AN 766: Intel® Stratix® 10 Devices, High Speed Signal Interface Layout Design Guideline

Contents

683132 | 2019.03.12

Send Feedback

Intel® Stratix® 10 Devices, High Speed Signal Interface Layout Design Guideline

This high speed signal interface design guideline helps you design best-in-class board layouts for high speed signals operating up to 28 Gbps.

These guidelines are based on the latest results of 3D board layout simulations and measurement.There are test coupons to validate this guideline methodology and evaluate the recommended layout guideline properly. Intel recommends that you read this guideline thoroughly and perform pre-layout and post-layout 3D simulations to confirm that your channel meets specifications.

Note: The content in this application note is based on currently available simulations and measurement data. It is subject to change pending new data. In the future revisions we will also focus on other high speed interfaces and backplane board layout design guidelines.

Intel® Stratix® 10 Devices and Transceiver Channels

Intel® Stratix® 10 devices vary by the number of supported transceivers channels. The figure below shows a magnified view of Intel Stratix 10 device F2397B package map.

Figure 1. Intel Stratix 10 F2397B Device Floor Plan with Magnified Transceiver Pins and an Application Example

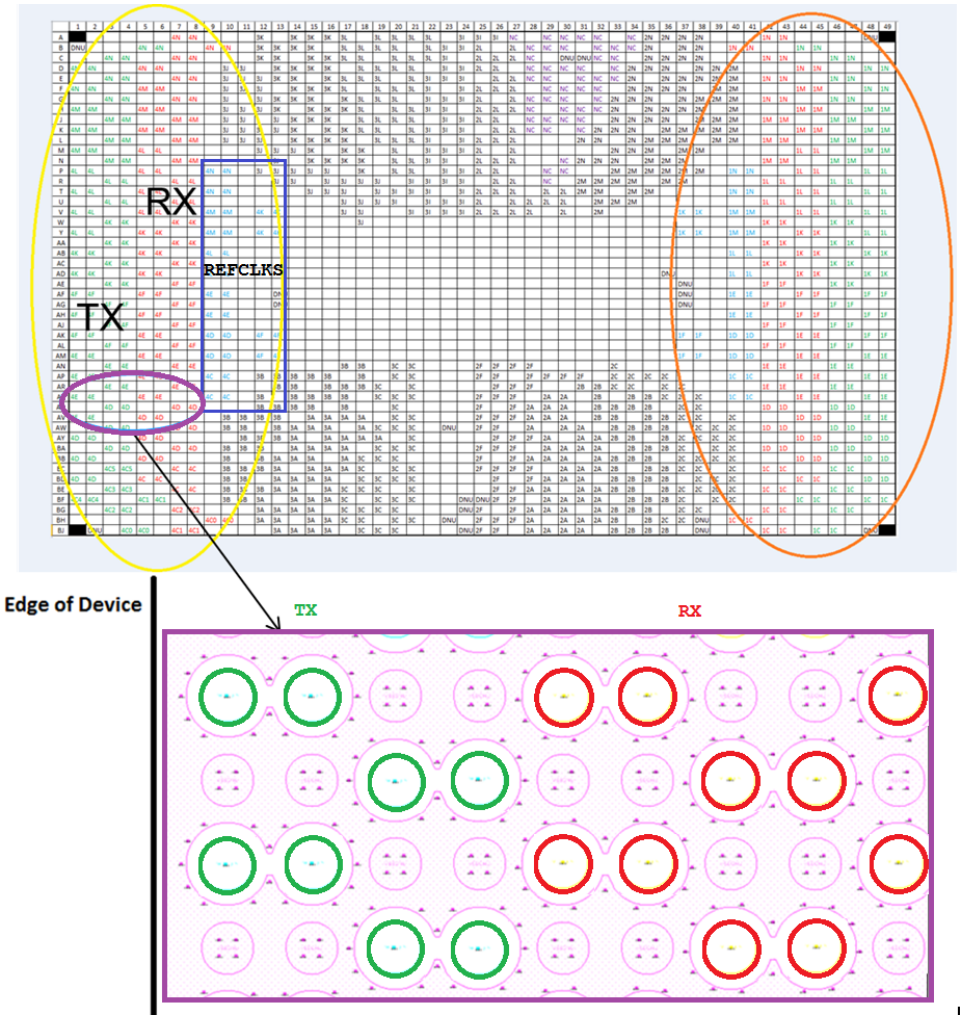

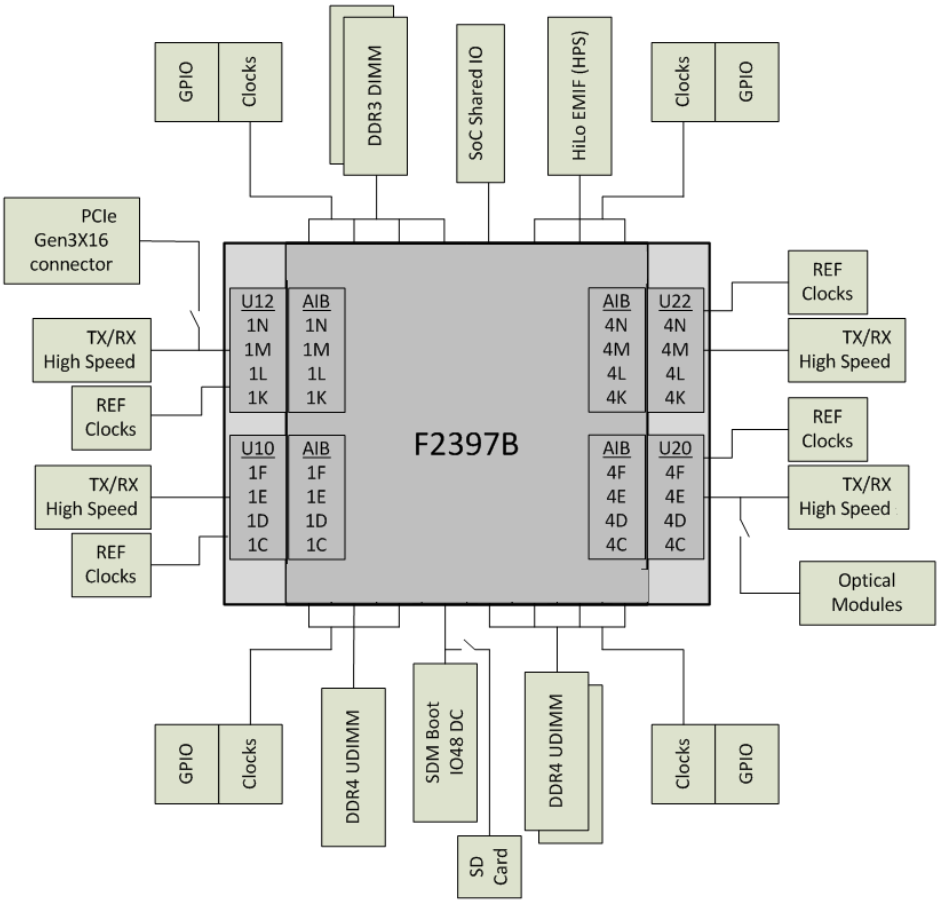

The transceivers pins are located on the edge of device. There are four transceiver dies inside the package that can support up to 96 transceiver channels for F2397B package. The total size for F2397B package (U50) is 50 mm by 50 mm. The BGA pitch is 1 mm.

RED pins are RX pairs and Green pins are TX pairs. (TX pairs are located on the edge of device, while RX pairs are located further into the device)

PCB Stackup Selection Guideline

For proper stackup selection for high speed signals in your PCB layout, follow these guidelines:

- Select a dielectric material with the lowest loss tangent and smaller dielectric constant, for example, the Megtron6 (df<0.002, epsr=3.1) is an appropriate choice.
  - When they become available after vendor characterization, dielectric materials such as Megtron 6N/6G or Tachyan 100G are good selections.
  - 25+G designs require special attention to material details including Fiberglass, Dielectric Matrix and Copper. The signal at higher data rate has higher frequency element and the wavelength goes on reducing. The change of fiber glass pattern, dielectric matrix pattern and copper pattern should be considered carefully. As for higher data rate (shorter signal wavelength), it appears to create more discontinuities and reflection with slight change. Please refer to *PCB Dielectric Material Selection and Fiber Weave Effect on High-Speed Channel Routing* for more information.
- Select smaller dielectric height for high speed signal routing.
  - It requires smaller trace width for trace impedance target. There is always a trade off between selecting wider trace width and shorter trace width. The wider width has less skin depth and lower insertion loss but takes more space for routing.
  - It also results in a smaller PCB height as well as smaller transition via height for achieving minimum impedance mismatching.
- Select enough stripline layers for all critical high speed signal routing.
  - Intel recommends stripline routing for all critical high speed signals (above 15 Gbps).
  - You can route all non-critical high speed signals (below 15 Gbps) on a microstrip layer.
  - Stripline routing has maximum isolation with other layers as long as both sides are reference planes. Intel does not recommend dual stripline routing unless the signal routing on both stripline layers are perpendicular. This means, longitudinal broadside coupling of differential pairs should be avoided.
  - Intel recommends Stripline preferred over microstrip. If microstrip routing is selected, Intel recommends removing the solder mask.
  - Stripline routing requires smaller trace width, which results in more space for signal routing.
- Selection of a ground/signal/ground stackup combination for critical high speed signals.
  - Selection of a ground/signal/ground combination may be feasible as long as the signal routing crossings on both stripline layers are perpendicular to minimize broadside coupling which results in cross-talk.
- Select enough power/GND layers to cover the power supply rails.

Estimating the insertion loss based on selected stackup material

Transmission line has various losses including the conductor loss, dielectric loss, surface roughness loss, skin depth loss etc. Table below shows various materials including their dielectric constant and loss tangents:

Table 1. Material Dielectric Constant and Loss Tangent

Material	εr	Tan(δ)
Typical FR4	4	0.02
GETEK	3.9	0.01
Isola 370HR	4.17	0.016
Isola FR406	4.29	0.014
Isola FR408	3.70	0.011
Megtron 6	3.4	0.002
Nelco 4000-6	4.12	0.012
Nelco 4000-13 EP	3.7	0.009
Nelco 4000-13 EP SI	3.2	0.008
Rogers 4350B	3.48	0.0037

The loss tangent mentioned in above table has been typically measured at 1 GHz based on the material data sheets.

Note: Intel recommends to refer to the manufactures latest data sheets

The average approximate PCB attenuation of only transmission at frequency f is based on below formula.

$$\text{atten}[\text{dB/in}] \sim \frac{1}{w[\text{mils}]}\sqrt{f[\text{GHz}]} + 2.3 \times f[\text{GHz}] \times \text{Df} \times \sqrt{\text{Dk}}$$

Equation (1)

Where:

W=the trace width in mil

f=the sine wave frequency in GHz, equivalent to Nyquist for specified data rate

Df=the dissipation factor (same as loss tangent)

DK=the dielectric constant

The formula above is divided into two parts: the first part is trace loss (including skin depth) and the second part is dielectric loss.

The graph in the figure *PCB Trace Attenuation Comparison per 1" trace length for various dielectric materials, while trace width is 5 mil, results up to 20 GHz* shows the average trace loss per inch for various materials in above table. This graph has been extracted based the assumption that W=5 mil.

Figure 2. PCB Trace Attenuation Comparison per 1 inch Trace Length for various dielectric materials, while Trace Width is 5 mil, results up to 20 GHz.

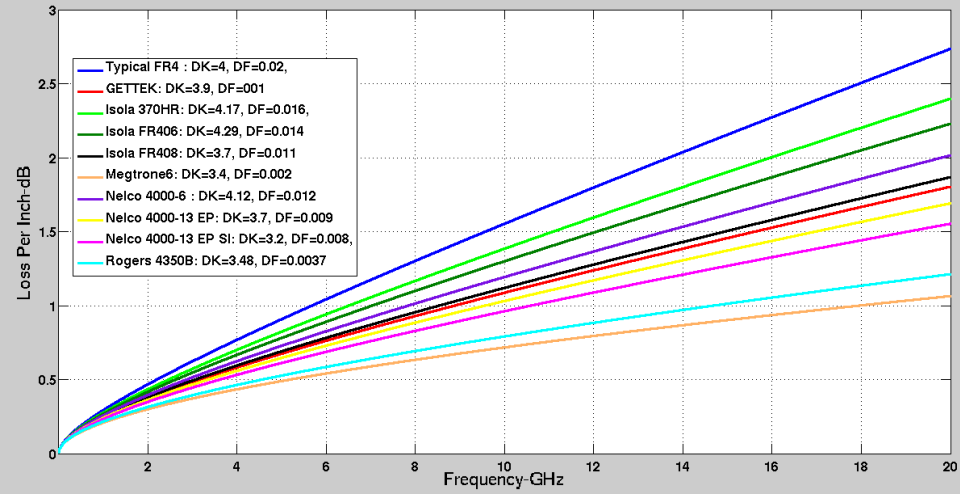

From the above figure, Megtron 6 has 0.85 dB average loss per inch at 28 Gbps (Nyquist is 14 GHz). On the other hand, Typical FR4 has approximately 2 dB loss at the same frequency.

Copper thickness has not been encounter into the above approximate PCB attenuation equation. The thicker copper width, the less trace resistance.

Intel recommends that the designers must consider an average of +/-5% variation into the loss obtained in figure *PCB Trace Attenuation Comparison per 1" trace length for various dielectric materials, while trace width is 5 mil, results up to 20GHz* due to some material tolerances by Fabrication Company.

An average surface roughness (approximately 2 µm) has been included into the approximate PCB attenuation equation for trace loss attenuation. For accurate loss calculation, Intel recommends the designers to have at least 2.5D CAD analysis on transmission loss attenuation considering actual surface roughness, copper thickness and frequency dependent dielectric materials.

Table 2. Average loss for trace per inch at 14 GHz for Megtron 4 vs Megtron 6 vs Tachyon100G

Material	MEG4	MEG6	Tachyon100G
Average Loss per inch @14 GHz	1.2 dB	0.85 dB	0.8 dB

Overall, MEG6 and Tachyon100G materials are the best options for 28 Gbps high speed signals routing.

For more information on the various weave compositions and material dielectric loss considerations and their influence on the channel performance, refer to *PCB Stackup Design Considerations for Altera FPGAs* and *PCB Dielectric Material Selection and Fiber Weave Effect on High-SpeedChannel Routing*.

Related Information

- PCB Dielectric Material Selection and Fiber Weave Effect on High-Speed Channel Routing

- PCB Stackup Design Considerations for Altera FPGA's

Recommendations for High Speed Signal PCB Routing

To achieve better performance for high speed channels, follow these guidelines:

- TX and RX signal routing must be isolated using separate stripline layers for critical high speed interfaces above 15 Gbps.
- Intel recommends that the RX signal routing layer be located above the respective TX signal routing layer. This means that the RX routing layer must be separated from TX routing layer. When FPGA's are located on top layer and all high speed vias are back drilled from bottom, Intel recommends RX layers on upper layers and TX layers on those layers below RX layers. The scenario will be opposite if FPGA's are located on bottom layer. In that case, RX layers are recommended on bottom layers and TX layers are recommended on those layers above RX layers.
  - RX signals are always weaker than TX signals. Obtaining shorter transition via length for RX signals reduces mismatching and reflection, and more RX signal power is received at the device.
  - Most of High Speed Interfaces require AC coupling caps on RX signal lanes. Intel recommends RX routing on upper layers close enough to top layer. By this, designer can achieve shorter signal via transition height and eventually reduce reflection on RX path.
  - AC caps can also be mounted on the bottom layer of the PCB. Intel recommends this for signal data rates below 15 Gbps. In this case, you can select the RX signal routing layer as a stripline as close as to bottom layer as possible.
- Ensure that you have a good contiguous and un-interupted ground reference plane for high speed signal routing.
  - Avoid using the power plane as reference plane for critical high speed signal routing above 15 Gbps.
  - Void regions along, underneath and above the high speed signal routing is prohibited.
  - Always maintain enough space from the edge of the signal trace to the edge of the void region to avoid mismatching due to lack of sufficient reference plane. Spacing should at least be the signal trace width. Intel recommends to have this space be equal or larger than trace width in break-out region.
  - To avoid cross talk between adjacent pairs, make sure to have enough space between pairs routed on the same layer. The rule is to keep at least 3x (the height of signal to reference plane or the width of signal trace, whichever is larger) for the space between adjacent pairs.
- You must back drill for all high speed signal transition vias.
  - Back drill can be applied from either top or bottom layer, depending on the FPGA and the connectors that have been mounted on top or bottom layers. This emphasizes which part of via as extra stub is needed to be removed.
- Remove all non-functional via pads for both signal and GND return path vias.

FPGA Fan-out Region Design

Signal Break-out Recommendations

Figure Figure 3 on page 10 shows that there are up to 4 transceiver pairs in one row for Intel Stratix 10 devices. Designers must assign at least 4 signals layers each separated by continuous ground planes for break-out routing. The number of layers assigned for transceivers routing is one of the key factors for stackup selection. Intel recommends that you follow FPGA break-out region guideline in the next section to achieve optimum performance.

Figure 3. Example of Recommended FPGA Break-Out Routing (Different colors stands for different layers)

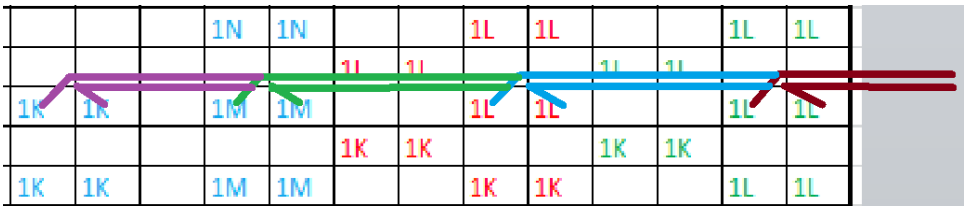

There are three options for FPGA fan-out region routing. Each one can be selected for high speed signal routing on PCB, based on their high priority and data rate.

Note: Intel recommends BGA pad diameter on PCB as 20 mil for Intel Stratix 10 devices. This is good for pads that are not near the corners of the device. However, for those five pads which are located closest to each corner of device, the recommended pad diameter is 24 mil with 20 mil solder mask opening for inner pins and 16 mil for corner pins.

Option 1: Via-In-Pad Topology

Intel recommends using via-in-pad technology for data rates above 15 Gbps. Use via-in-pad when you want to transfer signals from the BGA pad to the inner layers.

Figure 4. Option 1: FPGA Fan-out Configuration at Solder Ball Topology

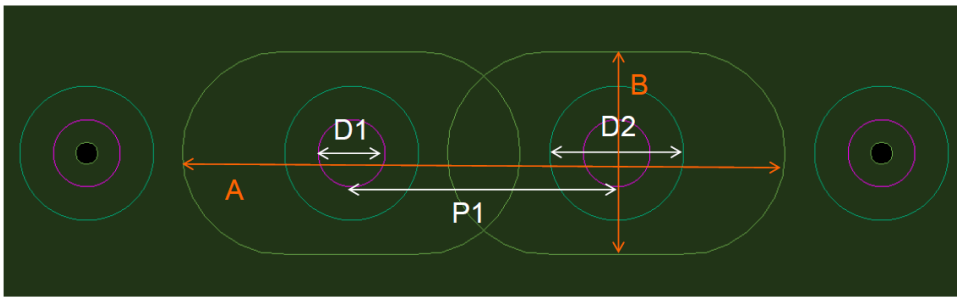

D1: Via drill hole diameter: 8 mil (for up to 1:12 stackup aspect ratio) or 10 mil for higher stackup aspect ratios.

