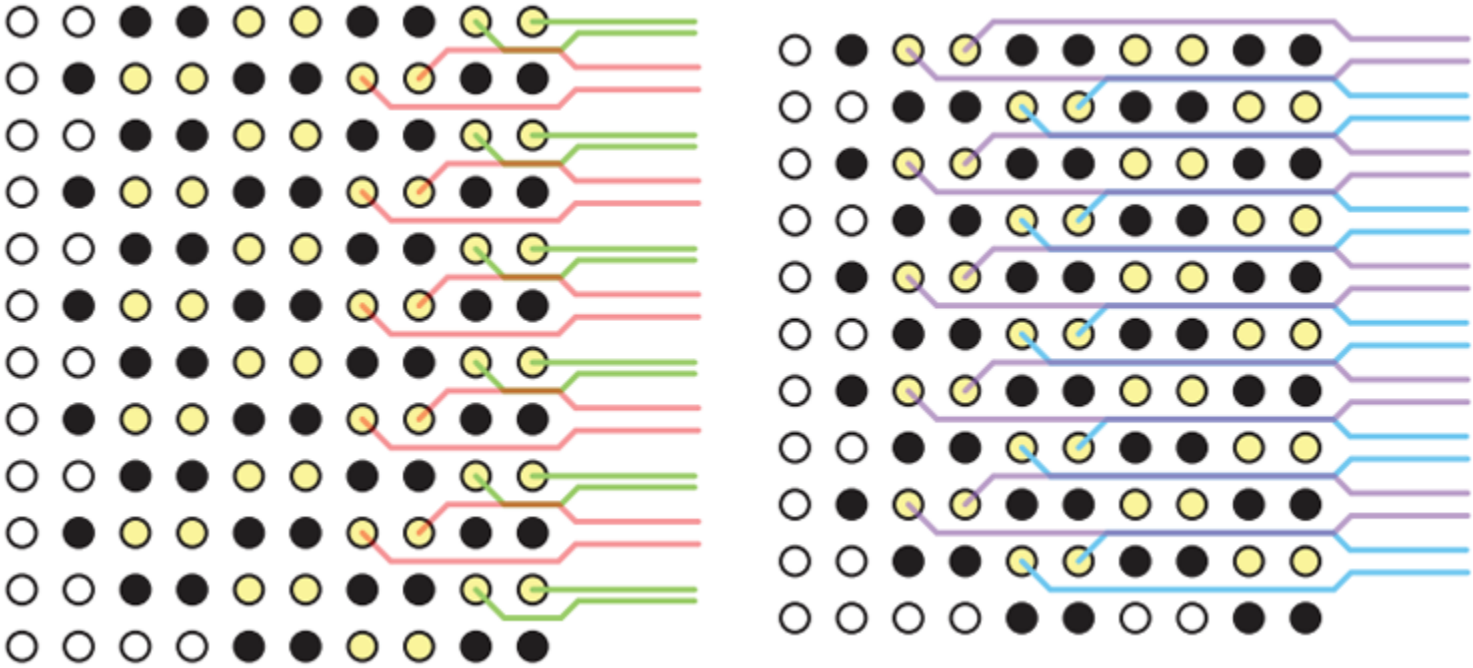

Figure 14: Dual Trace Breakout

Routing layers are separated by a GND plane.

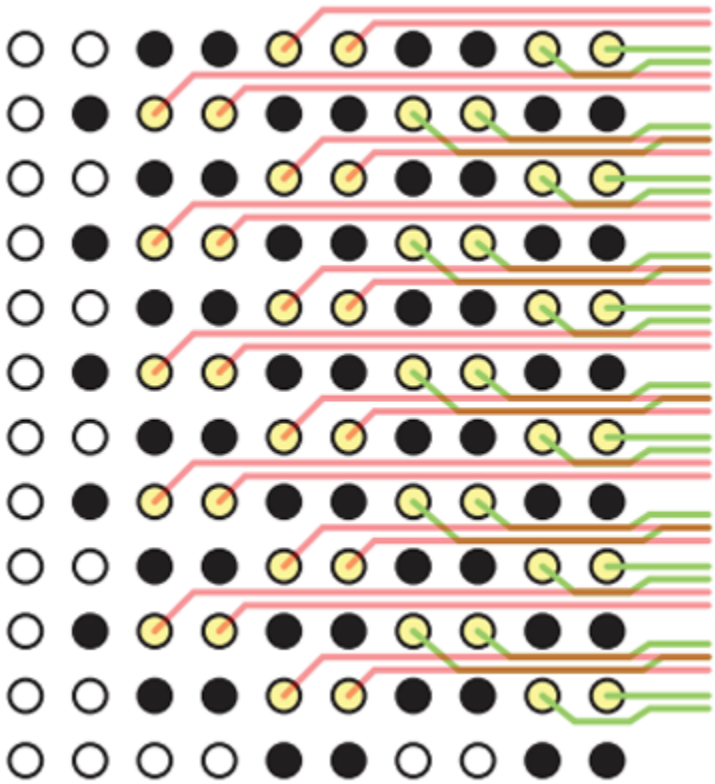

In the single trace breakout, more layers are required to fully route all the TX and RX transceiver pairs because only one trace is routed between the BGA via grid per layer. Because there is ample space for a maximum trace width of up to 13.37 mils (assuming typical minimum trace-to-copper clearances of 4 mils for the recommended 18 mil via pad size and 39.37 mil via-to-via center pitch), the trace width and resulting characteristic impedance is easily made uniform throughout the trace route.

In the dual trace breakout, two traces are used to reduce layer count requirement, but their maximum trace width is limited to 4.68 mils because of the same trace-to-copper clearance requirements. Because high-speed transceiver traces are usually routed using trace widths of 6 to 8 mils (or more) to reduce the impact of skin effect losses at higher frequencies, the reduction of trace width to accommodate dual trace breakout effectively increases the trace impedance. However, this increase is offset by the reduction of the trace-to-trace separation in the differential pair, so the net impedance remains unchanged.

Table 3: Net Impedance Effect of Trace Neck Down

Topology	Height from Reference Plane (H1/H2 in mils)	E_r	Trace Width (W in mils)	Trace Separation (S in mils)	Diff Zo (Ω)
Stripline	9/9	3.7	6	6.5	100

Topology	Height from Reference Plane (H1/H2 in mils)	E_r	Trace Width (W in mils)	Trace Separation (S in mils)	Diff Zo (Ω)
Stripline	9/9	3.7	4	4	101

Another benefit of the trace neck down is that the thinner trace width increases the trace inductance and helps to compensate for the higher capacitance of the BGA ball pad.

Related Information

PCB Breakout Routing for High-Density Serial Channel Designs Beyond 10 Gbps

Channel Routing Design

The channel routing design involves the selection of critical trace parameters that can affect performance.

Trace Width Selection

PCBs are becoming very constrained, and using fine trace widths down to 4 mils wide is commonly used to improve routability. However, for high-speed signals, narrow trace geometries increases conductor loss due to skin effect. As a result, routabilty must be properly balanced with trace width selection for better performance. Larger signal attenuation occurs for signals at 4 mils vs. 5 to 6 mils. For example, for a 28-Gbps signal, the difference in attenuation at the Nyquist frequency (14 MHz) is approximately 3 dB for 4 mils versus 6 mils wide trace.