D2: Via pad diameter: 18 mil (for up to 1:12 stackup aspect ratio) or 20 mil for higher stackup ratios.

P1: Standard via-to-via pitch: 1 mm.

A: Horizontal anti-pad: 90 mil.

B: Vertical anti-pad: 28 mil (for up to 1:12 aspect stackup ratio) or 30 mil for higher stackup aspect ratios.

Option 2: Dog-bone with GND Cutout at BGA Pad Topology

Intel strongly recommends this option for data rates above 15 Gbps if you do not want to use via-in-pad on the PCB. Use the dog-bone configuration when you use FPGA fan-out.

Figure 5. Option 2: Dog-bone with GND Cutout at BGA Pad Topology

- Use of dog-bone configuration in FPGA-fan-out.
- Circular diameter ground cut-out underneath of FPGA BGA pad on PCB used for device is 22 mil
  - Intel recommends BGA pad diameter on PCB for Intel Stratix 10 devices should be 20 mil
- D1: Via drill hole diameter: 8 mil (for stackup aspect ratios up to 1:12) or 10 mil for higher stackup aspect ratios.
- D2: Via pad diameter: 18 mil (for stackup aspect ratios up to 1:12) or 20 mil for higher stackup ratios.
- P1: Standard via-to-via pitch: 1 mm.
- A: Horizontal anti-pad: 90 mil
- B: Vertical anti-pad: 28 mil (for stackup aspect ratios up to 1:12) or 30 mil for higher stackup aspect ratios.
- PCB BGA pad to signal transition via pad trace length (center to center): 26 mil (for stackup aspect ratios up to 1:12) or 27 mil for higher stackup aspect ratios.
- Use of 47.5 Ω single-ended trace connecting the BGA pad to via pad. Since the GND reference plane underneath of this trace is already cutout, designers might need to go with maximum trace width possible to achieve 47.5 Ω single ended impedance (~20 mil trace width). This 47.5 Ω single ended impedance design is due to match with the targeted 95 Ω differential impedance characteristics design as recommendation for high speed signals routing on PCB. Refer to *Option 3: Micro-via Topology*.

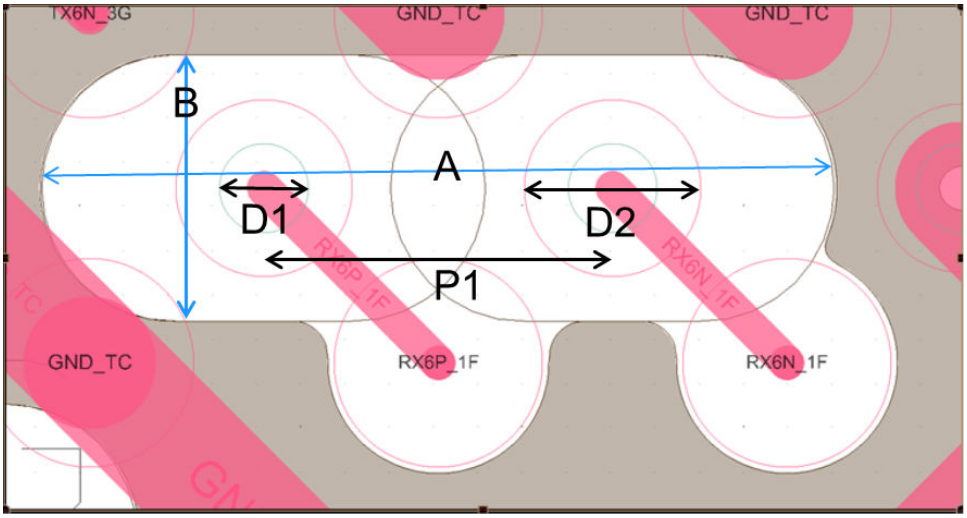

Related Information

Option 3: Micro-via Topology on page 12

Option 3: Micro-via Topology

Intel recommends this topology if you use a micro-via technology.

Figure 6. FPGA Fan-out Configuration at Solder Ball for Each Single-ended Lane

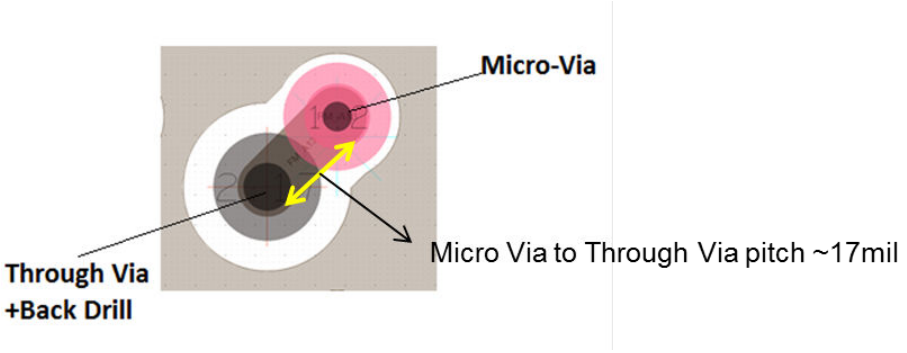

Topology specifications:

- Use of "Micro-Via or laser drilled-via" in combination with "Via-in-pad". Micro-via on FPGA Pad transfers the signal from top layer to the signal pad on the first GND reference layer underneath of top layer. Through via is then used to transit the signal to other layers.
- Micro-via dimensions:
  - Via hold/drill diameter: 5 mil
  - Via pad diameter: 10 mil
  - Via anti-pad diameter: 22 mil
- Through-via dimensions:
  - Via hold/drill diameter: 10 mil
  - Via pad diameter: 20 mil
  - Via anti-pad diameter: 30 mil
- Use of 47.5 Ω single ended trace impedance connecting the micro via pad to Through-via pad on GND reference plane. This 47.5 Ω single ended impedance design is due to match with the targeted 95 Ω differential impedance characteristics design as recommendation for high speed signals routing on PCB. Refer to *GND Cutout Under BGA Pads in Fan-out Configuration*.

Related Information

GND Cutout Under BGA Pads in Fan-out Configuration on page 13

GND Cutout Under BGA Pads in Fan-out Configuration

If you use the dog-bone fan-out configuration, Intel recommends that you have one GND reference plane cutout under the BGA pad to reduce capacitance in this area. Larger GND cutouts provide better impedance transition at the fan-out; however, the GND cutout in the fan-out is limited due to limited routing space.

Figure 7. Conventional Dog-bone and Recommended Dog-bone Comparison

Original design without a GND cutout under the BGA pad and with cutout.

Note: The cutout is 26 mil in diameter as socket is used for this device. If socket is not used, Intel recommends 22 mil.

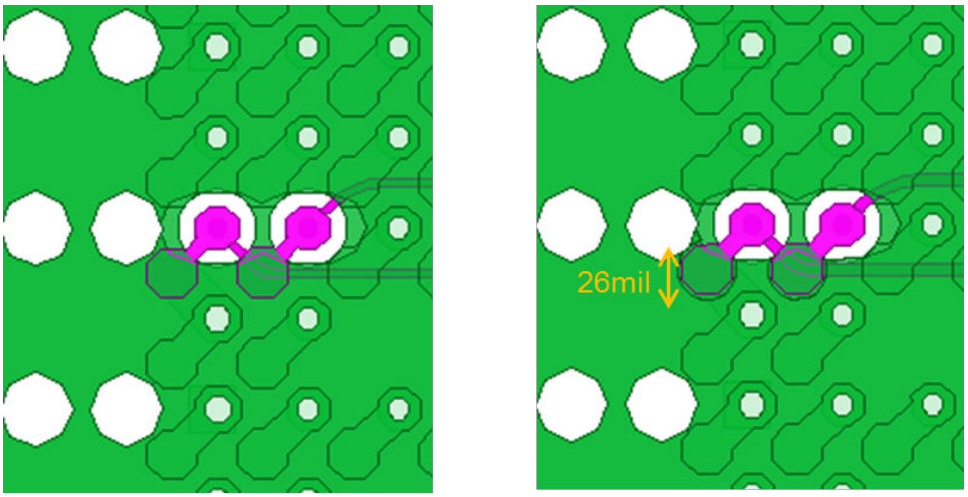

Figure 8. TDR Simulated Performances of the Entire Channel with Magnified FPGA BGA Area

Design uses 100 Ω and 90 Ω differential impedance routing.

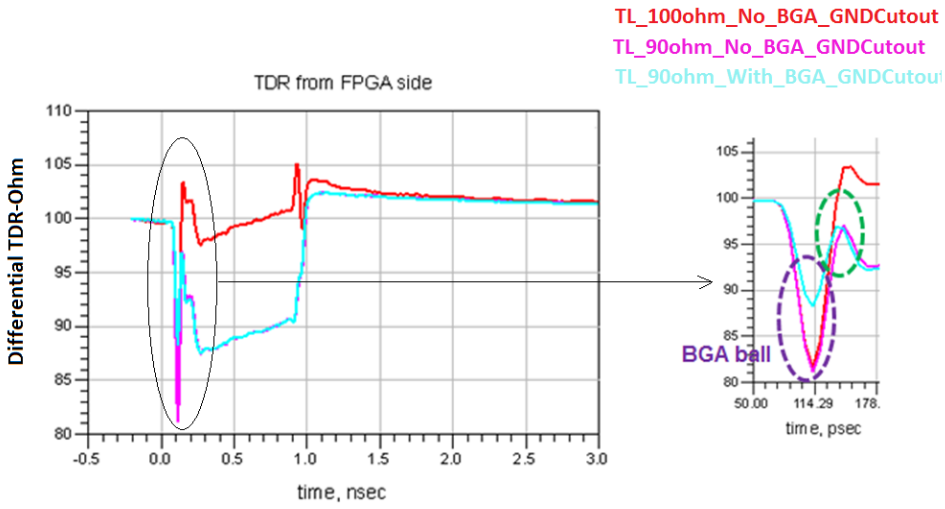

TL impedance impact effect with dog-bone fan-out configuration and with/without GND cutout underneath of BGA pad is approximately 10 Ω based on TDR performance in Figure 8 on page 14.

Figure 9. Simulated Performances of Insertion and Return Loss

Full channel simulation uses 100 Ω and 90 Ω differential impedance routing.

The red line indicates the original design and 100 Ω TL routing impedance.

The pink line indicates the original design and 90 Ω TL routing impedance.

The blue line indicates the original design with 26 mil diameter GND cutout under the BGA pad and 90 Ω TL routing impedance.

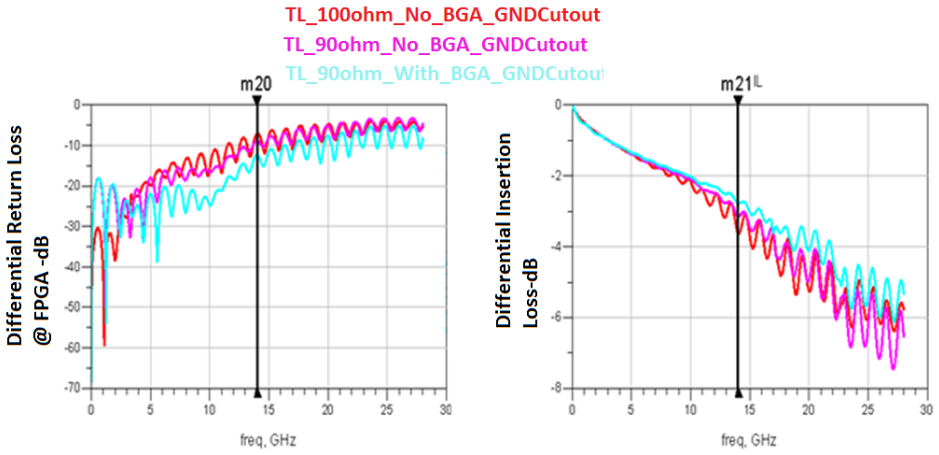

The improvement on insertion loss is 1 dB and approximately 5 dB on return loss at 14 GHz.

Comparison of Dog-bone with GND Cutout Under the BGA and Via-in-Pad Configurations

Figure 10. Comparison of TDR from the FPGA Pad of Via-in-Pad and Dog-bone Configurations

The pink line indicates the dog-bone configuration and the green line indicates the via-in-pad configuration.

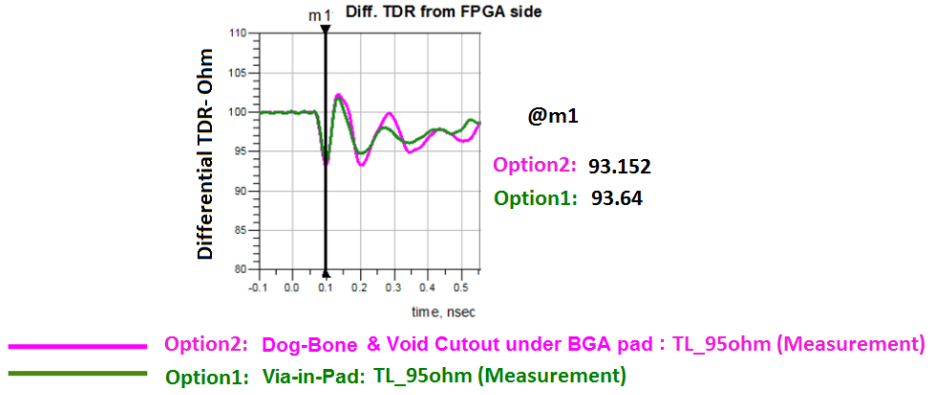

In terms of TDR, both configurations have similar performance.

Figure 11. Full channel scattering parameter performances for Dog-Bone Configuration with GND/Void cutout underneath of BGA (Option 2) vs Via-in-pad (Option 1) Configurations

Figure shows the identical channel but uses the via-in-pad or dog-bone configuration in the FPGA area.

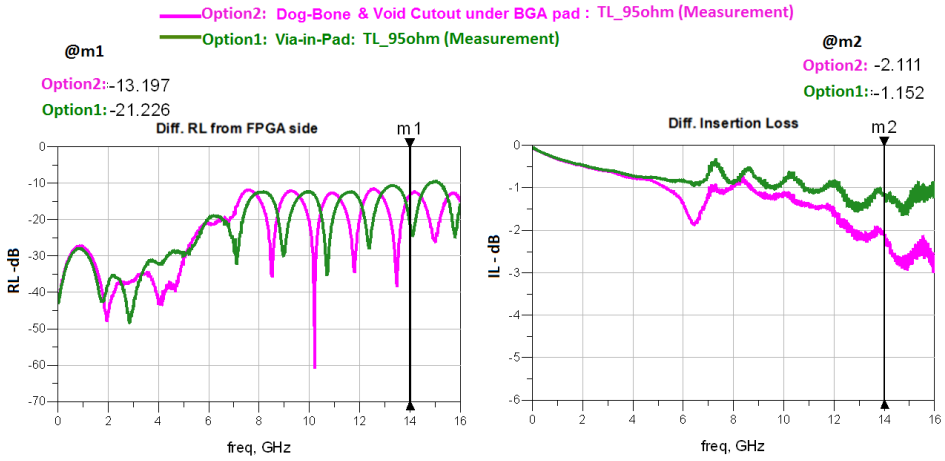

Note:

- Option 1 (Green): Via-in-pad.
- Option 2 (Pink): Dog-bone with GND/Void cutout under BGA pad configuration.

In this example, the via-in-pad shows slightly better performance. The other advantage of via-in-pad is that it provides more space for fan-out routing.

Figure 12. TDR Performance and TL Impedance Impact

The example uses a dog-bone fan-out configuration and one layer GND cutout under the BGA pad by various TL routing impedances.

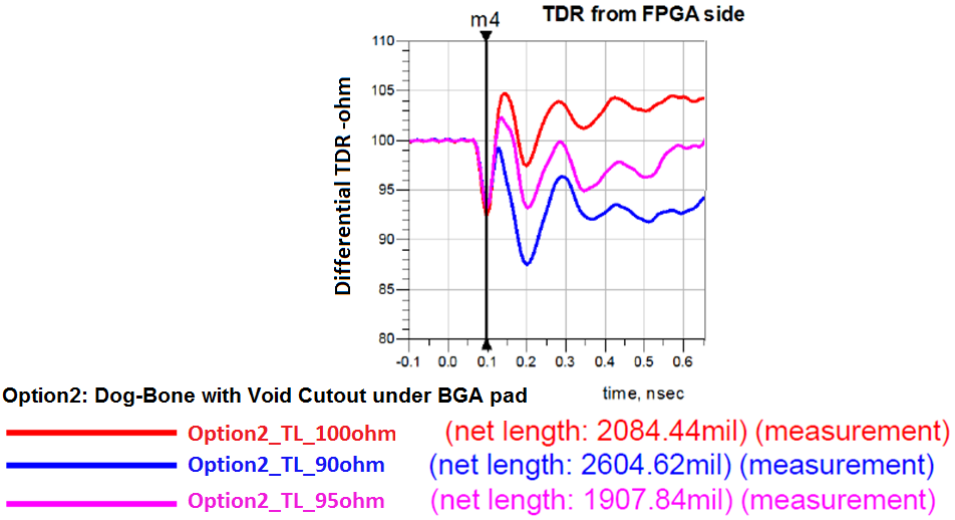

This example shows that slightly lowering the TL routing impedance ensures a smoother transition from the BGA pad, which results in less reflection on the channel. Intel recommends a 95Ω TL routing impedance for high speed serial interface (HSSI) channels.

Trace Shape Routing at the BGA Void Area (Tear Drop Configuration)

Impedance matching from the signal via pad to signal trace is an essential element for high speed interfaces. To avoid substantial discontinuity, Intel recommends using the following trace configuration for better transition (see the area highlighted in red).

Figure 13. Proposed tear drop configurations by PCB fabrications

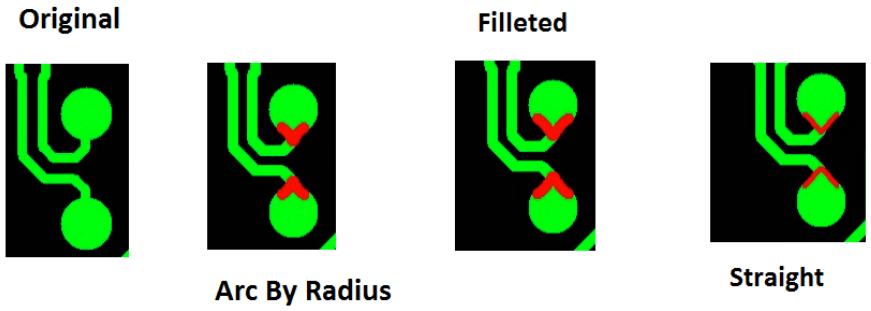

FPGA Fan-out Region Routing Recommendations

Intel recommends that you use a 95 Ω differential trace lane for both PCB and fan-out region routing. The tolerance of trace impedance can be within ±10% on the PCB.

The back-jog routing configuration is best for high speed signals. The advantage of the back-jog configuration is that it keeps the skew matching for both differential lanes in the break out region. Back-jog routing can be either single-ended back-jog or neck-down with back-jog in the fan-out region.

Figure 14. Recommended Single-ended Back-jog and Differential Neck-down with Back-jog Break-out Routing

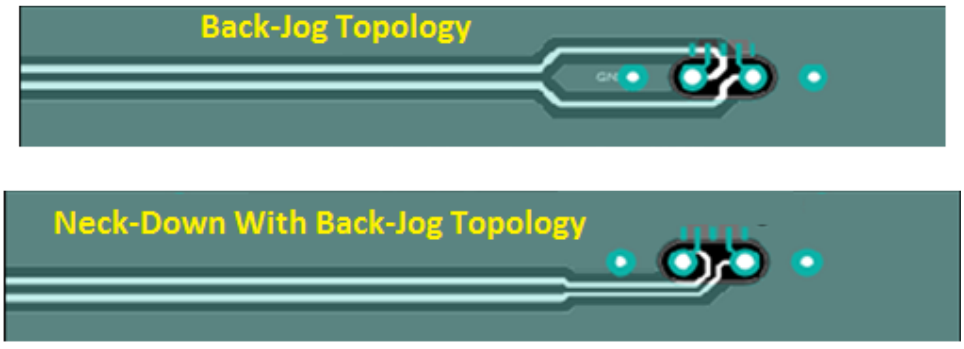

Simulations and measurements show that a traditional jog-out routing configuration in the break-out region degrades performance, specifically for return loss. The traditional jog-out routing requires skew-matching right at the device edge.

Figure 15. Conventional In-line Break-out and Differential Neck-down with Jog-out Break-out Routing

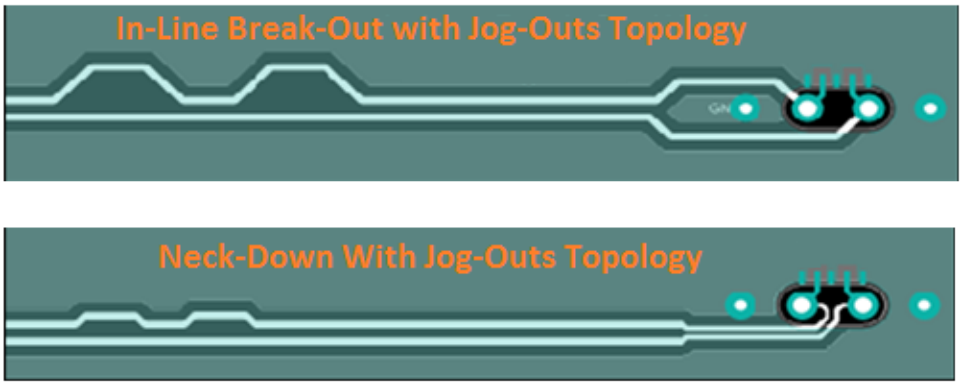

For signal data rates above 15 Gbps, Intel recommends that you use a back-jog break-out routing configuration. For data rates lower than 15 Gbps, you must use the jog-out break-out routing with skew matching at the device edge.

For both of these options, single ended routing in the fan-out region results in better insertion loss performance.

Observe these guidelines when routing in the fan-out region:

- Route one signal pair between 1 mm BGA pitch for maximum isolation between pairs.
- Keep the fan-out routing length less than 1 inch.
- Avoid routing lanes which are close to the edge of the void area. Maintain enough space between the trace edge and the void edge. A wide reference plane is always required for any high speed signal routing.

Figure 16. Example of Routing Away from the Void Area

The a and b space dimensions must be no smaller than the size of the trace width.

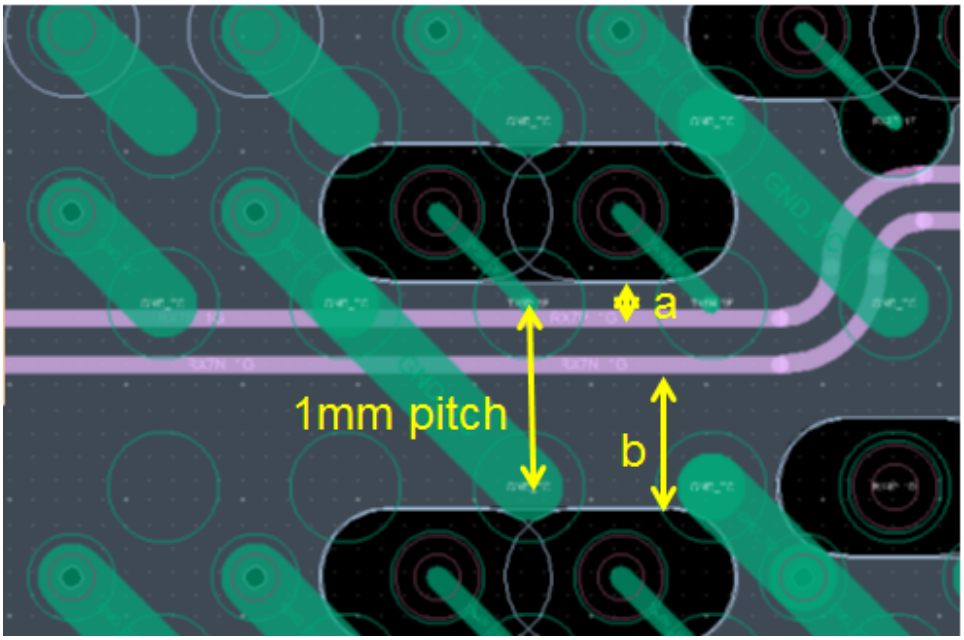

Comparison of Conventional and Recommended Break-out Routing Topologies

This evaluation uses a 24-layer stack-up.

Figure 17. Sample Portion of a 24-Layer Stackup for Break-out Routing Examples

Total stack-up height is 117 mil.

Layer		Cu Weight	Proposed Thickness (mils)	Construction	Assume copper density	Dk @ 10GHz	Df @ 10GHz	Ref.	OPCM Stack-up Information											
									Single End 35ohm +/- 5ohm		Single End 40ohm +/- 5ohm		Single End 50ohm +/- 10%		Differential 75ohm +/- 10%		Differential 85ohm +/- 10%		Differential 100ohm +/- 10% outer layers +/-7.5% inner layers	
									Target LW	Finished LW	Target LW	Finished LW	Target LW	Finished LW	Target LW/SP	Finished LW/SP	Target LW/SP	Finished LW/SP	Target LW/SP	Finished LW/SP
			0.70																	
L1	Top	1/3oz + Plate	2.20	HTE	50%			L2		9.2		7.4		5.0		5.9/4.0		4.5/4.0		3.5/5.0
	prepreg		2.86	R5670K 1078PC64		3.38	0.004													
L2	GND	0.5 oz	0.98	HVLP	75%															
	Core		4.00	R5775K 1028*2		3.37	0.004													
L3	SIG1	0.5 oz	0.98	HVLP	25%			L2 &4		7.8		6.3		4.2		5.9/3.5		4.0/4.0		3.5/4.5
	prepreg		4.00	R5670K 1035PC70 x 2		3.23	0.004													
L4	GND	0.5 oz	0.98	HVLP	75%															
	Core		4.00	R5775K 1028*2		3.37	0.004													
L5	SIG2	0.5 oz	0.98	HVLP	25%			L4 &6		7.8		6.3		4.2		5.9/3.5		4.0/4.0		3.5/4.5
	prepreg		4.00	R5670K 1035PC70 x 2		3.23	0.004													
L6	GND	0.5 oz	0.98	HVLP	75%															
	Core		4.00	R5775K 1028*2		3.37	0.004													
L7	SIG3	0.5 oz	0.98	HVLP	25%			L6 &8		7.8		6.3		4.2		5.9/3.5		4.0/4.0		3.5/4.5
	prepreg		4.00	R5670K 1035PC70 x 2		3.23	0.004													
L8	GND	0.5 oz	0.98	HVLP	75%															
	Core		4.00	R5775K 1028*2		3.37	0.004													
L9	SIG4	0.5 oz	0.98	HVLP	25%			L8 & 10		7.8		6.3		4.2		5.9/3.5		4.0/4.0		3.5/4.5
	prepreg		4.20	R5670K 1035PC70 + 1080P264		3.27	0.004													
L10	PWR1	2 oz	2.40	RTF	75%															
	Core		2.00	R5775K 1020*1		3.37	0.004													
L11	GND	2 oz	2.40	RTF	75%															
	prepreg		4.00	R5670K 1080PC64 x 2		3.38	0.004													
L12	PWR2	2 oz	2.40	RTF	75%															
	Core		4.00	R5775K 1028*2		3.37	0.004													

Conventional Jog-out Routings with Skew-matching Right After Break-out

This routing design is configured with the following characteristics:

- The stack-up is 24 layers with a thickness of 117 mil
- Eight signal layers
- Four PWR layers

- Material is Megtron6
- Via-in-pad topology with 8 mil finished drill
- 18-mil signal pad and 28 mil signal anti-pad
- Horizontal anti-pad is 68 mil (40 mil pitch + 28 mil anti-pad)
- Vertical anti-pad is 28 mil

Figure 18. Case1: Conventional Differential Routing with Neck-down and Jog-out

The transceiver pair of A and B have been routed on the layer 5 of stackup. Transceiver pair of C and D on the same row have been routed on a different layer. Only two signal layers are required for four transceiver pairs by using conventional differential routing with neck-down and jog-out routing.

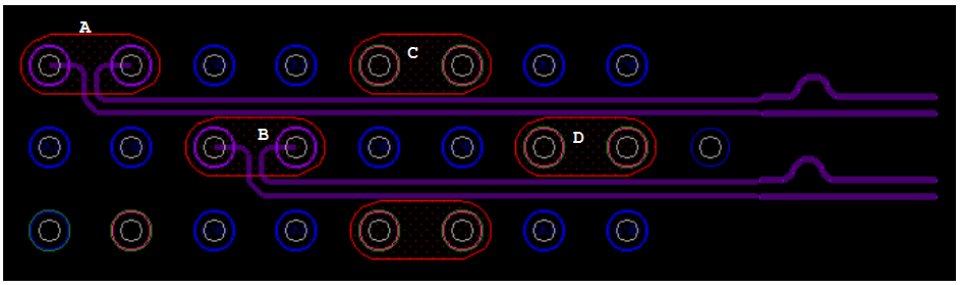

Figure 19. Case2: Conventional Single-ended In-line Breakout Routing with Jog-out

Yellow routing shows transceiver pair A routed on layer seven. Red routing shows transceiver pair B routed on layer nine. Due to single-ended break-out and lack of space, the transceiver pairs C and D on the same row have been routed on different layers. Four signal layers are required for four transceiver pairs by using conventional single-ended in-line breakout routing with jog-out.

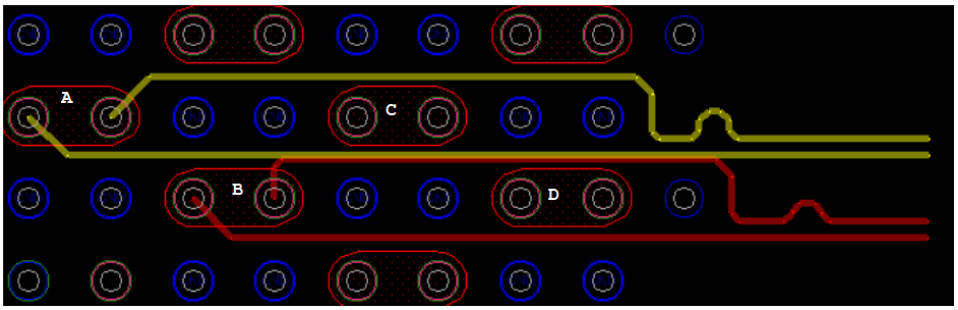

Because fabrication always has some layer-to-layer mismatch, this example implements a typical 5 mil layer-to-layer mismatch to the cases above. This allows you to observe the level of sensitivity to layer-to-layer fabrication mismatch. This layer-to-layer fabrication mismatch moves the routing passing by the GND void area and adds more discontinuity to the routing path.

Figure 20. Case 3: 5 mil Layer-to-layer Fabrication Mismatch on Conventional Differential Routing with Neck-down and Jog-out

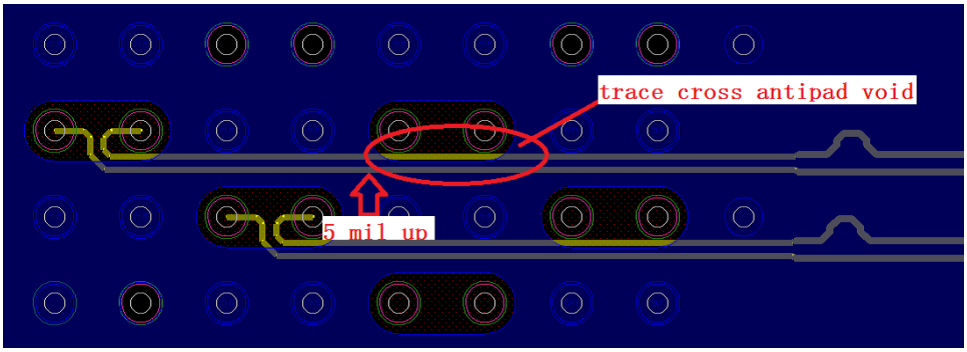

Figure 21. Case 4: 5 mil Layer-to-layer Fabrication Mismatch on Conventional Single-ended In-line Breakout Routing and Jog-out

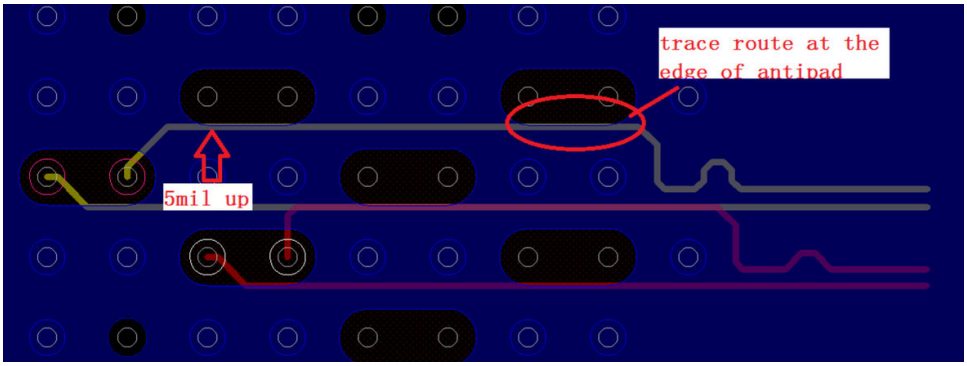

Performance Comparison Between Case 1 and Case 3

This section compares the performances between conventional differential routing with neck-down and jog-out with/without mismatch (Case 1 vs Case 3).

Figure 22. Simulated Differential Insertion Loss

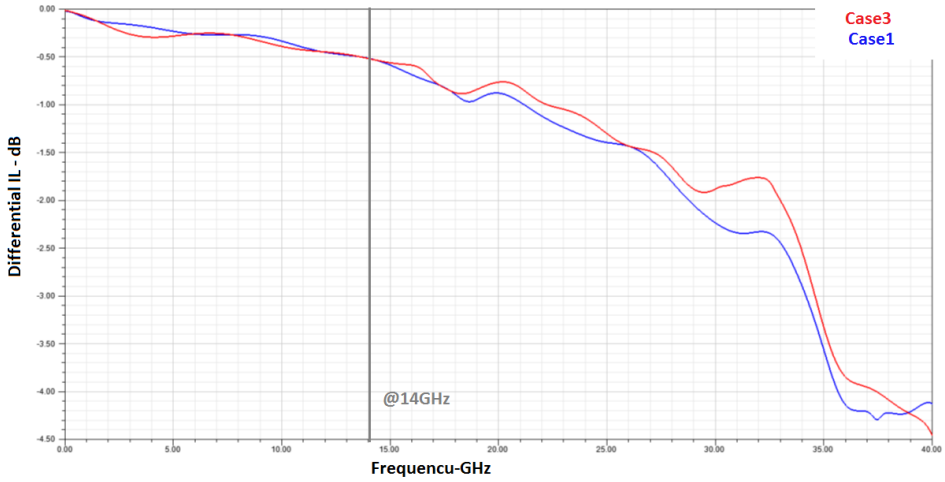

Figure 23. Simulated Differential Return Loss at the BGA Solder Ball

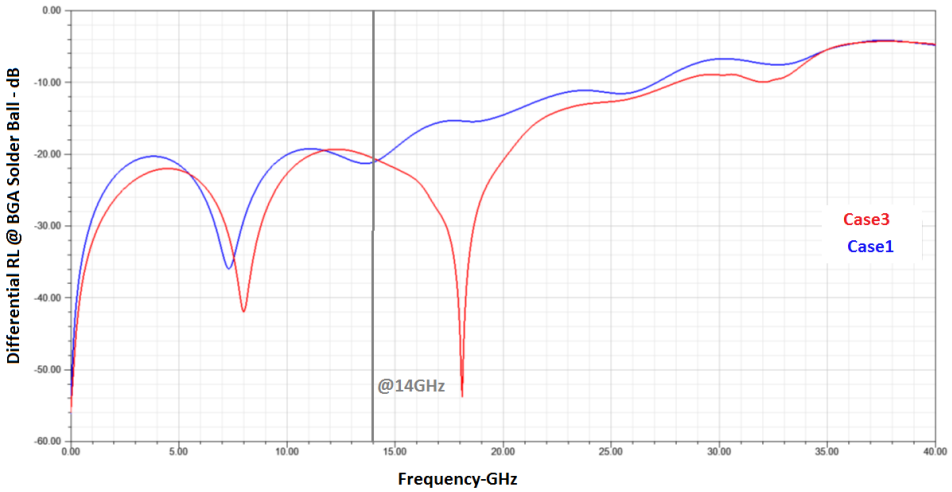

Figure 24. Simulated TDR Differential Impedance from the Trace End

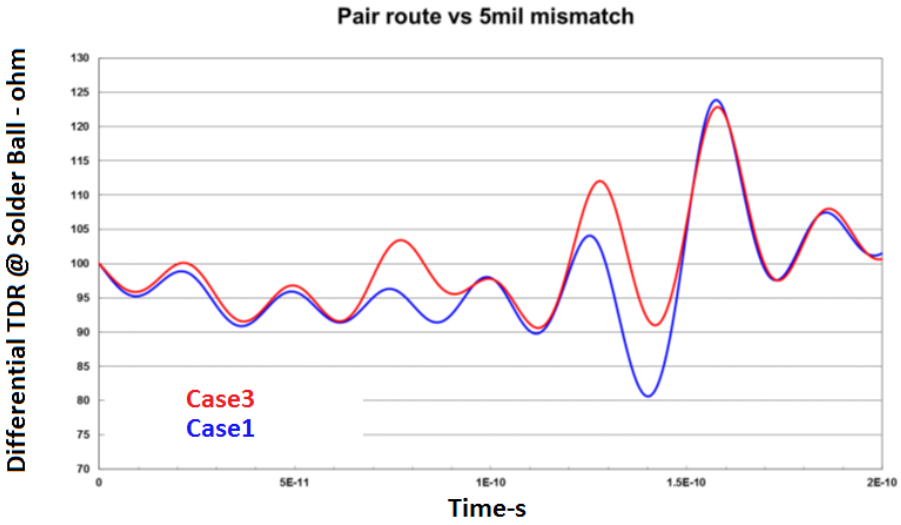

These performance results demonstrate that within 15 GHz bandwidth Case1 is robust enough to accommodate the layer-to-layer fabrication mismatch. The TDR impedance shows up to 7 Ω impedance mismatch due to the 5 mil layer-to-layer mismatch.