Figure 15: Trace Width vs. Signal Attenuation

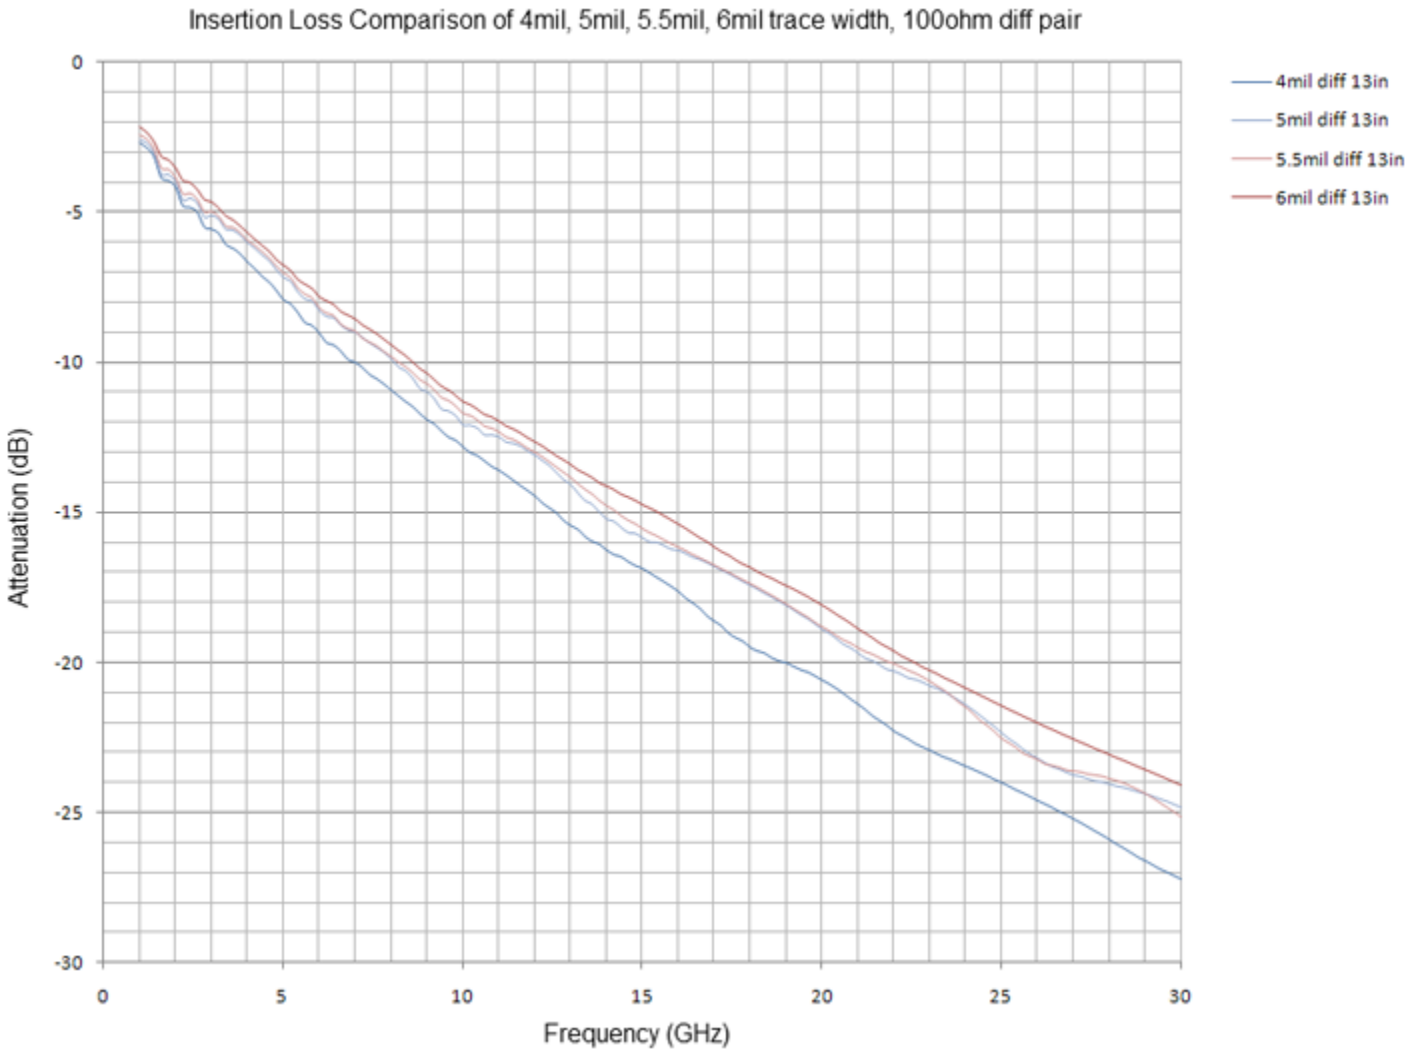

Note: For high-speed transceiver signals, use trace widths of 6 mils or more to reduce conductor loss.

Note: Limit the use of 4 mil trace widths to the BGA breakout area and keep their trace length as short as possible.

Loose vs. Tight Coupled Traces

The decision to use loose vs. tight coupling is mainly a trade-off between routing density and impedance control.

Table 4: Loosely vs. Tightly Coupled Trace Routing

Routing	Advantage	Disadvantage
Loosely Coupled	• Thinner dielectrics required for the same trace width • Less sensitivity to trace-to-trace variations provides better impedance control	Consumes more area vs. tight coupling
Tightly Coupled	• Higher routing density • Smaller trace width for the same trace impedance • Better common mode noise rejection	Impedance control highly sensitive to trace-to-trace variations

For example, differential pair length matching typically requires serpentining of one leg of the differential pair to maintain P to N skew. For loosely coupled traces, the serpentining does not drastically alter the differential impedance of the trace. However, for tightly coupled traces, the change in the trace-to-trace separation can significantly change the nominal differential impedance of the pair beyond the ±10% tolerance. When applying serpentine routing, it is best to deskew after each bend or node that causes the trace lengths to be mismatched. Doing so helps reduce common mode noise incurred along the signal path.

Figure 16: Differential Pair Length Matching

Table 5: Loosely vs. Tightly Coupled Routing Impedance Control

Routing Topology	Dielectric Constant (E_r)	Trace Width (W)	Trace Separation (S)	Height above reference plane (H)	Zdiff (Ω)
Loosely coupled microstrip	3.7	6 mils	12 mils	4 mils	100
Loosely coupled microstrip	3.7	6 mils	18 mils	4 mils	102
Tightly coupled microstrip	3.7	6 mils	6 mils	4.8 mils	100
Tightly coupled microstrip	3.7	6 mils	12 mils	4.8 mils	**112**

Note: Loosely coupled traces are easier to route and maintain impedance control but take up more routing area. Tightly coupled traces saves routing space but can be difficult to control impedance.

Crosstalk Control

Crosstalk is induced noise current resulting from mutual capacitive (C_m) and mutual inductive (L_m) coupling on a victim trace due to switching activity from nearby aggressor trace or traces. The current coupled from C_m travels along the victim trace in both the forward and reverse direction with the same polarity. Similarly, the current from Lm travels forward and backwards in opposite polarity. As a result, crosstalk can be separated into two distinct components referred to as near-end crosstalk (NEXT) and far-end crosstalk (FEXT). In NEXT, the coupled noise current is the sum of the induced currents from C_m and L_m as those currents are the same polarity. Conversely, for FEXT, the current is the difference of C_m and L_m due to the polarity difference. For signals entirely contained within a homogeneous dielectric material (such as stripline), the capacitive and inductive forward crosstalk are equal and cancel. For non-homogeneous dielectrics (such as microstrip), the inductive component tends to be larger and the resulting coupled noise is negative.

Figure 17: NEXT and FEXT Coupling Components

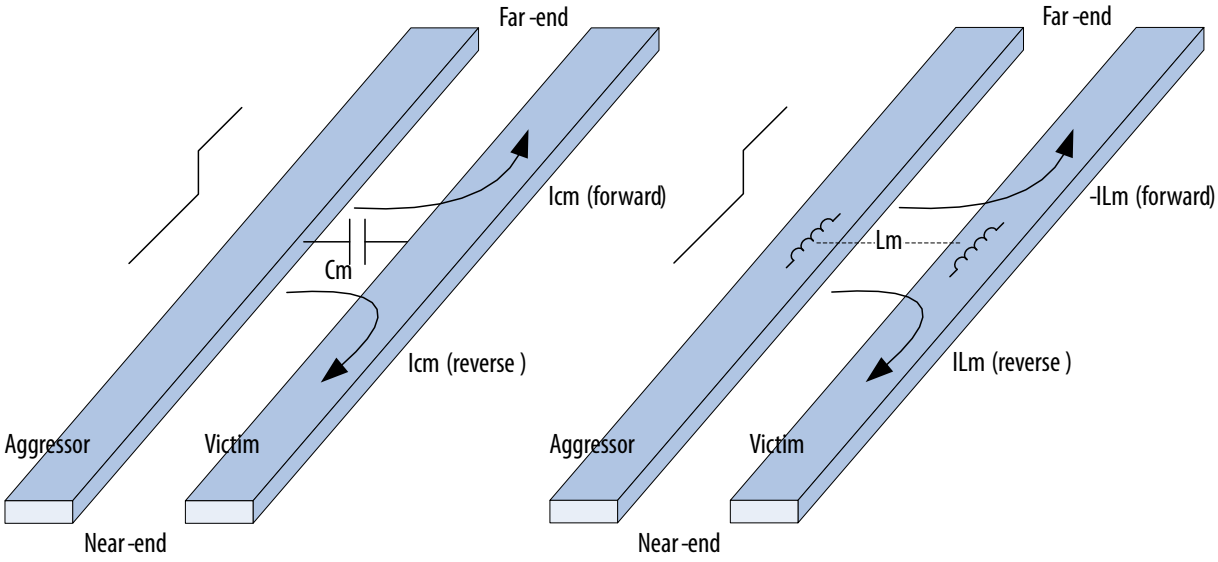

Crosstalk control usually involves reducing signal edge rates and maintaining enough trace-to-trace separation to reduce the mutual capacitive and mutual inductive coupling energy. In high-speed transceiver designs running at many gigabits per second, reducing the signal edge rate is usually not an option since the unit interval time (UI) is very small. Therefore, crosstalk control for high-speed transceiver designs is mainly determined by PCB layout spacing constraints to keep the transceiver traces far enough apart to minimize the coupling effect. For very high-speed traces, it is desirable to keep the coupling noise to less than 1% of the source signal if possible.

Edge Coupling

Consider two differential pairs routed on the same signal layer as shown in **Figure 9** with a height (H) from the reference plane.