Performance Comparison Between Case 2 and Case 4

This section compares the performances between conventional single-ended in-line routing with jog-out with/without mismatch (Case 2 vs Case 4).

Figure 25. Simulated Differential Insertion Loss

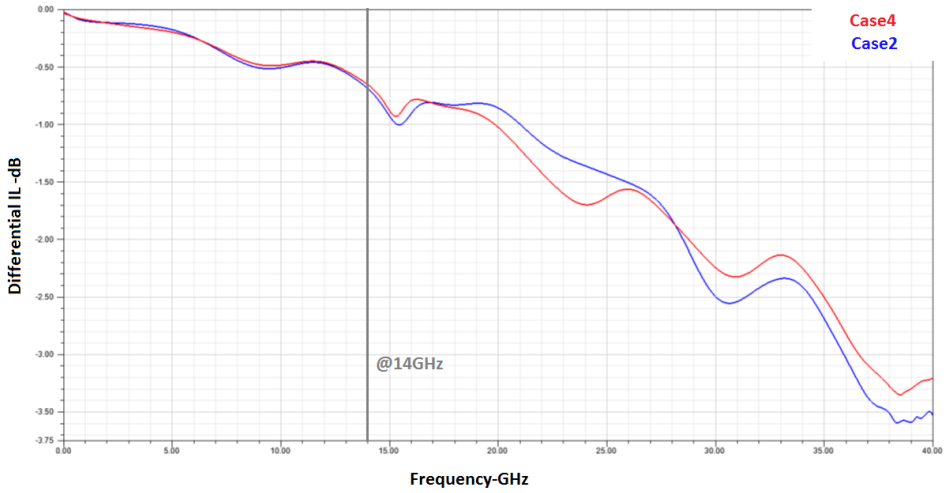

Figure 26. Simulated Differential Return Loss at BGA Solder Ball

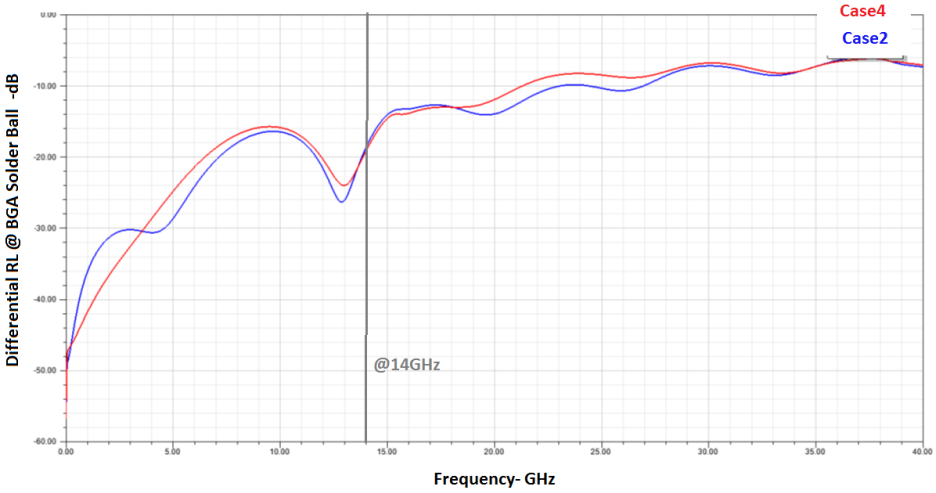

Figure 27. Simulated TDR Differential Impedance from the Trace End

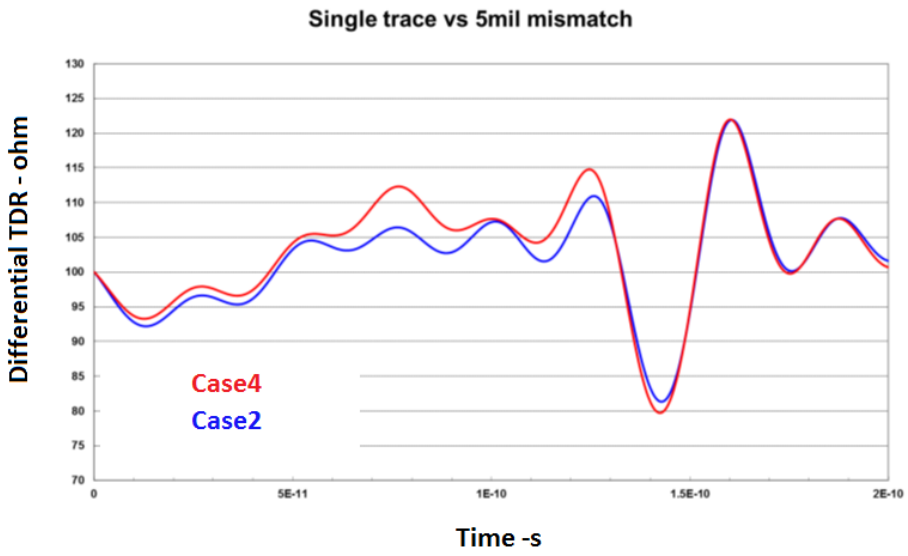

These performance results demonstrate that within 15 GHz bandwidth, Case 2 has a more robust layer-to-layer fabrication mismatch. The TDR impedance shows up to a 5Ω impedance mismatch occurring due to the 5 mil layer-to-layer mismatch.

Intel recommends that the break-out routing is either differential neck-down with back-jog or single-ended bak-jog routings due to the fixed skew matching at the BGA area. The following figures demonstrate the recommended break-out routing performances.

Figure 28. Case 5: Differential Routing with Neck-down and Back-jog

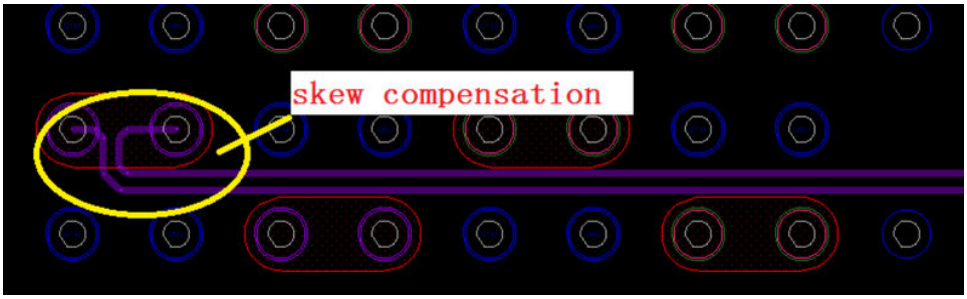

Figure 29. Case 6: Single-Ended Routing with Back-jog

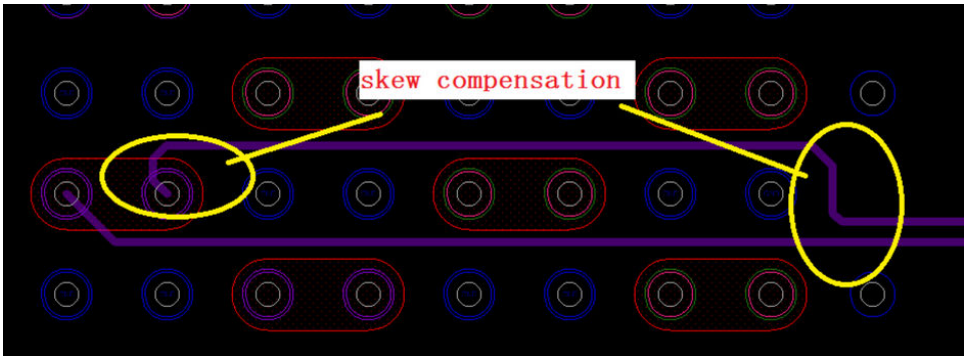

Performance Comparison Between Case 1 and Case 5

This section compares the performances between conventional differential routing with neck-down and jog-out and differential routing with neck-down with back-jog (Case 1 vs Case 5).

Figure 30. Simulated Differential Insertion Loss

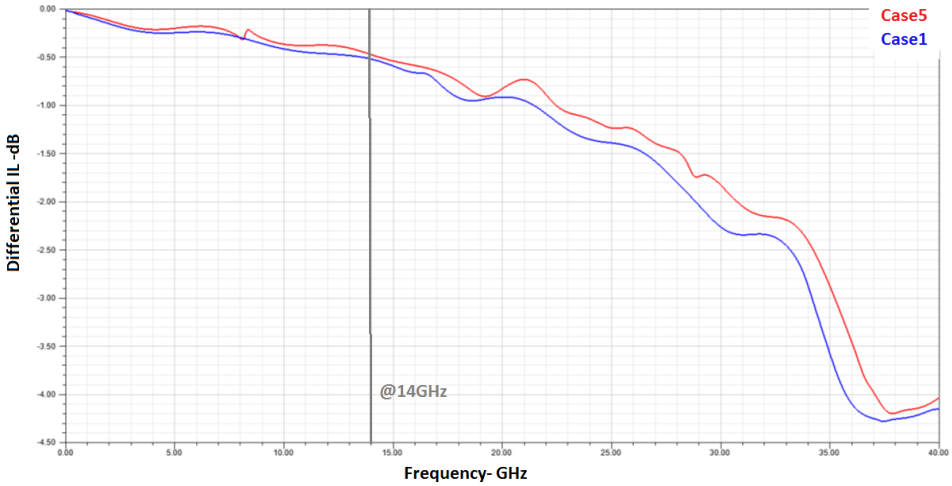

Figure 31. Simulated Differential Return Loss at the BGA Solder Ball

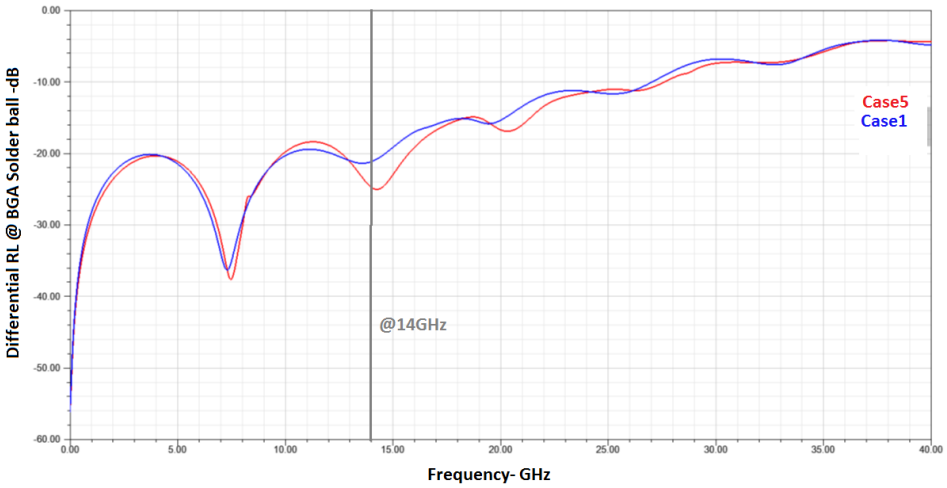

Comparing the conventional differential and jog-out routing configurations with the recommended differential routing with back-jog shows up to 0.1 dB insertion loss improvement. It also exhibits up to a 5 dB return loss improvement within a 15 GHz bandwidth when using the back-jog configuration.

Performance Comparison Between Case 2 and Case 6

This section compares the performances between conventional single-ended in-line breakout routing with jog-out and single-ended routing with back-jog (Case 2 vs Case 6).

Figure 32. Simulated Differential Insertion Loss

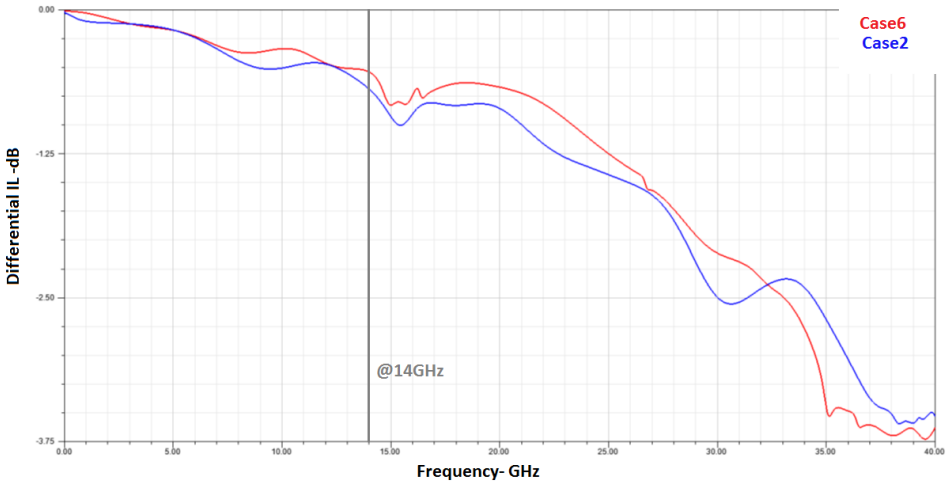

Figure 33. Simulated Differential Return Loss at the FPGA Solder Ball

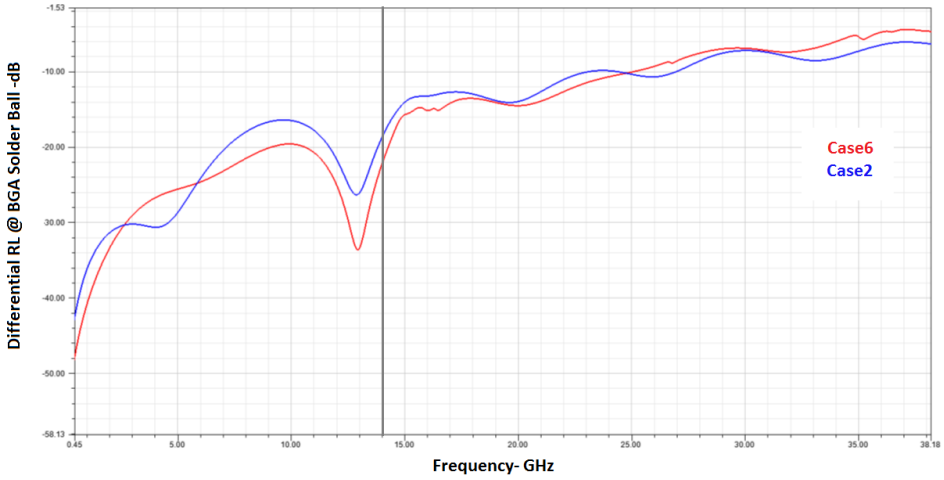

Single-ended break-out routing is less sensitive to layer to layer mismatch.

Comparing the conventional single-ended routing with jog-out and the recommended single-ended routing with back-jog, the single-ended routing with back-jog shows an insertion loss improvement of up to 0.25 dB. It also exhibits a return loss improvement up to 7 dB within a 1 -GHz bandwidth.

The differential break-out routing with back-jog shows 0.1 db insertion loss improvement and up to 6 dB return loss improvement within 15 GHz bandwidth compared to single-ended break-out routing with back-jog.

In addition, differential break-out routing with back-jog has slightly better performances above 15 GHz compared to single-ended routing with back-jog.

Intel recommends use of single-ended routing with back-jog due to less sensitivity over the layer-layer mismatch, if customers use low number of transceiver channels or they have enough signal layers for routing. Due to this, differential break-out routing with back-jog is preferred to single-ended break-out routing with back-jog, because differential break-out routing requires half of routing layers used for single-ended break-out routing.

The Impact of Via Height and Via Anti-pad Over Insertion Loss

This section demonstrates how much via length and anti-pad size change the insertion loss. The results of this examination are based on the following criteria:

- A 24-layer stackup
- A via-in-pad with a 10 mil drill (finished size) hole
- 18 mil pad

Figure 34. 24-Layer Stack-up Information for investigations on insertion loss impact using various via lengths and anti-pad sizes.

Layer		Cu Weight	Proposed Thickness (mils)	Construction	Assume copper denisty	Ref.	Single End 50ohm +/- 10%		Differential 95ohm +/- 10% outer layers +/-5% inner layers		Diffential 100ohm +/- 10% outer layers +/-5% inner layers	
							Target LW	Finished LW	Target LW	Finished LW	Target LW/SP	Finished LW/SP
			0.70									
L1	Top	1/3oz + Plate	2.20	HTE	50%	L2		5.0		3.5/4.0		3.5/5.0
	prepreg		2.80	R5670G 1080RC64								
L2	GND	0.5 oz	0.60	HVLP	75%							
	Core		4.00	R5775G 3313*1								
L3	SIG1	0.5 oz	0.60	HVLP	25%	L2 & 4		4.2		3.8/4.1		3.5/4.5
	prepreg		4.00	R5670G 1035RC70 x 2								
L4	GND	0.5 oz	0.60	HVLP	75%							
	Core		4.00	R5775G 3313*1								
L5	SIG2	0.5 oz	0.60	HVLP	25%	L4 & 6		4.2		3.8/4.1		3.5/4.5
	prepreg		4.00	R5670G 1035RC70 x 2								
L6	GND	0.5 oz	0.60	HVLP	75%							
	Core		4.00	R5775G 3313*1								
L7	SIG3	0.5 oz	0.60	HVLP	25%	L6 & 8		4.2		3.8/4.1		3.5/4.5
	prepreg		4.00	R5670G 1035RC70 x 2								
L8	GND	0.5 oz	0.60	HVLP	75%							
	Core		4.00	R5775G 3313*1								
L9	SIG4	0.5 oz	0.60	HVLP	25%	L8 & 10		4.2		3.8/4.1		3.5/4.5
	prepreg		4.30	R5670G 1035RC73 x 2								
L10	PWR1	2 oz	2.40	RTF	75%							
	Core		3.00	R5775G 1078*1								
L11	GND	2 oz	2.40	RTF	75%							
	prepreg		3.10	R5670G 1035RC70 x 2								
L12	PWR2	2 oz	2.40	RTF	75%							
	Core		3.00	R5775G 1078*1								
L13	PWR3	2 oz	2.40	RTF	75%							
	prepreg		3.10	R5670G 1035RC70 x 2								
L14	GND	2 oz	2.40	RTF	75%							
	Core		3.00	R5775G 1078*1								
L15	PWR4	2 oz	2.40	RTF	75%							
	prepreg		4.30	R5670G 1035RC73 x 2								
L16	SIG5	0.5 oz	0.60	HVLP	25%	L15 & 17		4.2		3.8/4.1		3.5/4.5
	Core		4.00	R5775G 3313*1								
L17	GND	0.5 oz	0.60	HVLP	75%							
	prepreg		4.00	R5670G 1035RC70 x 2								
L18	SIG6	0.5 oz	0.60	HVLP	25%	L17 & 19		4.2		3.8/4.1		3.5/4.5
	Core		4.00	R5775G 3313*1								
L19	GND	0.5 oz	0.60	HVLP	75%							
	prepreg		4.00	R5670G 1035RC70 x 2								
L20	SIG7	0.5 oz	0.60	HVLP	25%	L19 & 21		4.2		3.8/4.1		3.5/4.5
	Core		4.00	R5775G 3313*1								
L21	GND	0.5 oz	0.60	HVLP	75%							
	prepreg		4.00	R5670G 1035RC70 x 2								
L22	SIG8	0.5 oz	0.60	HVLP	25%	L21 & 23		4.2		3.8/4.1		3.5/4.5
	Core		4.00	R5775G 3313*1								
L23	GND	0.5 oz	0.60	HVLP	75%							
	prepreg		2.80	R5670G 1080RC64								
L24	BOT	1/3oz + Plate	2.20	HTE	50%	L23		5.0		3.5/4.0		3.5/5.0
	Soldermask		0.70									
Finished Thickness (mils)			115.20									

Figure 35. Differential Via and Anti-pad Configuration

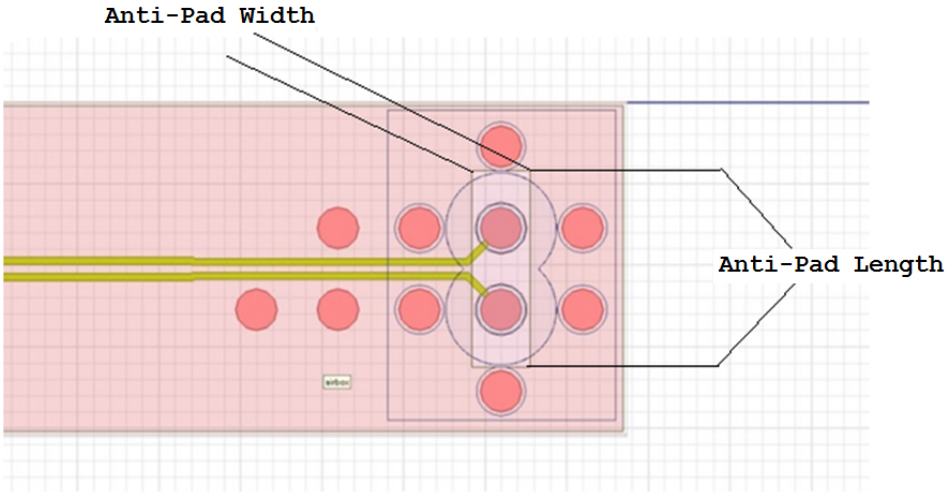

Table 3. Impact of Via Length on Insertion Loss

In this case, the anti-pad (GND cutout for each GND layer up to the backdrill) is fixed at 95 mil x 28 mil.