Table 6: Microstrip NEXT vs. Aggressor Separation

Microstrip routing requires a separation of 6H and 7H to properly manage the crosstalk coupling to less than 1% for NEXT and FEXT, respectively. However, to achieve 1% of crosstalk coupling using stripline routing, requires only 5H separation for NEXT and 2H separation for FEXT.

Separation (D)	Isolation (dB)	Coupled Voltage from a 1-V Aggressor (mV)	Coupling (%)
1H	22	79	7.9
2H	27	45	4.5
3H	31	28	2.8
4H	35	18	1.8
5H	36	16	1.6
6H	**41**	**9**	**0.09**

Table 7: Microstrip FEXT vs. Aggressor Separation

Microstrip routing requires a separation of 6H and 7H to properly manage the crosstalk coupling to less than 1% for NEXT and FEXT, respectively. However, to achieve 1% of crosstalk coupling using stripline routing, requires only 5H separation for NEXT and 2H separation for FEXT.

Separation (D)	Isolation (dB)	Coupled Voltage from a 1-V Aggressor (mV)	Coupling (%)
6H	34	20	2.0
7H	**41**	**9**	**0.09**

Table 8: Stripline NEXT vs. Aggressor Separation

Separation (D)	Isolation (dB)	Coupled Voltage from a 1-V Aggressor (mV)	Coupling (%)
1H	24	63	6.3
2H	32	25	2.5
3H	36	16	1.5
4H	39	11	1.1
5H	**40**	**10**	**1.0**

Table 9: Stripline FEXT vs. Aggressor Separation

Separation (D)	Isolation (dB)	Coupled Voltage from a 1-V Aggressor (mV)	Coupling (%)
1H	32	25	2.5
2H	**40**	**10**	**1.0**

Note: Use stripline routing to avoid FEXT concerns. Use stripline traces with 5H differential pair-to-pair separation to minimize NEXT to 1%.

Note: If microstrip routing is required, used 6H-7H differential pair-to-pair separation to avoid NEXT and FEXT issues.

Broadside Coupling

For differential pair traces routed on adjacent signal layers, broadside coupling is much stronger if there is any overlapping of traces. This problem usually occurs under the congested BGA area where routing is constrained.

Figure 18: Broadside Coupled Traces

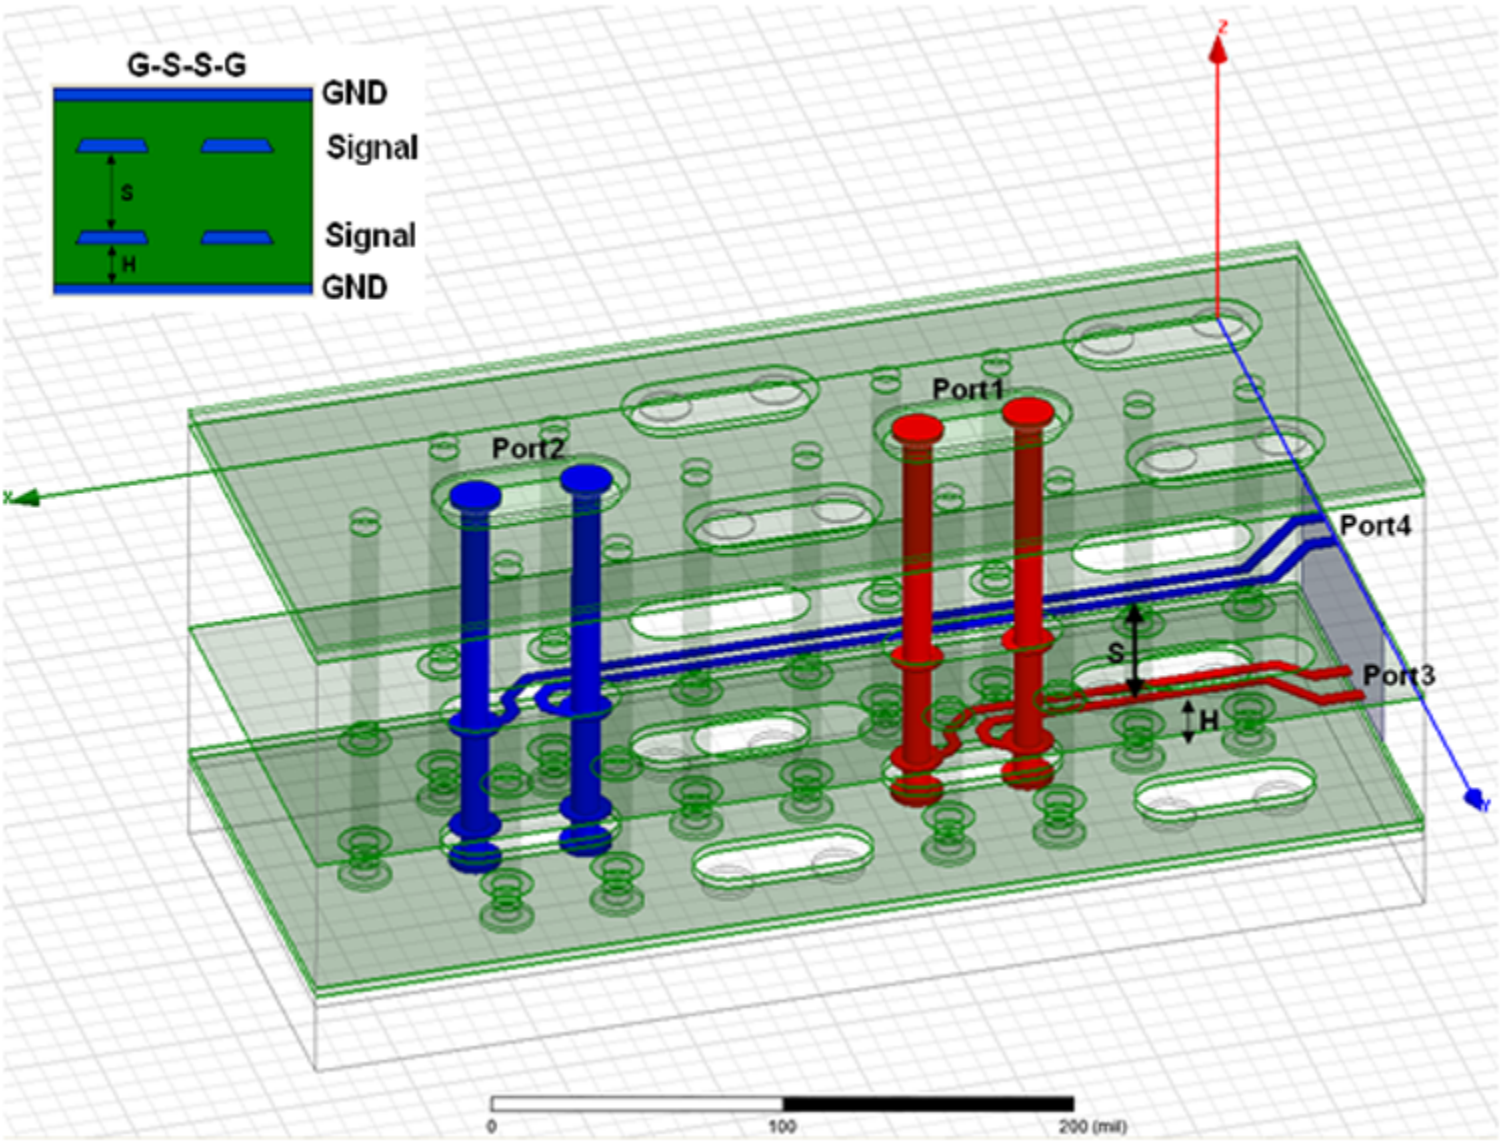

Simply increasing the vertical separation of broadside coupled traces is not practical because the required vertical separation (S) would need to be 10H for 1% coupling. This causes board thickness to quickly become unreasonable. Instead, broadside crosstalk minimization techniques should follow edge coupling rules by maintaining a 5H horizontal separation (D) between pairs on adjacent layer. For the congested BGA area, the suggested routing implementation helps to avoid trace overlap.

Figure 19: Broadside Coupling Avoidance

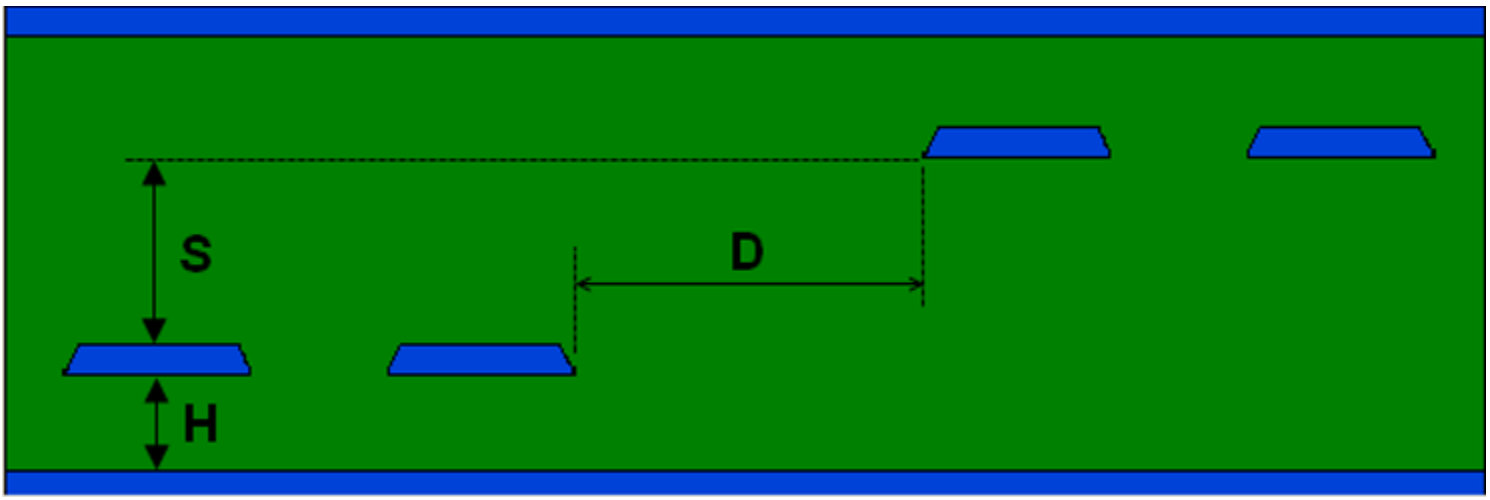

Figure 20: Suggested Routing Under BGA Area

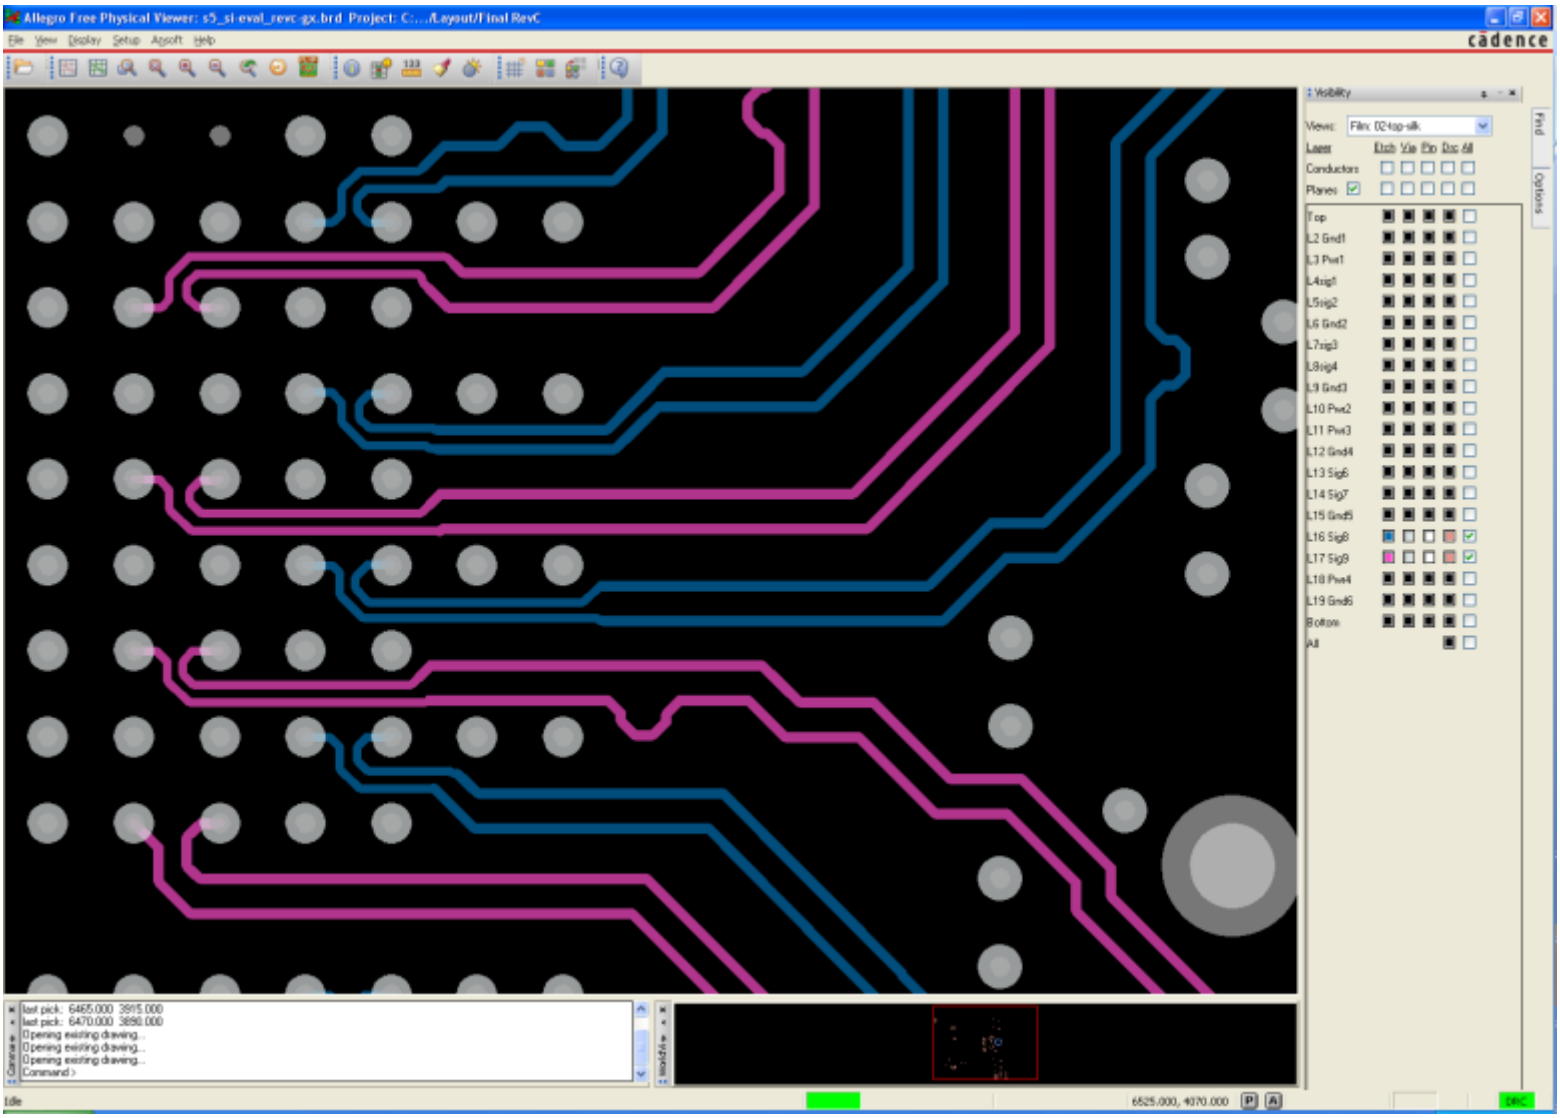

Transparent Via Design

Unless you are routing all transceiver channels with only microstrip traces on the top layer, you must use vias in the design to transition layers. Unfortunately, the characteristic impedance of differential vias are lower than 100 Ω. Generally it is in the range of 80 to 85 Ω. This impedance mismatch causes reflections that degrade the channel performance. To better match the impedance of the via with the 100 Ω differential traces requires optimization techniques that minimize the parasitic capacitance (C_{via}) and inductance (L_{via}) of the via.

Tip: You can minimize C_{via} using the following optimization techniques:

- Reduce the via capture pad size
- Eliminate all non-functional pads (NFP)
- Increase the via anti-pad size to 40 or 50 mils

Tip: You can minimize L_{via} using the following optimization techniques:

- Eliminate and / or reduce via stubs
- Minimize via barrel length by routing the stripline traces near the top surface layer and applying backdrilling

For example, consider an 8-layer PCB board that uses standard via with 10-mil drill diameter, 20-mil capture pad diameter, and 30-mil anti-pad diameter. Optimizing this via by successively applying the techniques above to minimize C_{via} and L_{via} improves the impedance of the via, and its insertion and return loss.