Case	Insertion Loss @ 14 GHz	Return Loss @14 GHz,~dB
Via length from top to L3 (7.4 mil)	-0.1730 dB	-29
Via length from top to L5 (16.6 mil)	-0.1768 dB	-28
Via length from top to L7 (25.8 mil)	-0.1828 dB	-27

Figure 36. Impact on Differential Insertion Loss by Via Length (SDD21)

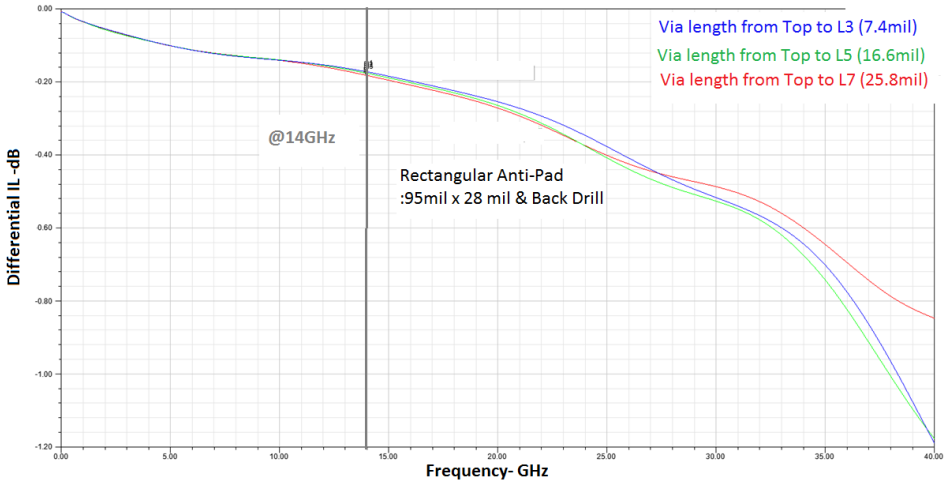

Figure 37. Impact on Differential Return Loss by Via Length (SDD11)

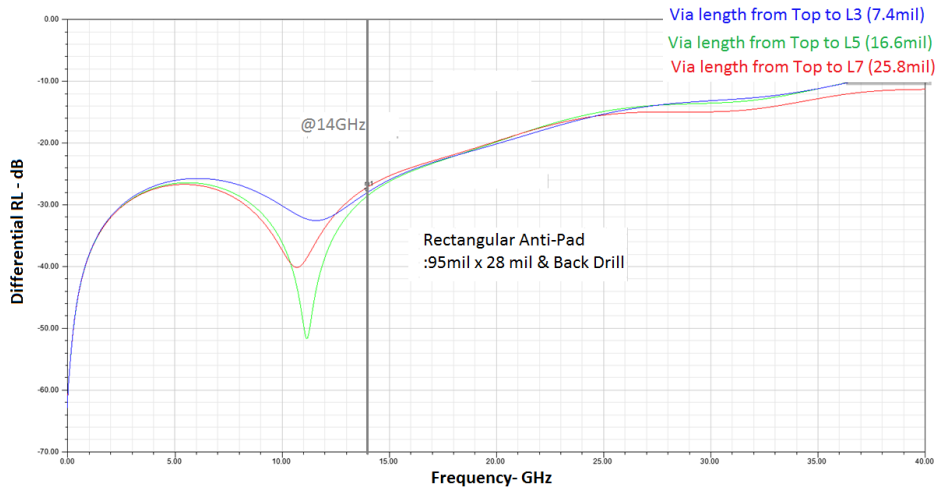

Table 4. Impact of Via Anti-pad Length on Insertion Loss

The via length from the top to layer 3 is 7.4 mil.

Case	GND02 Cutout Size	GND04 Cutout Size	IL 14 GHz
1A	68 mil x 28 mil	68 mil x 28 mil	-0.1988 dB
2A	95 mil x 28 mil	95 mil x 28 mil	-0.1730 dB

Table 5. Impact of Via Anti-pad Width on Insertion Loss

For various GND cutouts or (or anti-pad), the via length from the top to layer 5 is 16.6 mil.

Case/Cutout	GND02	GND04	GND06	IL @14 GHz
1B	95 mil x 28 mil	95 mil x 28 mil	95 mil x 28 mil	-0.1768 dB
2B	95 mil x 50 mil	95 mil x 50 mil	95 mil x 50 mil	-0.1741 dB

The signal anti-pad size results in these tables demonstrates that the larger the anti-pad, the less insertion and return loss results. Because space is always a consideration, you must balance the signal anti-pad size with the need for proper break-out routing.

AC Coupling Capacitor Layout and Optimization Guidelines

It is possible to use both the 0402 and 0201 capacitor sizes on boards as AC coupling capacitors on transceiver links.

Figure 38. AC Capacitor Placement and GND Cutout on PCB with Stripline Routing on Both Ends

Structural detail where trace routing is stripline and the AC capacitor is mounted on the top or bottom layers.

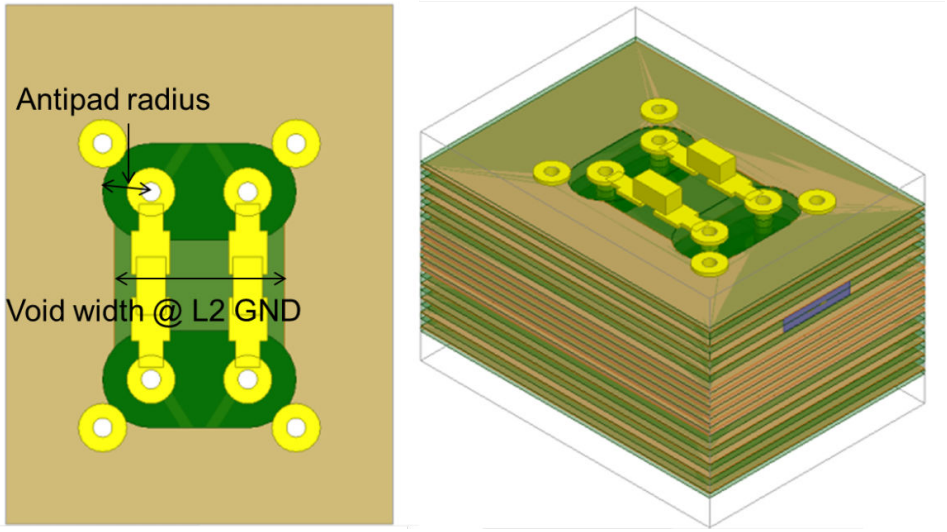

The structural detail includes the following specifications:

- 10 mil drill hole and 20 mil pad
- The cap is mounted on the top layer and the trace breakout is routed on layer 7
- 10 mil stub length
- 0201 capacitor size copper block
  - 24 mil x 12 mil x 12 mil
- 0402 capacitor size copper block
  - 40 mil x 20 mil x 14 mil

The board stack-up configuration includes the following specifications:

- 24 layers and a thickness of 117 mil
- 8 signal layers, 4 PWR lanes
- Megtron6 material

0201 AC Capacitor

Sweeping Anti-pad Radius

Figure 39. Differential Return Loss and TDR Impedance Performances for 0201 AC Capacitors by Various Anti-pad Sizes

Figure shows a fixed rectangular GND cutout under the capacitors while changing only the signal via anti-pad radius.

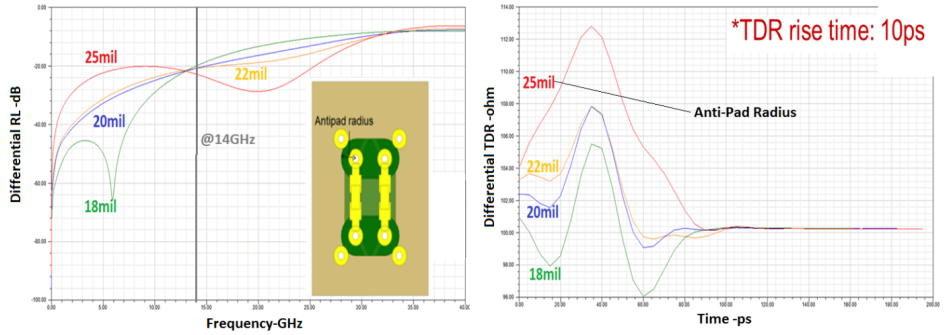

Based on this TDR response, a larger via anti-pad eventually increases the impedance of the entire structure. Using a 22 mil anti-pad radius is an optimum solution for this case and results in the least mismatching.

Sweeping Void Width only on the First GND Plane Under the Capacitor

Keeping the via anti-pad radius fixed while adjusting the void width on the first GND plane under the AC capacitors also impacts the structure's impedance.

Figure 40. Differential Return Loss and TDR Impedance Performances for 0201 AC Capacitors by Various GND Cutout Widths

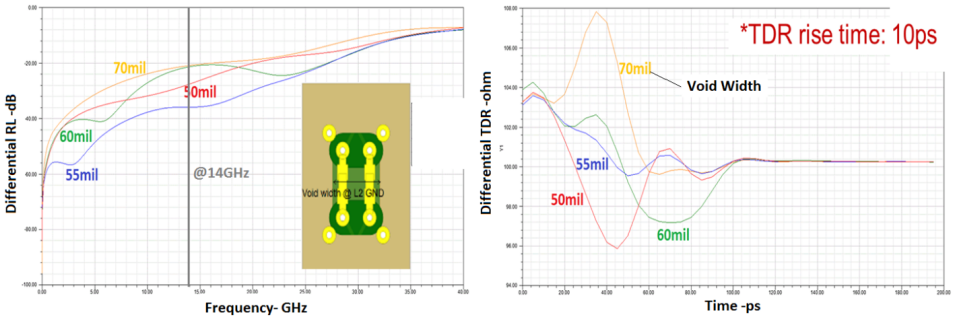

Changing the void width from 50 mil to 70 mil impacts the structure's impedance and return loss. A 55 mil void width is the optimum solution for this case and results in the least mismatching.

0402 AC Capacitor

Sweeping Anti-pad Radius

Figure 41. Differential Return Loss and TDR Impedance Performances for 0402 AC Capacitors by Various Anti-pad Sizes

Figure shows a fixed rectangular GND cutout under the capacitors while changing only the signal via anti-pad radius.

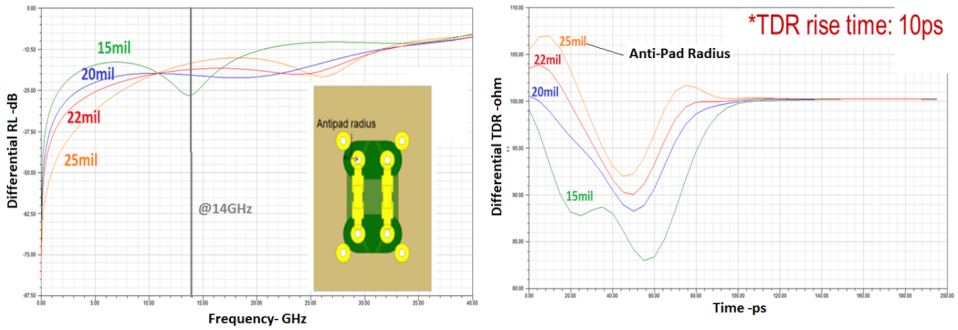

The large via anti-pad increases the impedance of the structure. Using a 22 mil via anti-pad is an optimum solution for the 0402 capacitor and results in the least mismatching with the rest of the system.

Sweeping Void Width Only on the First GND Plane under the Capacitor

Keeping the via anti-pad radius fixed while adjusting the void width on the first GND plane under the AC capacitors also impacts the structure's impedance.

Figure 42. Differential Return Loss and TDR Impedance Performances for 0402 AC Capacitors by Various GND Cutout Widths

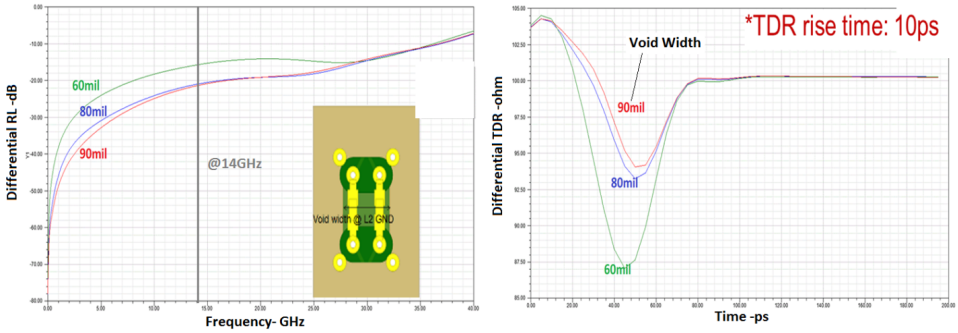

Changing the void width from 60 mil to 90 mil impacts the structure's impedance and return loss. A 90 mil void width is the optimum solution for this case and results in the least mismatching.

Adding a Trace Reference to the 0201 AC Capacitor

A small plane is added as a reference plane on the anti-pad void area on both sides of the GND planes for the stripline to lower the impedance of the breakout trace. This section examines the impact of the added reference plane width.

Figure 43. Configuration of the Added Trace Reference for the 0201 AC Capacitor GND Cutout

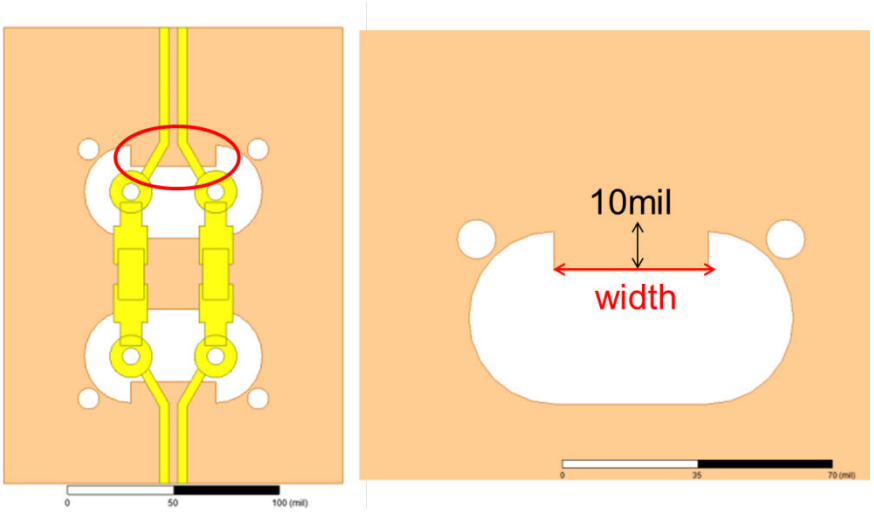

Figure 44. Differential Return Loss and TDR Performances by Different Widths of Adding Trace Reference for the 0201 AC Capacitor GND Cutout

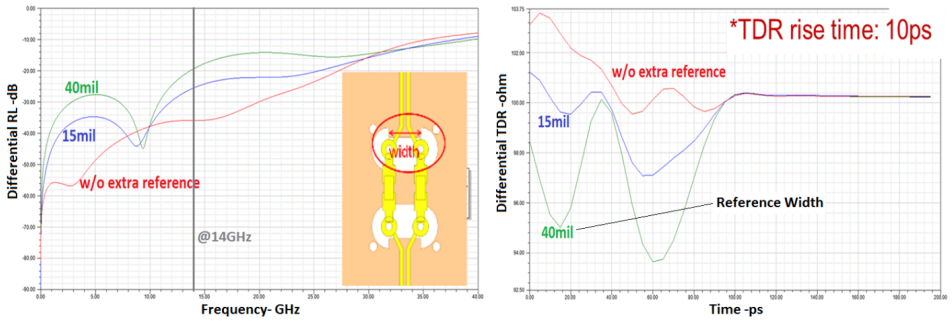

The extra reference plane for the stripline trace lowers the impedance fluctuation while connected to the signal via. A 15 mil width is a good solution for lowering impedance mismatching for this case.

If the AC capacitor is located close to the connector, it is better to use microstrip routing to reach the connector from the AC capacitor.

Figure 45. 0201 AC Capacitor Placement and GND Cutout on the PCB

This figure shows stripline routing on one end and microstrip routing on the other.

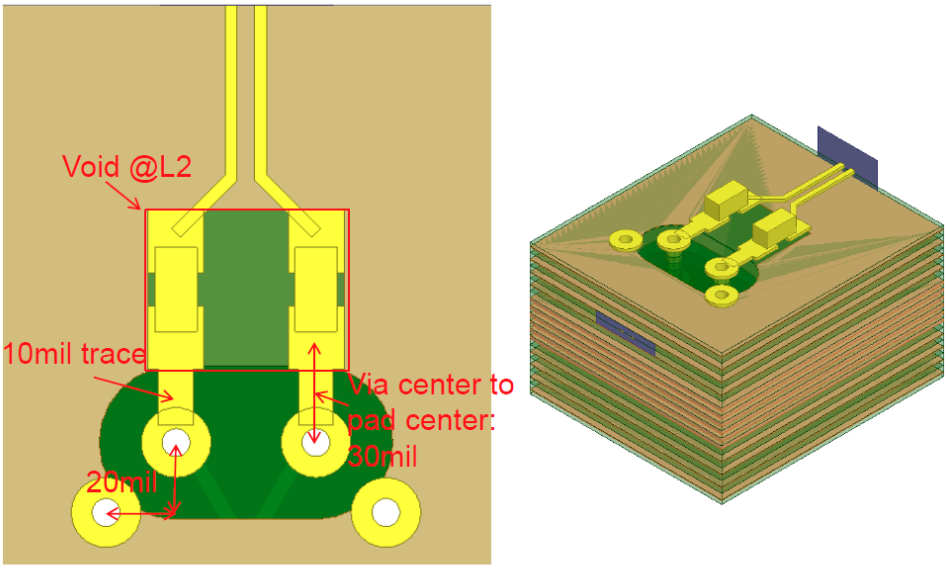

The top layer breakout performance above is preferred because only one layer transition via is included. The improved performance depends on the AC capacitor's close proximity to the connector.

Adding a Trace Reference to the 0402 AC Capacitor

Figure 46. Configuration of the Added Trace Reference for the 0402 AC Capacitor GND Cutout

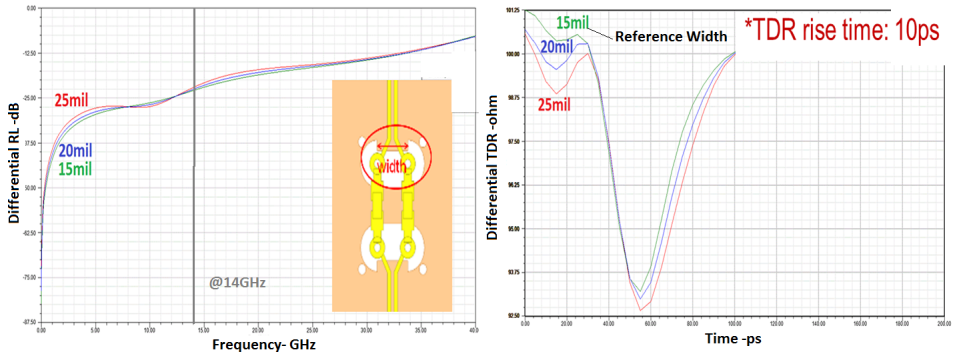

For the 0402 capacitor, the optimum width is approximately 20 mil to lower the impedance mismatching in this structure.

PCB AC Capacitor Placement Layout Recommendation

Figure 47. Recommended 0402 AC Capacitor Layout

For this layout, there is no dependency on the stackup.

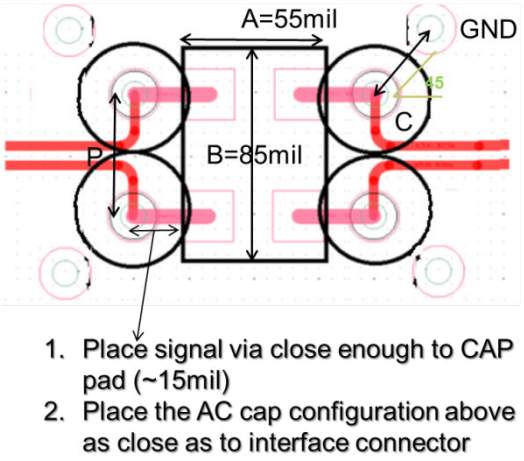

Follow these guidelines for 0402 AC capacitor layout:

- AC capacitors are mounted on either the top or bottom layers for RX lanes (for some specific interfaces)
- If the AC capacitor is mounted on the top layer, RX routing must be stripline with back-drill as close to the top layer as possible to reduce the AC capacitor signal via height
- If the AC capacitor is mounted on the bottom layer, RX routing must be stripline as close as possible to the bottom layer with a back-drill of the via from the top
- A x B GND cutout under the AC capacitor (only one layer GND cutout)
- Signal-to-signal via pitch P = 40 mil
- Signal-to-GND via pitch C = 30 mil
- Signal/GND via drill diameter = 10 mil
- Signal/GND via pad diameter = 20 mil
- Signal anti-pad diameter = 45 mil

Figure 48. Recommended 0201 AC Capacitor Layout

For this layout, there is no dependency on the stackup.

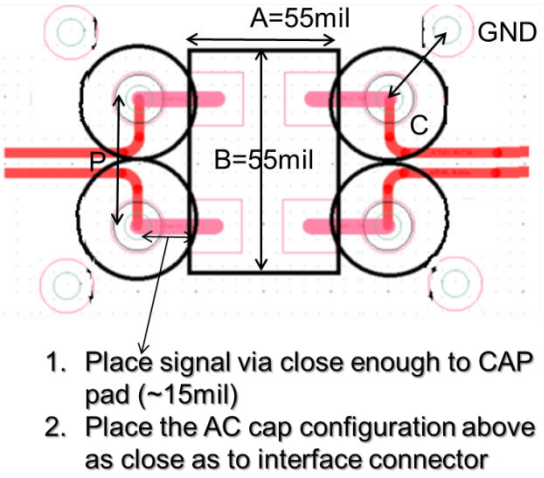

Follow these guidelines for 0201 AC capacitor layout:

- A x B GND cutout under the AC capacitor (only one layer GND cutout)
- Signal-to-signal via pitch P = 40 mil
- Signal-to-GND via pitch C = 30 mil
- Signal/GND via drill diameter = 10 mil
- Signal/GND via pad diameter = 20 mil
- Signal anti-pad diameter = 45 mil

Where possible, Intel recommends using the 0201 AC capacitor rather than the 0402 AC capacitor for improved differential return loss and TDR performance.

Figure 49. Differential Return Loss and TDR Comparison Between the 0402 and 0201 AC Capacitor Layouts

Figure shows results after optimization of the GND cutout and signal via anti-pad. Both structures use a 22 mil signal via anti-pad radius. The 0201 AC capacitor has a 55 mil GND cutout width and the 0402 AC capacitor has a 90-mil GND cutout width on the first GND layer under the capacitors.

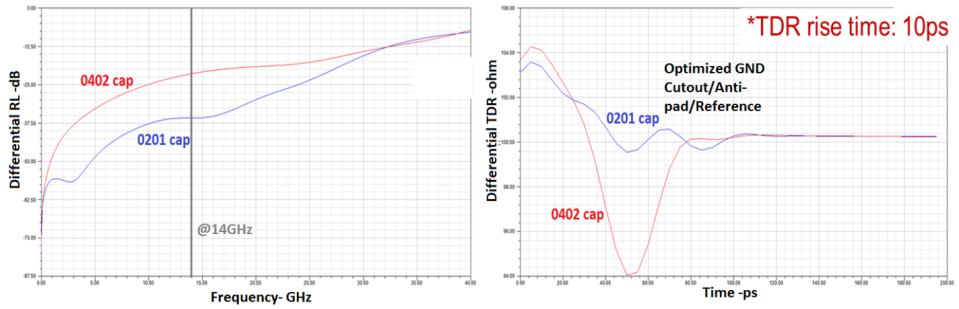

What to Avoid for AC Capacitor Configuration

The differential signal vias must be coupled to reduce the inductance for each individual signal via. Intel recommends tightly-coupled AC capacitors over single-ended AC capacitor placement.

Figure 50. Avoided vs Recommended AC Capacitor Layout Configuration Comparison

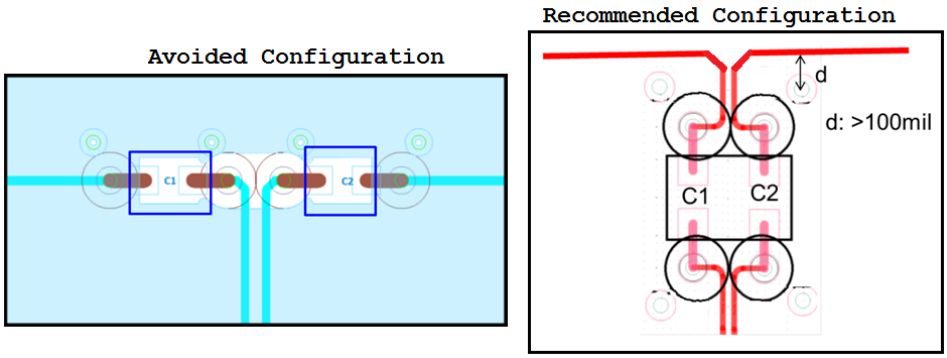

CFP2/CFP4 Connector Board Layout Design Guideline

CFP2 Host Connector, Module Assembly, and Pinout

Figure 51. CFP2 Host Connector Assembly and N X 25 Gbps Pin Map

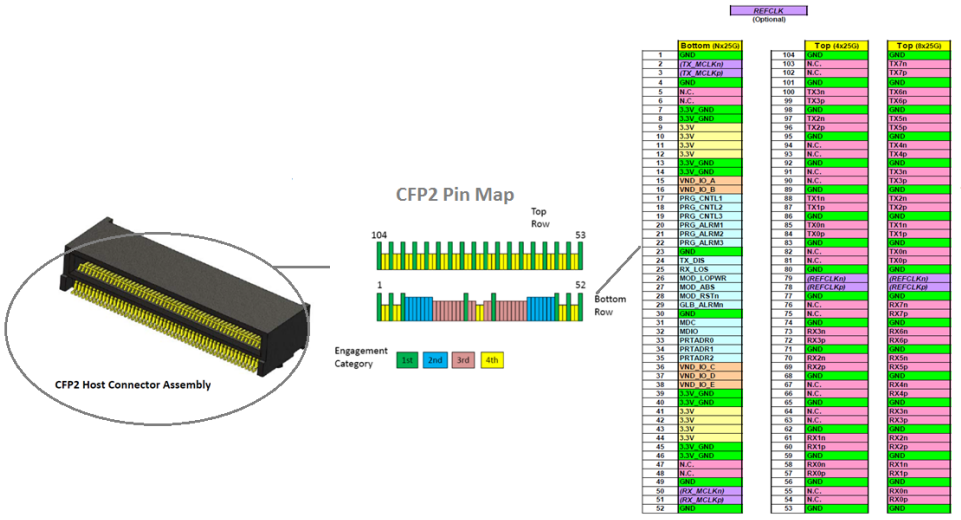

The CFP2 module supports up to 10 channel at up to 25 Gbps. It also has future support for up to 8 channels at up to 50 Gbps. This high speed electrical interface will be AC coupled within the CFP2 module. The 25 Gbps specification is defined in the OIF-CEI-28G-VSR.

Note: For more information, refer to the CEI-28G-VSR working clause specification. Document number OIF2010.404.08.

CFP4 Host Connector, Module Assembly, and Pinout

The CFP4 module high speed electrical interface supports the following configurations:

- 4 TX lanes + 4 RX lanes, each at 25 Gbps
- 4 TX lanes + 4 RX lanes, each at 10 Gbps

Figure 52. CFP4 Host Connector Assembly and N X 25 Gbps Pin Map

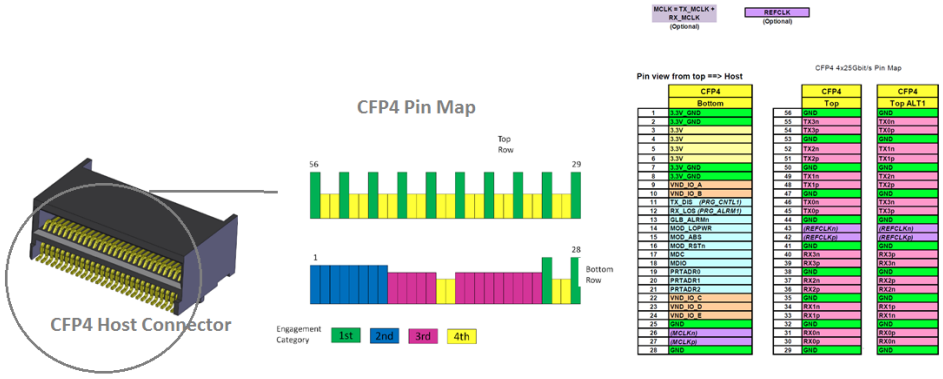

The high speed electrical interface will be AC-coupled within the CFP4 module. The 25 Gbps specification is defined in the OIF-CEI-28G-VSR.

Note: For more information, refer to the CEI-28G-VSR working clause specification. Document number OIF2010.404.08.

Recommended PCB Design Guideline at the CFP2/CFP4 Connector

Figure 53. Recommended CFP2/CFP4 Connector Layout on the PCB

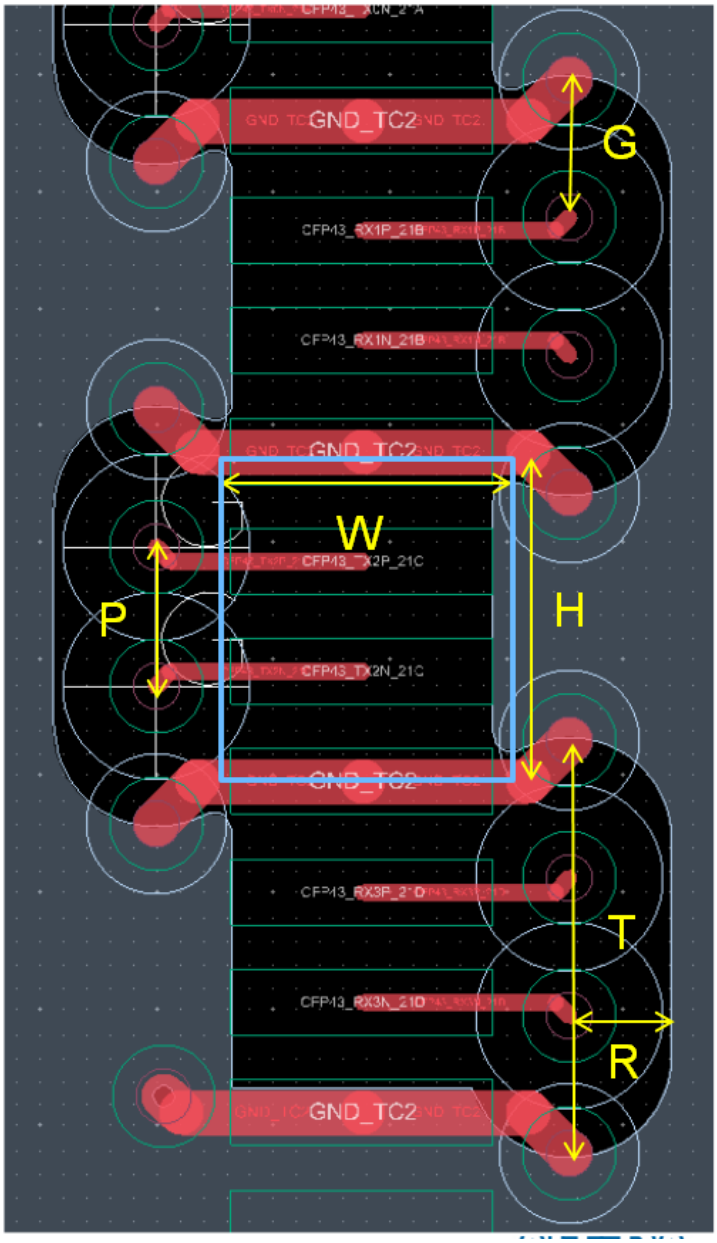

Intel recommends that you observe the following design guidelines:

- For CFP4: rectangular (W=62 mil x H=60 mil) cutout on both layers GND02/ GND04 [1] under the CFP4 pads
- Signal and GND vias, drill diameter = 10 mil, Via pad diameter = 20 mil
  - All signal vias are back-drilled
  - Each signal via must have one single GND via
- Signal anti-pads: T = 90 mil, R = 22.5 mil
- G (Signal-to-GND via pitch) = 30 mil

[1] This refers to the first and second GND layers under the connector. If the stack-up height is less than 1:12, you can also use an 8 mil finished via drill and 18 mil via pad.

Note: The actual drill is 10 mil, but the copper filling inside makes it an 8 mil finished drill.

Intel recommends that you have both the signal via and GND via located closely enough to the connectors' signal and GND pads, respectively, to avoid cavity resonance and higher frequencies.

Figure 54. Recommended CFP2/CFP4 Connector Fan-out Routing on the PCB

Routing layers are differentiated by different colors.

The differential lanes in green can be routed on any signal layer.

The differential lanes in blue can only be routed on signal layers where the GND reference layers are not GND02/GND04 because they cross the cutout area under the connector. You can route these lanes from the opposite direction (similar to the green lanes) if there is space for routing. In this case, all signal layers can be used for signal routing.

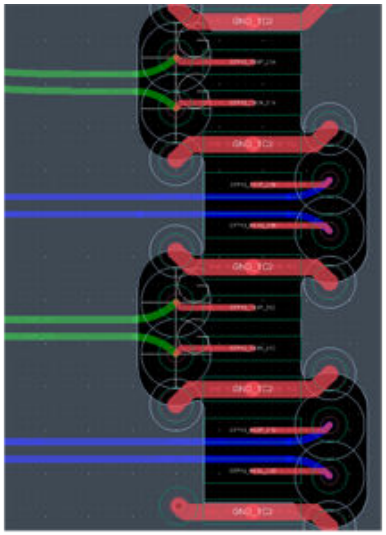

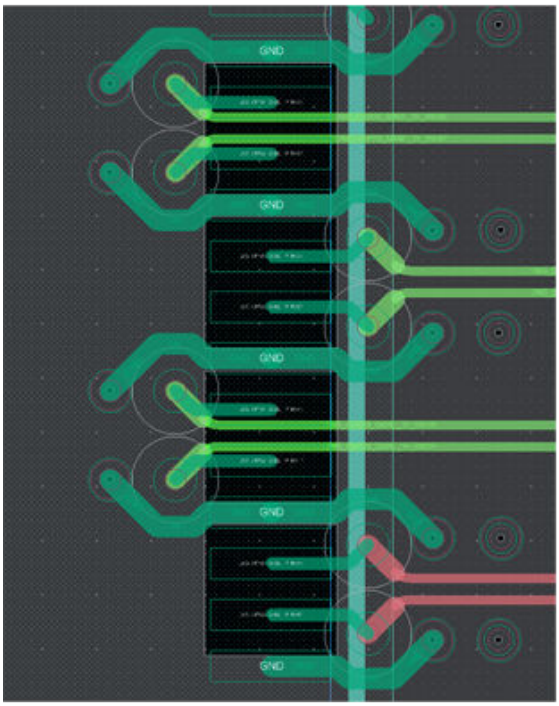

Make sure that you have proper GND reference plane for signal routing.