Figure 21: TDR of Standard vs. Optimized Via

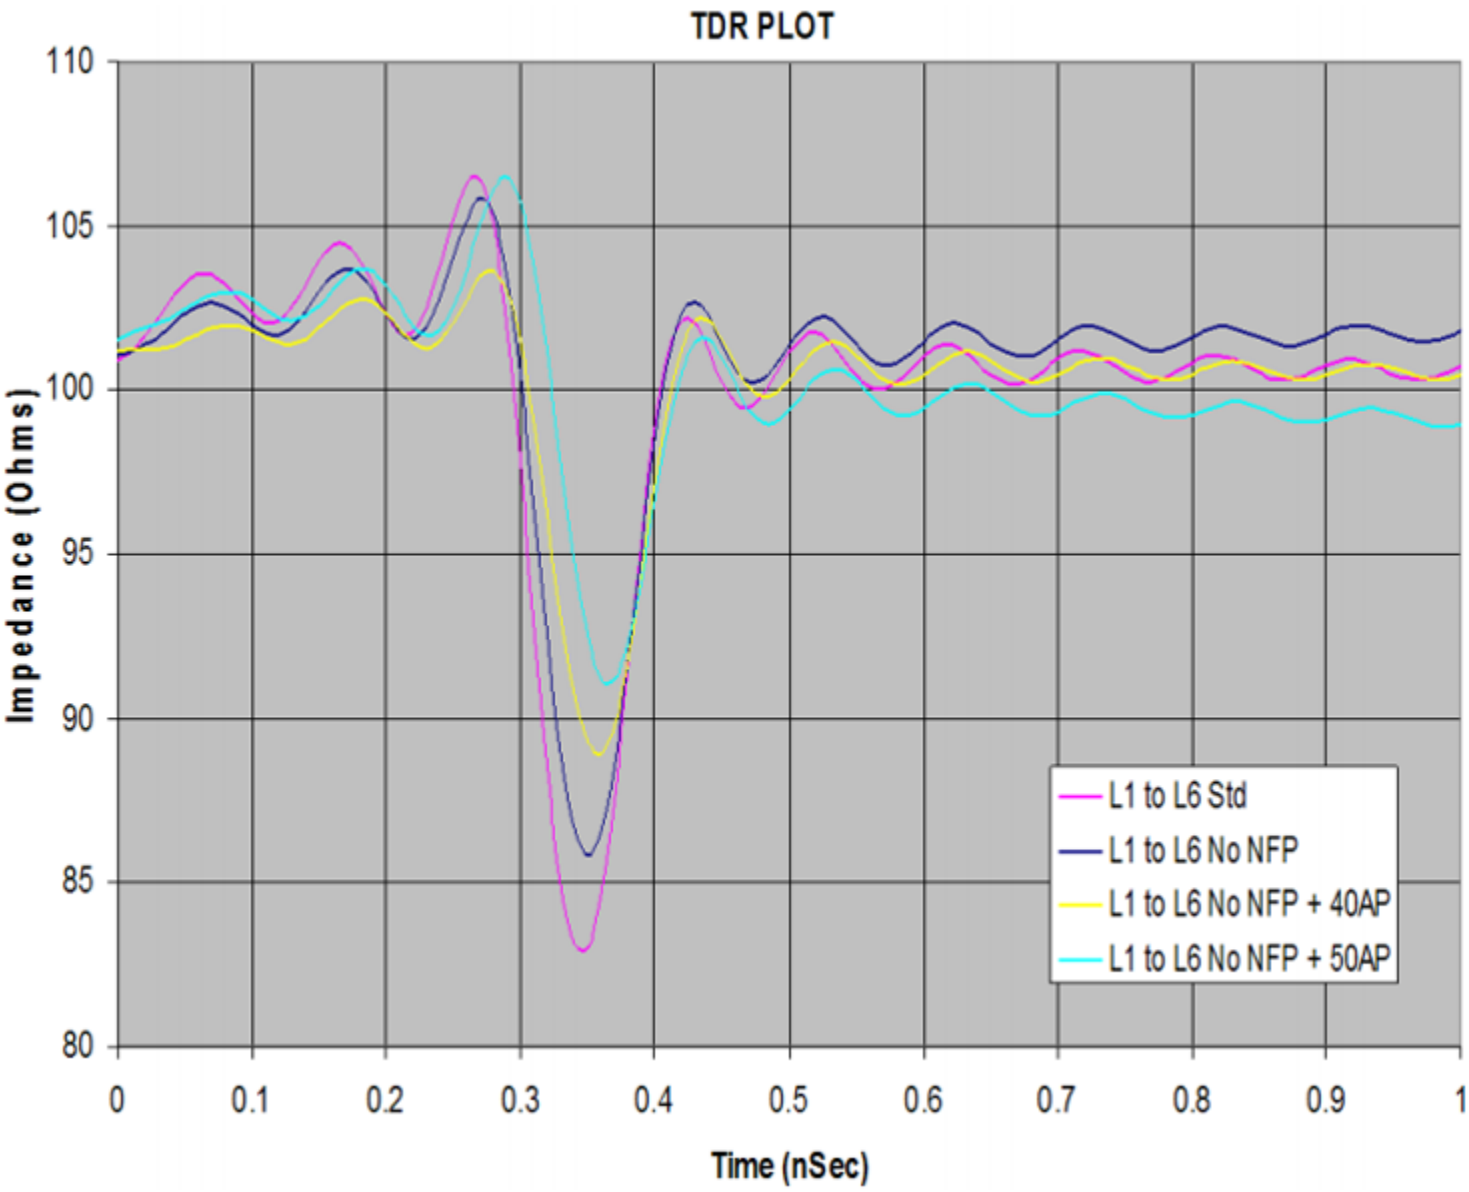

Figure 22: Insertion and Return Loss of Standard vs. Optimized Via

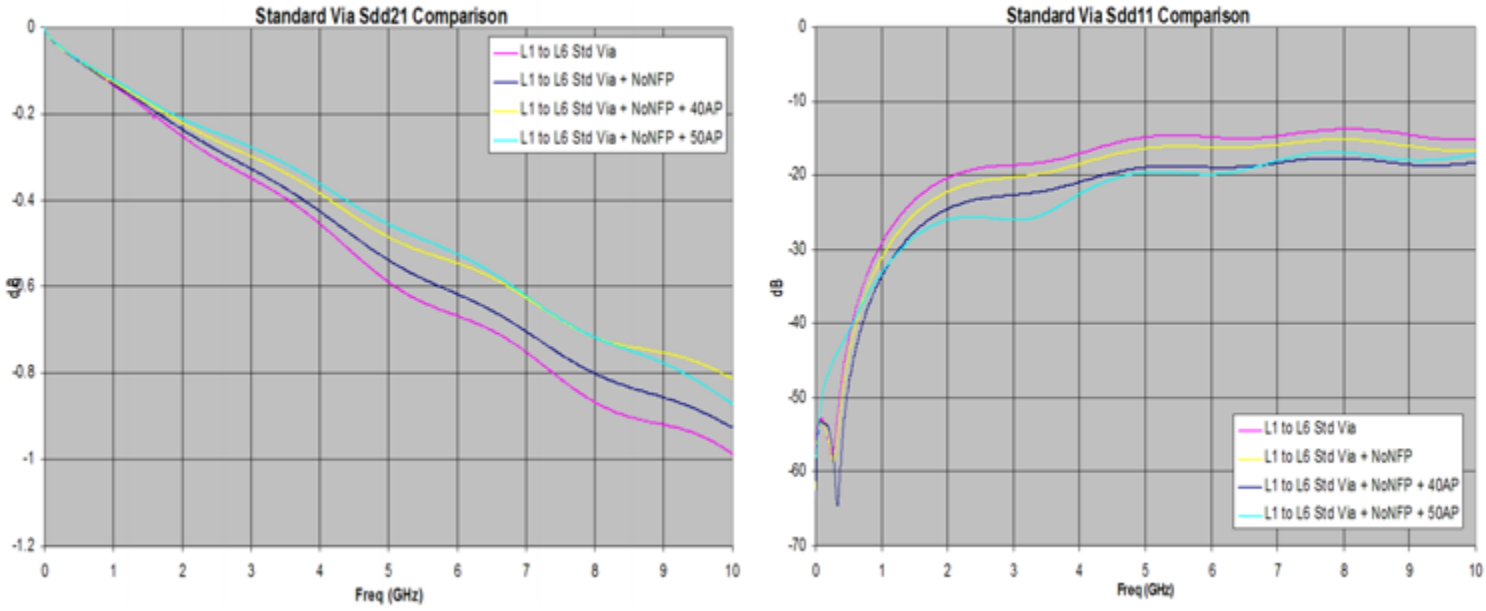

Additional improvements besides minimizing C_{via} and L_{via} involves providing a better AC return path by adding a ground via next to each signal via as well as applying backdrilling to remove any left over via stub.

Note: Add ground return vias within 35 mils of each signal via to further improve the insertion and return losses of the via.

Related Information

Via Optimization Techniques for High-Speed Channel Designs

Blocking Cap Optimization

Transceiver channels often incorporate DC blocking capacitors to control the common mode voltage at the receiver. However, the presence of the blocking capacitors in the channel creates an abrupt discontinuity where the trace meets the capacitor. Similar to via optimization, the layout footprint for the blocking caps can be optimized to minimize their impact on the channel. Because the larger capacitor pad results in lowering its characteristic impedance, one way of increasing this impedance to better match the trace impedance is to increase the distance to the reference by making cut-outs underneath the body of the capacitor footprint.