Observe these CFP2/CFP4 guidelines for better performance at 28 Gbps on the main channel:

- Match the length for each pair (between P and N lanes). Both P and N lanes must be in phase to recover the data. The skew matching in a pair is 2 ps.
- Length matching between pairs is not mandatory unless it is specified by designers.
- For optimized FPGA break-out layout design, refer to *FPGA Fan-out Region* chapter.
- Always use minimal routing length from the FPGA to the connector to achieve minimum insertion loss. Refer to *PCB Stackup Selection Guideline* chapter for Stackup and material selection and *Recommendations for High Speed Signal PCB Routing* chapter for HSSI PCB routing.
- Ensure that the insertion loss and return of the channel is within specifications. Refer to *Electrical Specifications* chapter for specifications.

Related Information

- PCB Stackup Selection Guideline on page 6
- Recommendations for High Speed Signal PCB Routing on page 9
- FPGA Fan-out Region Design on page 10
- Electrical Specifications on page 87

CFP4 Design Example and Optimization Performance

CFP4 Connector Layout: Signal Via, Trace Routing Impact, and Optimization

This design example shows the impact of a via anti-pad diameter, trace width at the void area, and the main transmission line impedance at the CFP4 connector. Traditional designs for signal via anti-pad at CFP4 have been for a 50 mil diameter. Reducing via anti-pad diameter to 40 mil shows more impedance matching and less reflection. Changing the trace width on the void area also enhances TDR, IL, and RL performances. Combining both these approaches to reduce the signal via anti-pad diameter and increase the trace width on the void area eventually improves reflections in this area. As mentioned in the fan-out for a 90Ω TL routing impedance, you can see the results of this approach at the connector to see the impact on IL and RL.

Figure 55. Original and Optimized Anti-pad, Trace Width, and TL Impedance Configurations at the CFP4 Connector Area on the PCB

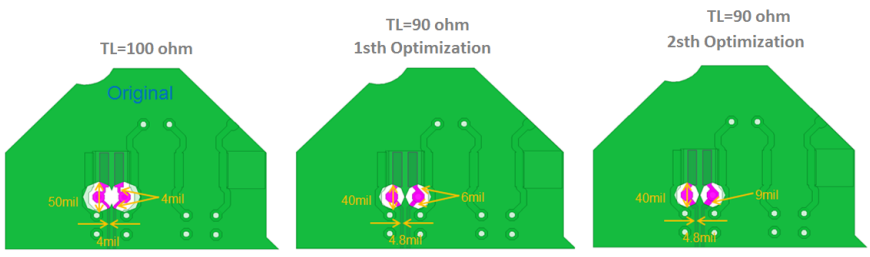

Figure 56. TDR and Full Channel IL and RL Performances

Original and optimized anti-pad, trace width, and TL impedance configurations and the CFP4 connector area on the PCB.

The red line indicates a 4 mil trace width on the void area with a 50 mil anti-pad, and 100Ω TL impedance.

The dark blue line indicates a 6 mil trace width on the void area with a 40 mil anti-pad, and a 90Ω TL impedance.

The pink line indicates a 9 mil trace width on the void area with a 40 mil anti-pad, and a 90Ω TL impedance.

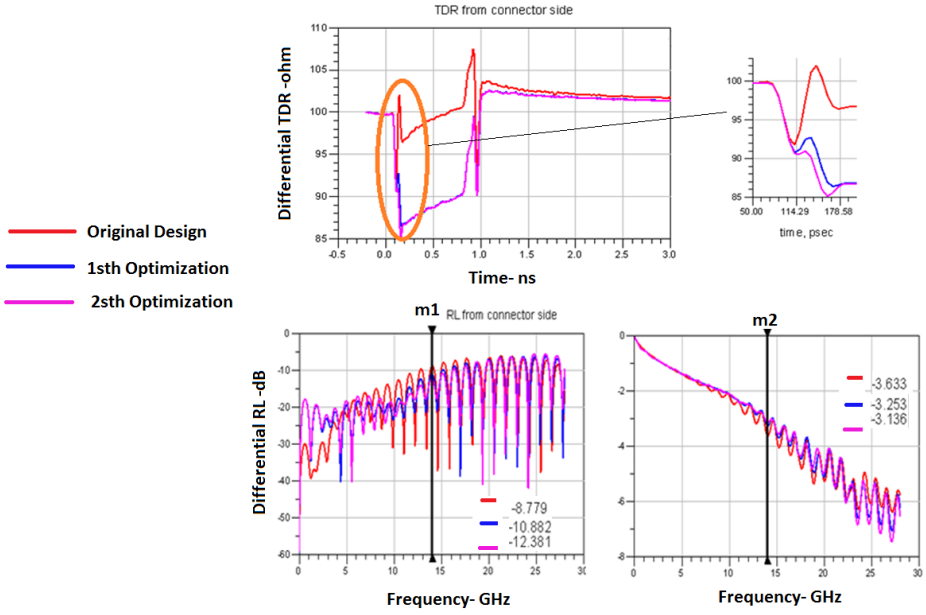

The IL and RL improvement is about 0.5 dB and 3.6 dB at 14 GHz, respectively.

Two Different Break-out Routings at the CFP4 Connector Area

Figure 57. Two Different Break-out Routings at the CFP4 Connector Area

The 3C configuration is what has been previously recommended. The 3G configuration is a different way of routing while the signal vias are moved away from the connector pads to allow more space for break-out routing.

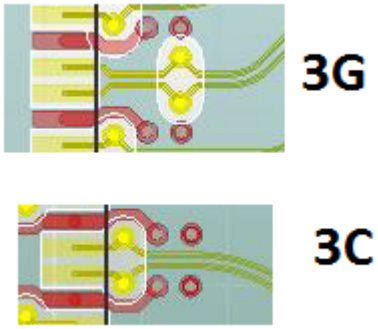

The main PCB routing differential impedance is designed for 95 Ω as was previously recommended.

Figure 58. TDR Performance from the CFP4 Connector Pad on the PCB

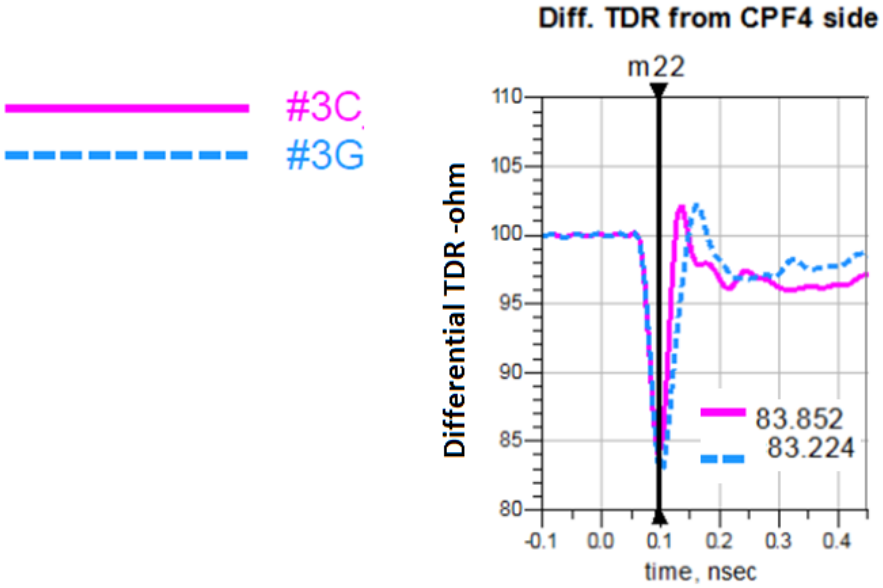

The measured TDR from the connector pads show no appreciable difference.

The total length of 3C is 60 mil larger than 3G.

Note: Intel recommends that you not have two GND vias assigned for one signal via as seen in 3C above.

CFP4 Connector Routing Topologies Design Example

In the figure below, notice that the CFP4 connector has not been included in simulations. The standard 24-layer stack-up at 117 mil thickness is used for this design example. The stack-up material is Megron6 and has eight signal layers and four PDR layers.

Figure 59. Proposed Design Layout Dimensions at the CFP4 Connector

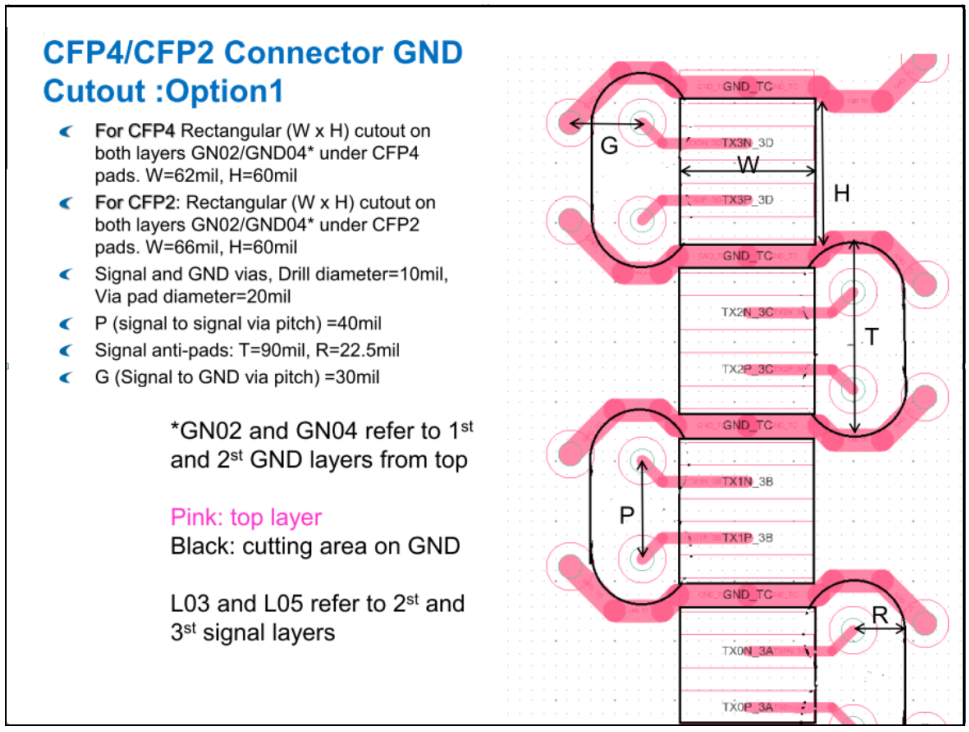

Figure 60. Side and Top Views of the Design Example

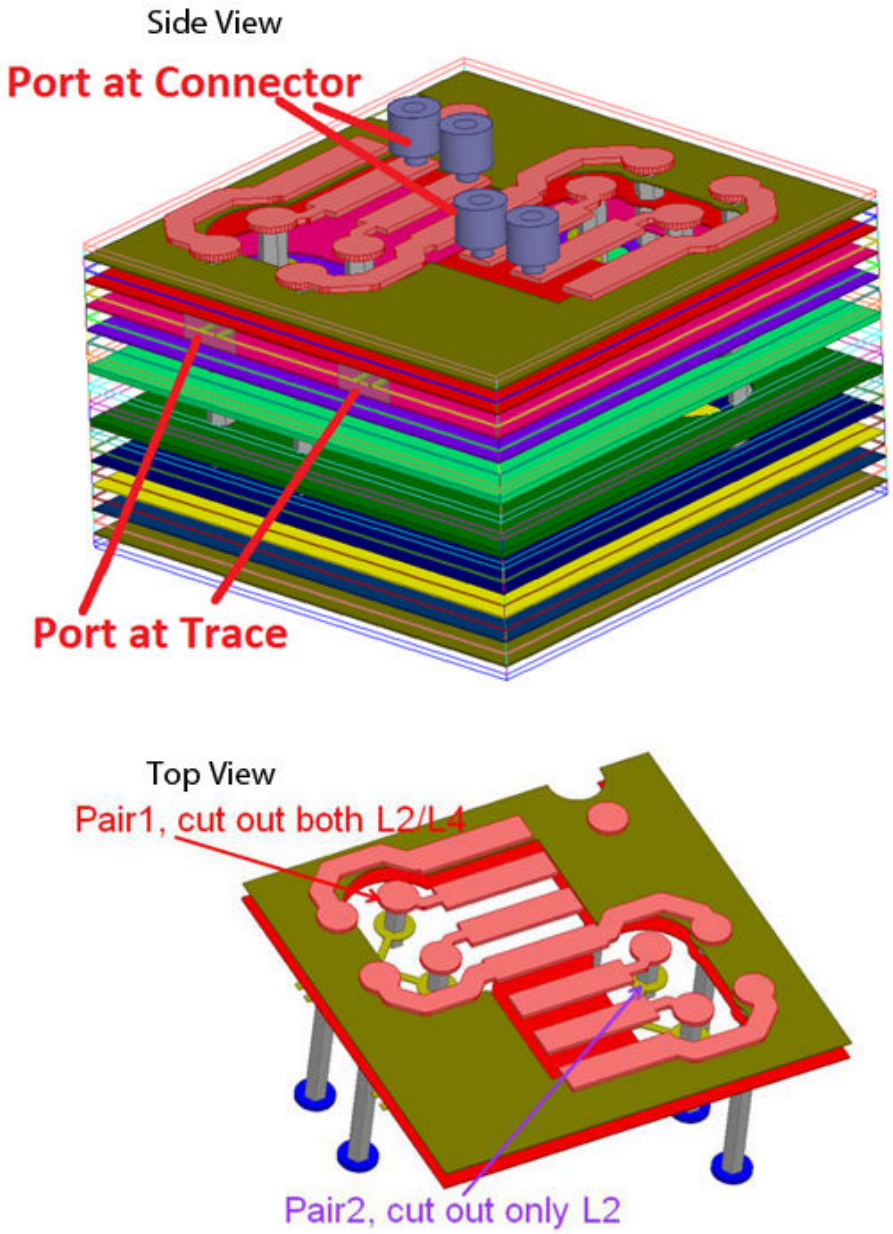

The channel trace routing is on layer 5 for both of the pairs above. Pair1 has two GND cutouts (on layers 2 and 4) under the connector pads, while pair1 has only one GND cutout on layer 2.

The reason for not having a GND cutout on layer 4 for pair 2 is because trace routing needs a reference GND plane while routed and passed under the connector pads. This aids in obtaining a minimum channel path as opposed to breaking out of the opposite direction and looping back, which creates a longer channel path.

The following figures compare the performances of pair 1 and pair 2.

Figure 61. Differential Insertion Loss of Pair 1 and Pair 2

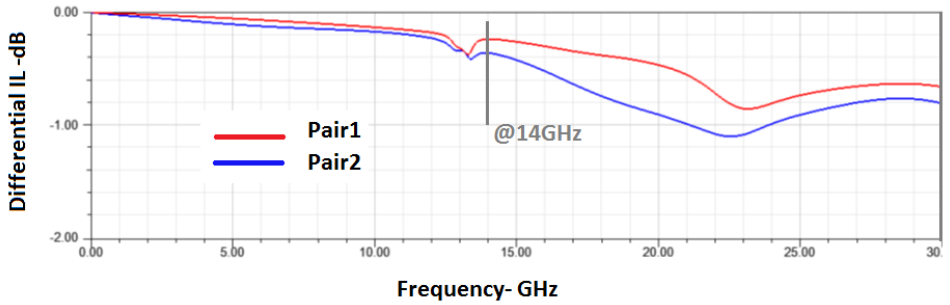

Figure 62. Differential Return Loss from the CFP4 Connector Pad

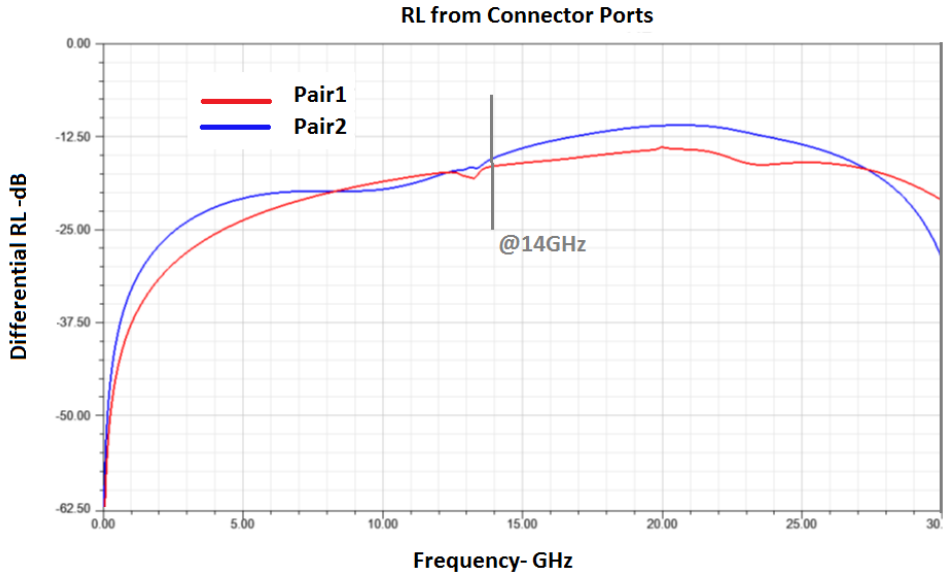

Figure 63. TDR Differential Impedance from the CFP4 Connector Pads

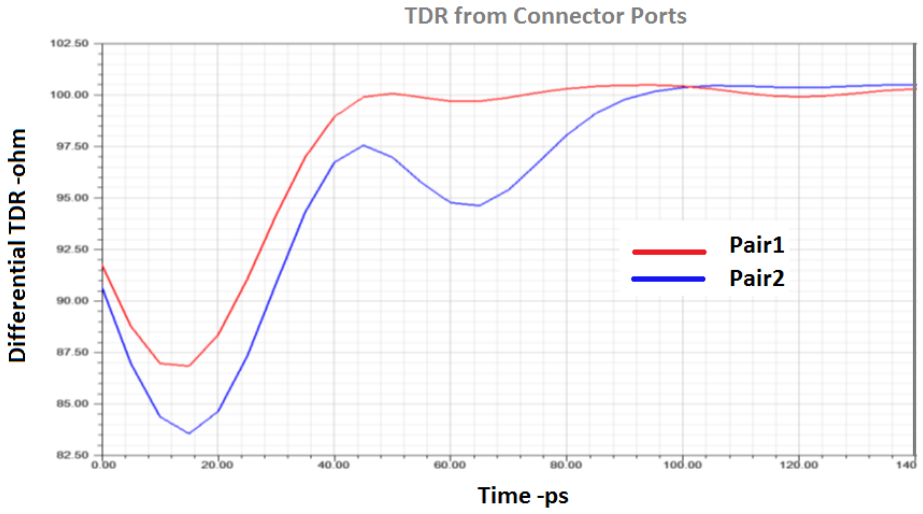

A quick comparison shows that pair 2 has up to 0.2 dB more insertion loss within a 15 GHz bandwidth. This is due to the lack of a GND cutout on layer 4 under the connector pads in addition to the extra routing length under the connector pads. Pair 2 also has up to 2.5 dB return loss with a 15 GHz bandwidth.

The TDR differential impedance also shows less impedance mismatch for pair 1.

In the following figure, the layout configuration has been changed from pair 2 (in the previous example) to pair 3.