Figure 23: DC Blocking Capacitor Plane Cut-out

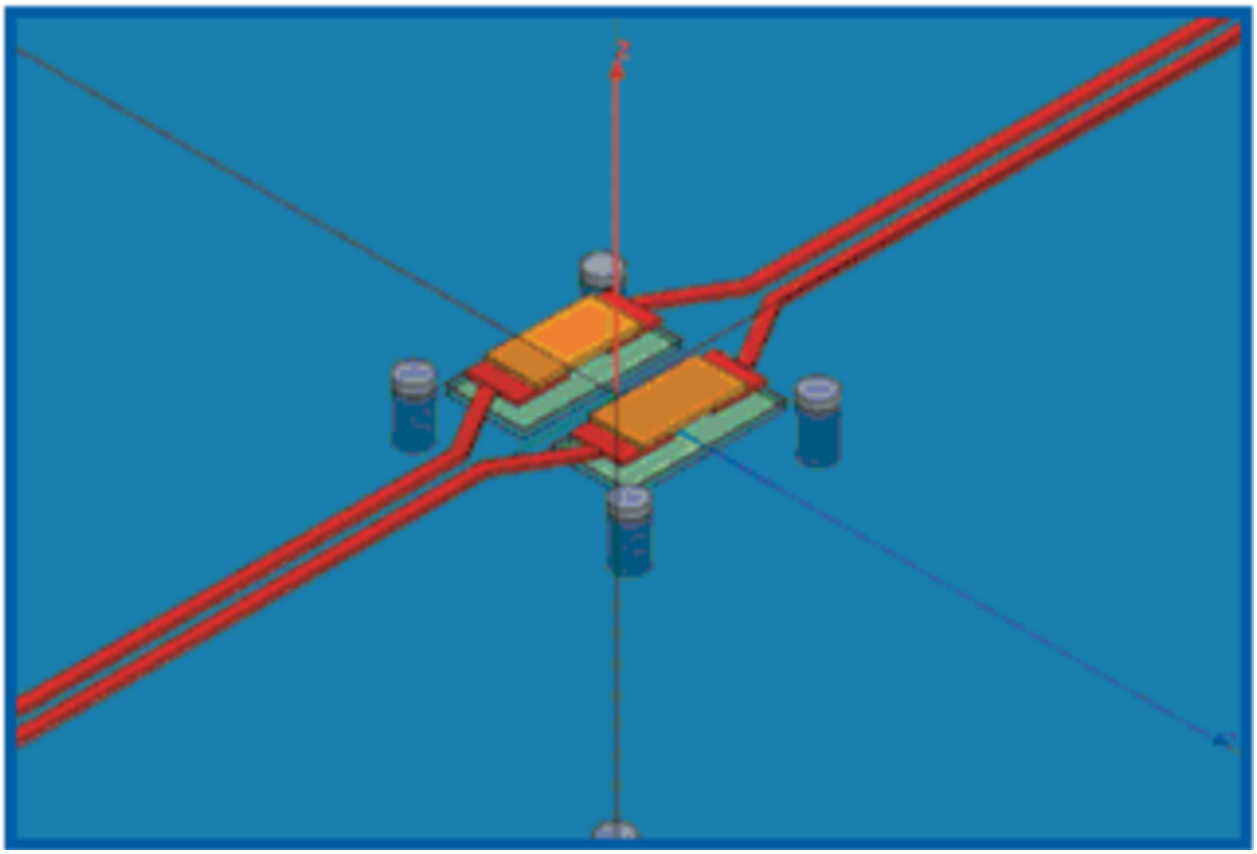

By cutting out the first reference plane directly below the capacitor, the impedance increases as it references the second plane further away. However, if this second reference plane is close to the first reference plane, the increase may still not be enough. In this case it also becomes necessary to cut out the second, third, or even more successive planes underneath to further increase the impedance.

Figure 24: Additional Plane Cut-outs Underneath DC Blocking Cap

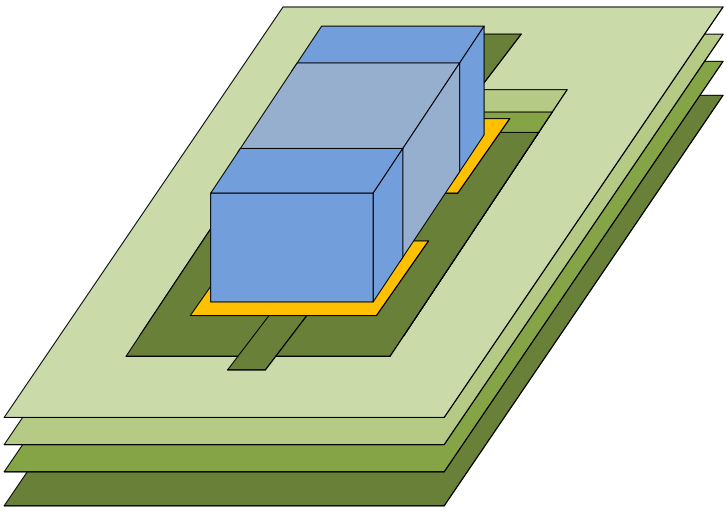

Normally, determining the proper plane cutout size and the number of layers below the capacitor to cut is determined by extensive 3-D simulations. However, a formulaic approach based on simulations for determining this cutout is also possible.

Figure 25: DC Blocking Capacitor Compensation

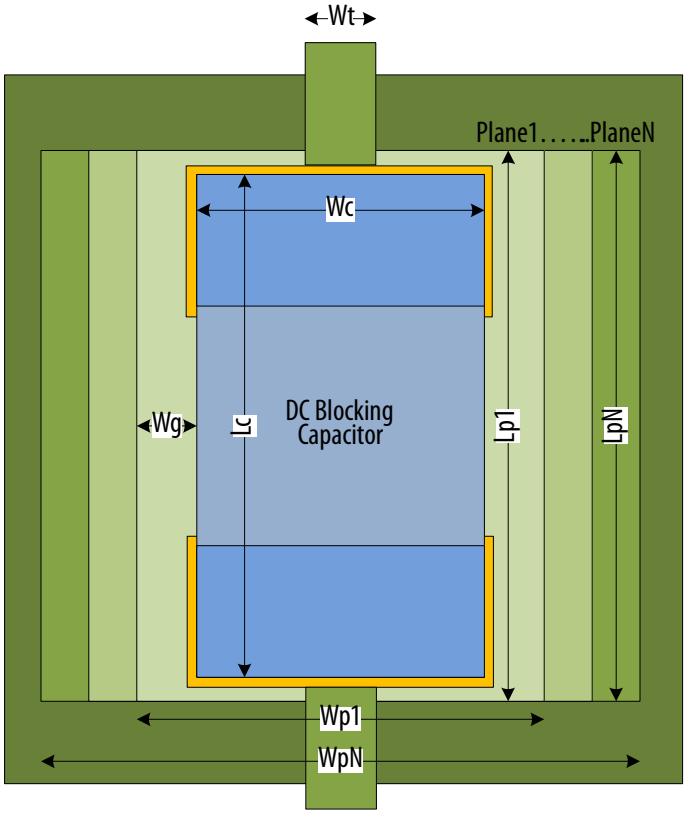

Where,
Wt = Trace width
Wc = Component width
Lc = Component length
Wg = 0.7 x (Wc -Wt)
Wp1 = Wc +2Wg
Lp1 = Lc +2mil
WpN = Wp1 + 10mil x (N-1)
LpN = Lp1

1. Cut out any plane underneath the capacitor whose proximity is within 0.75 Wc.
2. Set the side gap of the cut-out for plane 1 (Wg) = 0.7 (Wc – Wt).
3. Set the cut-out width of plane 1 (Wp1) = Wc+2 Wg.
4. Set the cut-out length of plane 1 (Lp1) = Lc+2 mils.
5. Set the cut-out width of successive plane N (WpN) = Wp1 + 10 (N-1).
6. Set the cut-out length of successive plane N (LpN) = Lp1.

The following example compares the time-domain reflectometer (TDR) results of the DC blocking capacitor layout with and without the plane cutout improvements. With the plane cutouts properly applied using the above guidelines, the large discontinuity at the trace to DC blocking capacitor junction is eliminated.

Figure 26: DC Blocking Capacitor Layout with and without Plane Cutouts

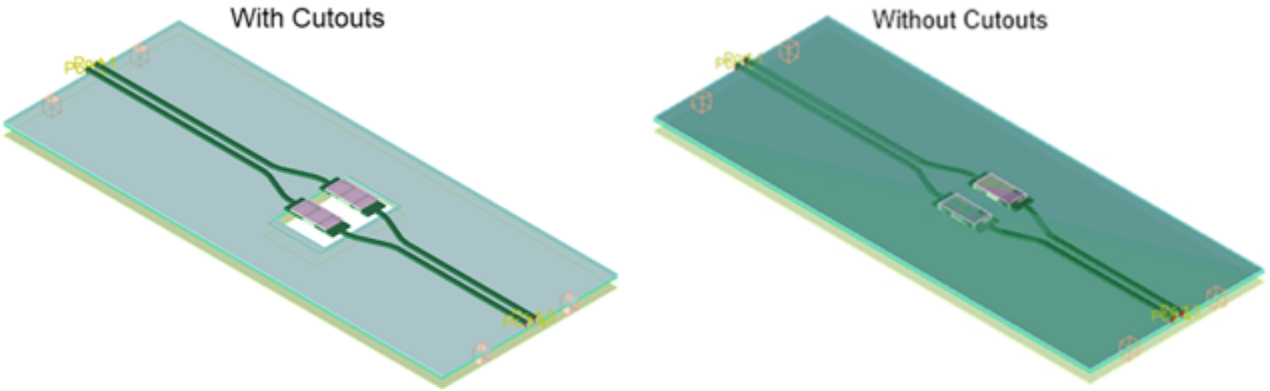

Figure 27: TDR Plot for DC Blocking Capacitor with and without Plane Cutouts

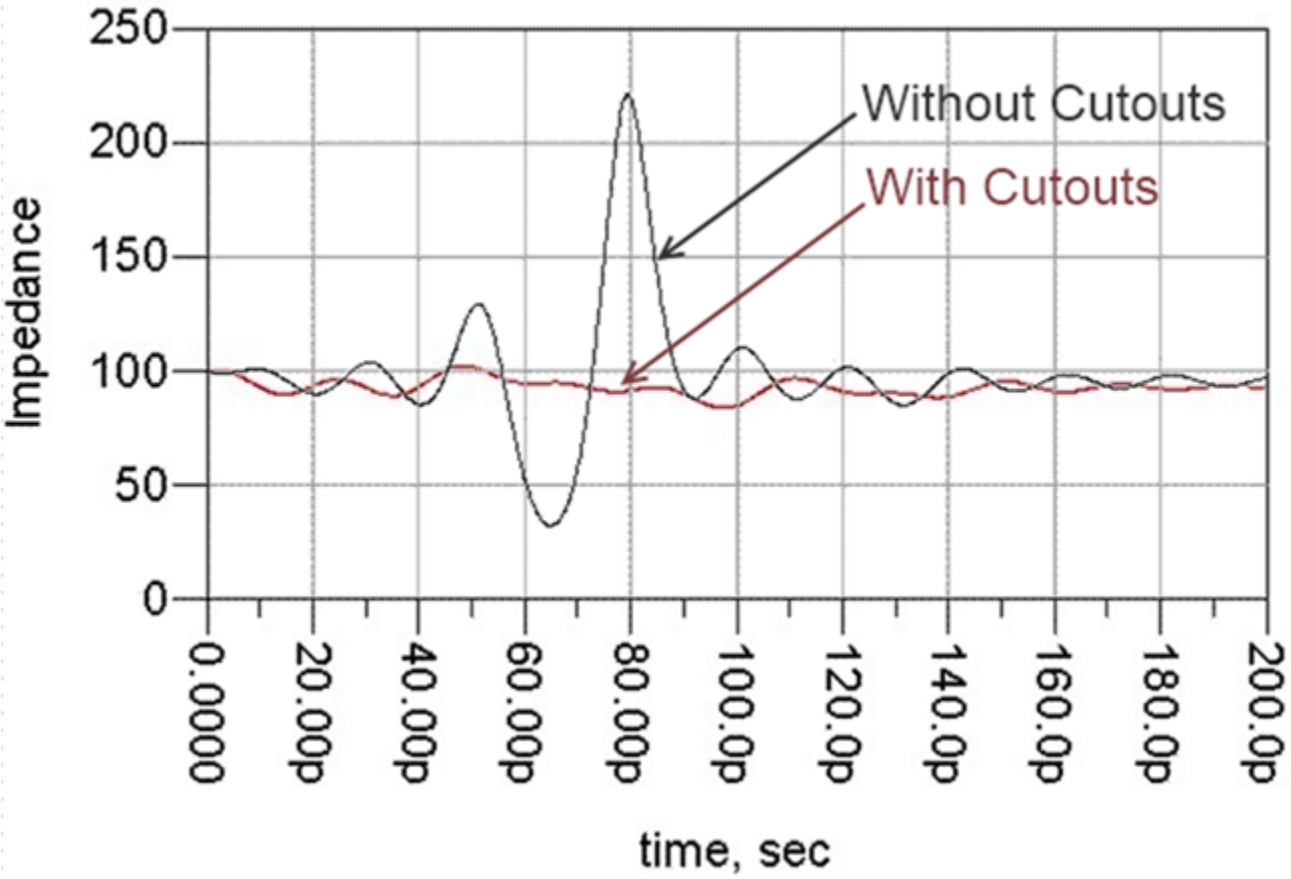

Related Information

Optimizing Impedance Discontinuity Caused by Surface Mount Pads for High-Speed Channel Designs

Connectors Optimization

Another component commonly found in the path of transceiver channels are backplane connectors and optical interface modules such as SFP+, XFP, and CFP2 among many others. For these components, the manufacturer may have specific design recommendations for the best connector performance, and the designer should follow those recommendations. If no recommendation is provided by the manufacturer, perform simulations to determine the best layout optimization. However, because most backplane and optical connector systems use press-fit through-hole via or surface mount pad technology respectively, the concept for optimization is very similar to that done for transparent via design and DC blocking capacitor compensation. The impact of impedance mismatch limits performance. For these cases, you can use the same methodology to optimize both the vias and surface mount pads in the connectors to better match the connector impedance to the 100-Ω trace impedance.

Figure 28: Example Backplane Connector Layout

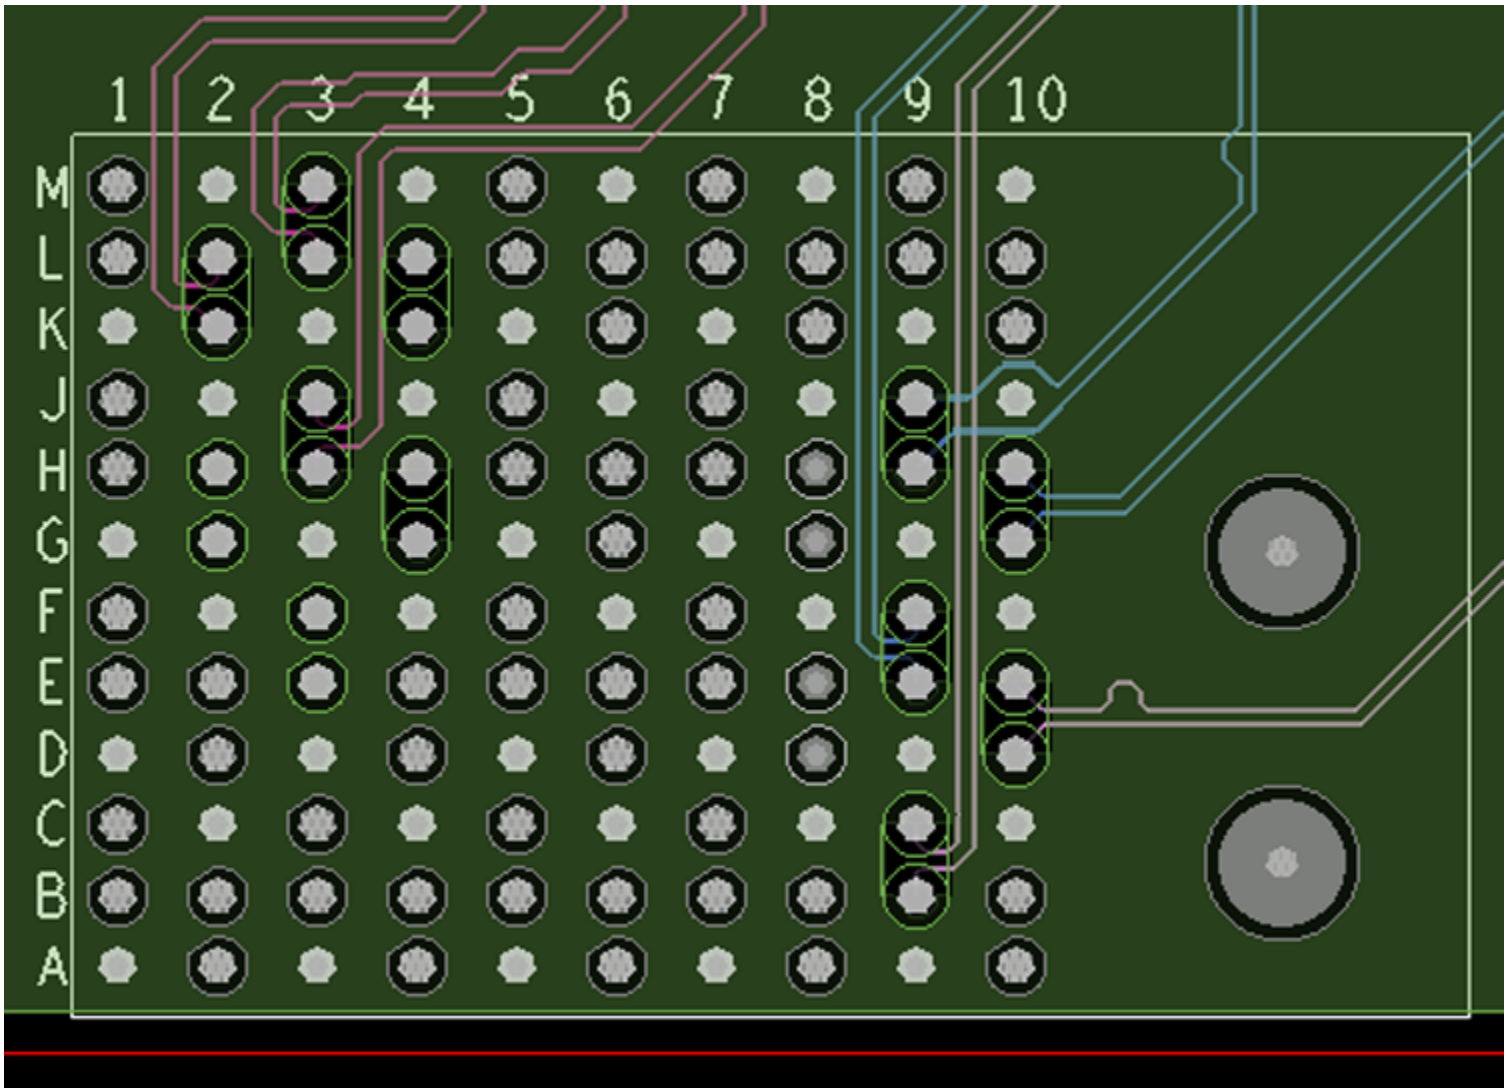

Figure 29: Example XFP Optical Module Connector Layout

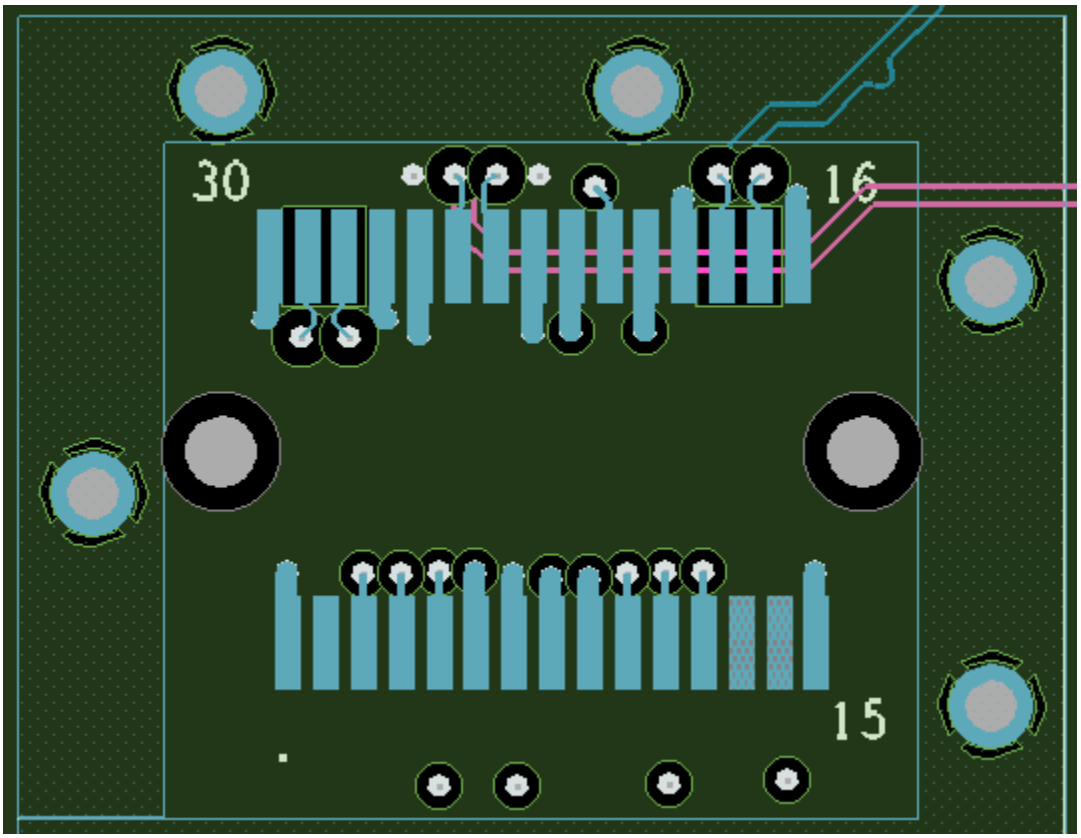

Notice that the via and surface mount pads show evidence of optimization such as increased via anti-pad size and plane cut-outs underneath the surface mount pads.

Note: Use manufacturer layout recommendations for connectors if available. In the absence of any specific manufacturer recommendations, designers can apply the transparent via optimization and DC blocking capacitor compensation technique to connector systems that use through-hole via and surface mount pad technologies, respectively.

Summary

A typical end-to-end transceiver link is comprised of three main components:

- PCB material
- Stack-up design
- Channel design

High-speed channel design involves understanding the first-order factors that impact performance in each one of those components. These first-order factors are:

- Signal attenuation
- Discontinuities and impedance control
- Crosstalk

Some recommended design choices to oppose those effects are summarized below.

PCB Material Selection

- Select the material with the lower loss tangent to reduce signal attenuation from dielectric losses.
- Always choose lower ε_r with a flat frequency response for best signal performance and to reduce signal dispersion that adds phase jitter.
- Always choose more densely woven fiberglass style for the core and prepreg material surrounding the high speed signal layers for more uniform ε_r that minimizes impedance and signal velocity variations.
- Choose sparsely dense fiberglass styles for power layers and slower general purpose signal layers to reduce PCB cost.
- Use wider traces and choose rolled copper foils over traditional electrodeposited (ED) copper foils in the PCB construction to mitigate conductor loss.
- Certain simulation tools may not include loss effects from surface roughness. In these cases, the correction factor (K_{SR}) must be added to get realistic prediction of the actual loss.

Stackup Design

For the same trace width and copper thickness considerations, stripline results in less signal attenuation compared with microstrip.

Channel Design

- For high-speed transceiver signals, use trace widths of 6mils or more to minimize conductor loss.
- Limit use of 4-mil trace widths to the BGA breakout area and keep their trace length as short as possible.
- Loosely coupled traces are easier to route and maintain impedance control but take up more routing area.
- Tightly coupled traces saves routing space but can be difficult to control impedance.
- Use stripline routing to avoid FEXT concerns.
- Use stripline traces with 5H differential pair-to-pair separation to minimize NEXT to 1%.
- If microstrip routing is required, used 6H-7H differential pair-to-pair separation to avoid NEXT and FEXT issues.
- C_{via} optimization techniques
  - Reduce the via capture pad size
  - Eliminate all non-functional pads (NFP)
  - Increase the via anti-pad size to 40 or 50 mils

- L_{via} optimization techniques:
  - Eliminate and / or reduce via stubs
  - Minimize via barrel length by routing near the stripline traces near the top surface layer and applying backdrilling
- Add ground return vias within 35 mils of each signal via to further improve the insertion and return losses of the via.
- Use manufacturer layout recommendations for connectors if available. In absence of any specific manufacturer recommendations, designers can apply the transparent via optimization and DC blocking capacitor compensation.

References

- Base Materials for High Speed, High Frequency PC Boards, Rick Hartley, PCB&A, Mar 2002
- High-Speed Digital Design, Handbook of Black Magic, Howard Johnson & Martin Graham, Professional Technical Reference, Upper Saddle River, NJ 07458, 1993
- Signal Integrity – Simplified, Dr. Eric Bogatin, Prentice Hall Professional Technical Reference, Upper Saddle River, NJ 07458, Dec 2006
- A Survey and Tutorial of Dielectric Materials Used In the Manufacture of Printed Circuit Boards, Lee W. Ritchey, Speeding Edge, Circuitree Magazine, Nov 1999
- High-Speed Digital Systems Design, A Handbook of Interconnect Theory and Design Practices. Stephen W. Hall, Garrett W. Hall, James A. McCall, New York, John Wiley & Sons, Inc., 2000
- Discussions on Non Functional Pad Removal / Backdrilling and PCB Reliability, Bill Birch, PWB Interconnect Solutions Inc., 103-235 Stafford Road West, Nepean, Ontario, Canada
- A Practical Method for Modeling PCB Transmission Lines with Conductor Surface Roughness and Wideband Dielectric Properties, Tao Liang, Stephen Hall, Howard Heck & Gary Brist
- Non-Classical Conductor Losses Due to copper Foil Roughness and Treatment, Gary Brist, Stephen Hall, Sidney Clouser, & Tao Liang
- CEI-25G-LR and CEI-28G-VSR Multi-Vendor Interoperability Testing White Paper, Optical Internetworking Forum, March 2012

Document Revision History for the AN 672: Transceiver Link Design Guidelines for High-Gbps Data Rate Transmission

Document Version	Changes
2020.01.29	Fixed the formatting and an error in the "Copper Roughness Approximation Equations" in the *Dielectric Constant* topic.
2017.05.08	Rebranded as Intel.
2017.02.02	• Changed the K_{SR} equation and added clarification to the R_a note in the *Copper Roughness Approximation Equations* figure. • Clarified the units of measure for WpN in the "DC Blocking Capacitor Compensation" figure.
2013.02.15	Initial release