Figure 64. Side and Top Views of the Design Example

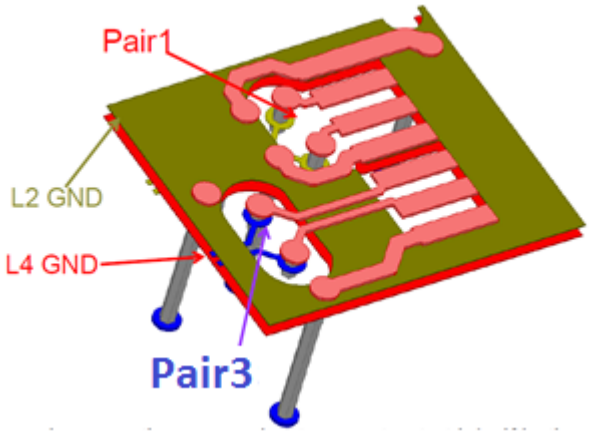

In the example above, pair 3 signal vias have been moved to the left side. This allows an additional GND cutout on layer 4 for pair 2 below the CFP4 connector pads. In this configuration, pair 3 can be routed on signal layer 3. The total routing length for both pair 1 and pair 3 are now equal.

The following figures compare the performances of pair 2 and pair 3.

Figure 65. Differential Insertion Loss

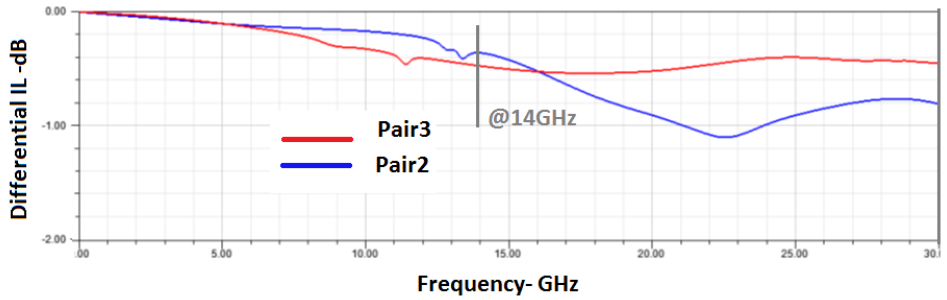

Figure 66. Differential Return Loss at the CFP4 Connector Pads

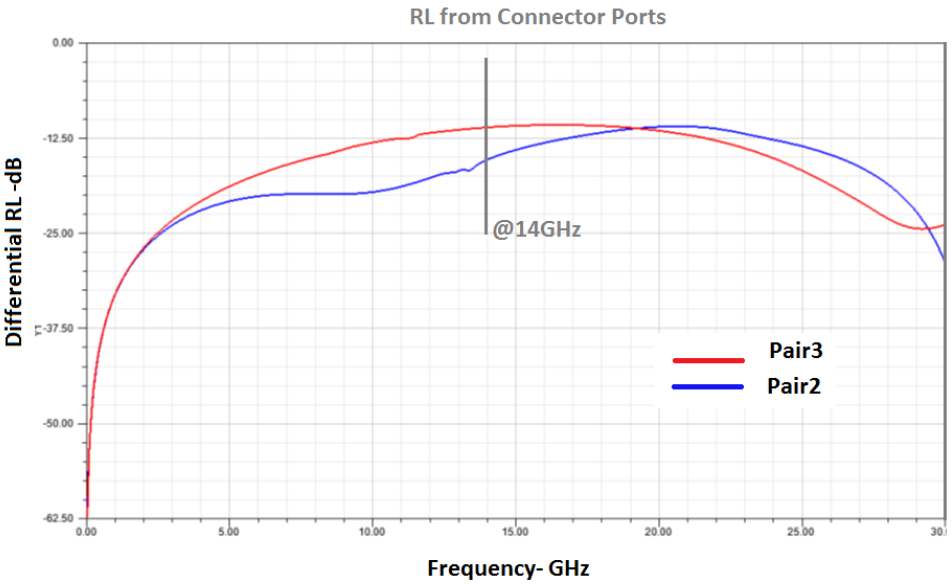

Figure 67. TDR Differential Impedance at the CFP4 Connector Pads

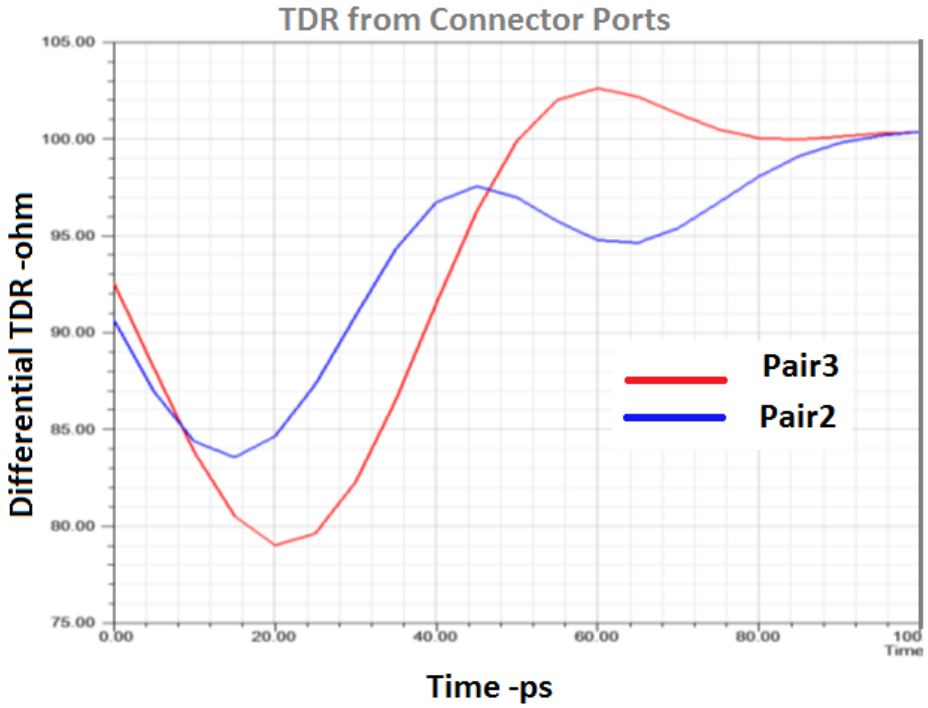

The insertion and return loss results above show better performance for pair 2 within a 15 GHz bandwidth. This is due to a larger GND cutout below the CFP4 connector pads (combined rectangular cut out and signal via anti-pads). This is also observed from TDR differential impedance results. However, above 15 GHz, pair 2 in Figure 60 on page 46 exhibits greater performance degradation than pair3 in Figure 64 on page 48.

Full CFP4 Channel Analysis Design Example (Excluding the Connector)

This section illustrates a full FP4 channel simulation.

Figure 68. Full Channel CFP4 Design Example

Figure shows FPGA to CFP4 connector, excluding the connector.

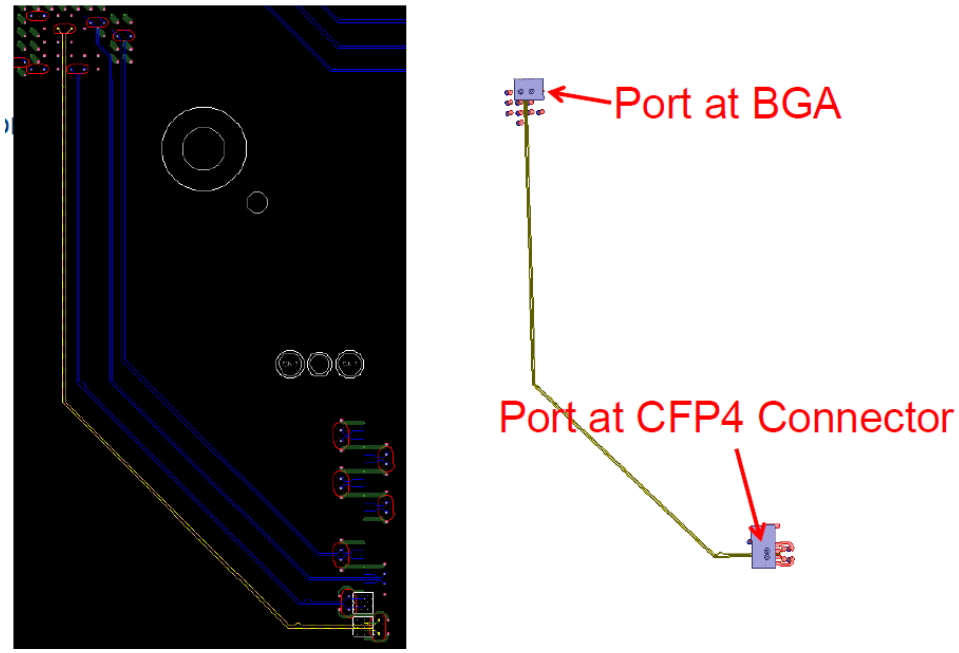

The total channel length is approximately 2.4 inch from the BGA to the connector pads.

The main routing is stripline on layer 5. The connector break-out configuration is similar to what you can see at pair 2 in Figure 60 on page 46.

Figure 69. Full CFP4 Channel Insertion Loss Performance

Figure excludes the CFP4 connector.

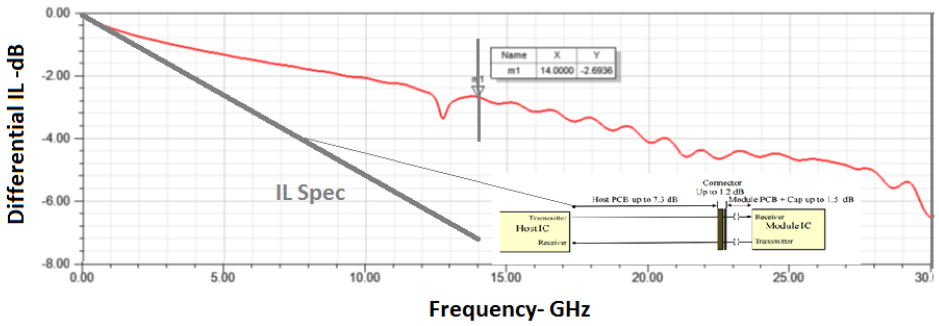

Figure 70. Full CFP4 Channel Return Loss Performance

Figure excludes the CFP4 connector.

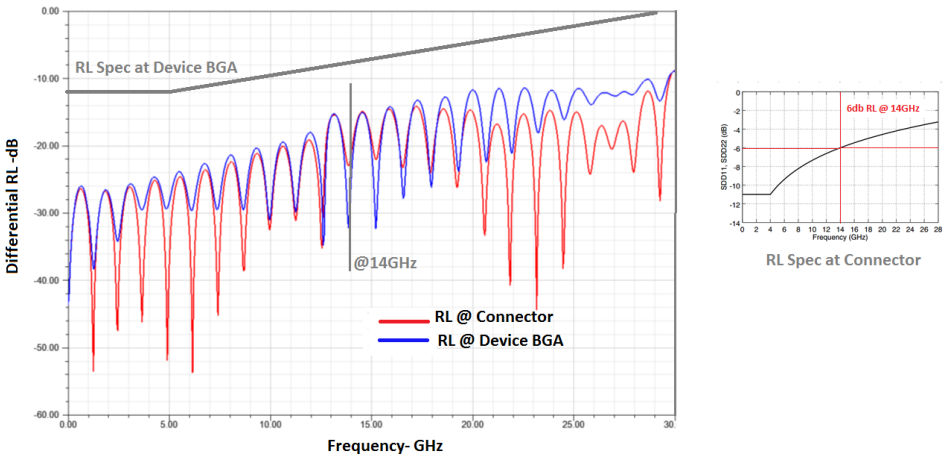

Comparing the scattering performances (SDD21 and SDD11/22) above with the host-to-module specification in CEI-28G-VSR shows that both insertion and return loss meet the specifications.

CFP2 Connector Area Layout

These are the differential TDR results from simulation including the PCB portion connected to the CFP2 connector and host compliance board. The TDR result is based on the following configuration:

- No GND cutout
- One single GND layer cutout
- Two GND layer cutouts below the CFP2 high speed signal pads

The host PCB routing has a 100 Ω impedance.

Figure 71. Simulation Structure of Partial HCB, CFP2 Connector, and Host PCB with TDR Differential Impedance Results

The TDR differential impedance results are from the host PCB while the HCB port is terminated.

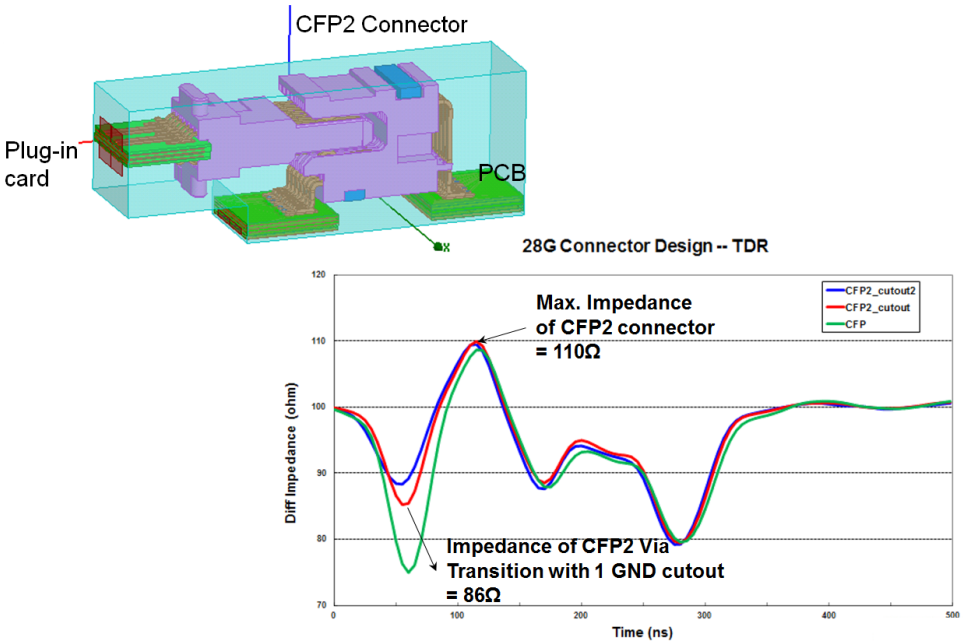

The following figure shows the actual measured differential TDR on an Intel Arria® 10 device populated SI board for CFP2 full channel. Two GND layers are cut out below the high speed connector pads. The reduced differential impedance is approximately 8.5 Ω at the CFP2 connector transition to the PCB.

Figure 72. Intel Arria 10 device PCB Layout at the CFP2 Connector

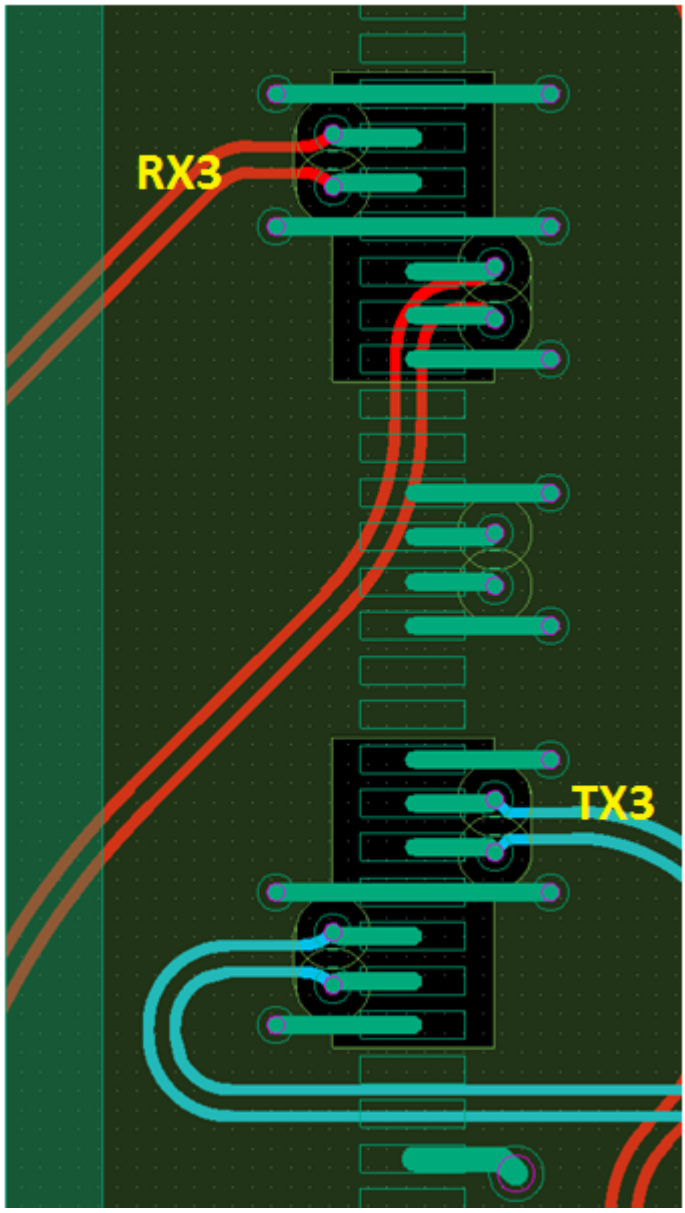

Figure 73. TDR Differential Impedance for Intel Arria 10 device CFP2 Full Channel Including CFP2 Connector and HCB

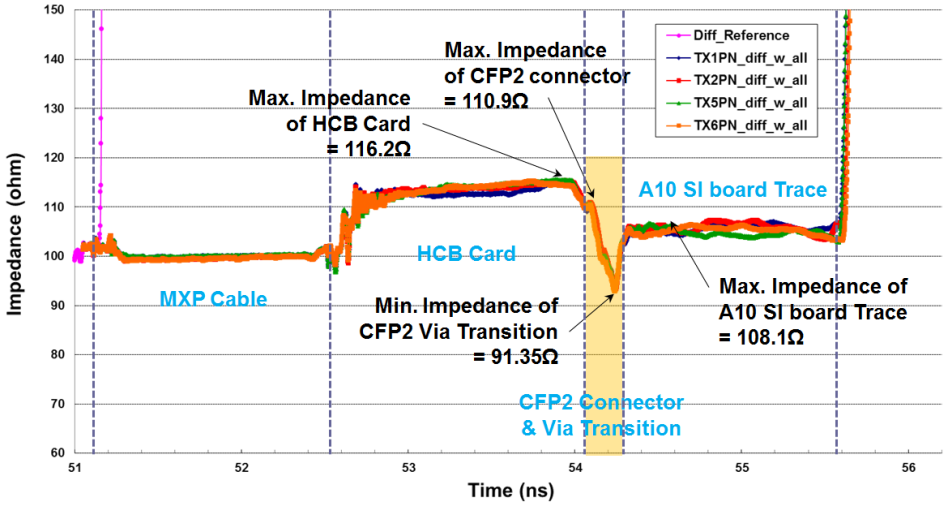

Figure 74. Insertion and Return Loss Performances of the Intel Arria 10 device CFP2 Channel

This figure uses the example in Figure 72 on page 54 with bare host board only.

Measured SP performances for one single TX pair and one single RX pair which both meet the CFP2 specification.

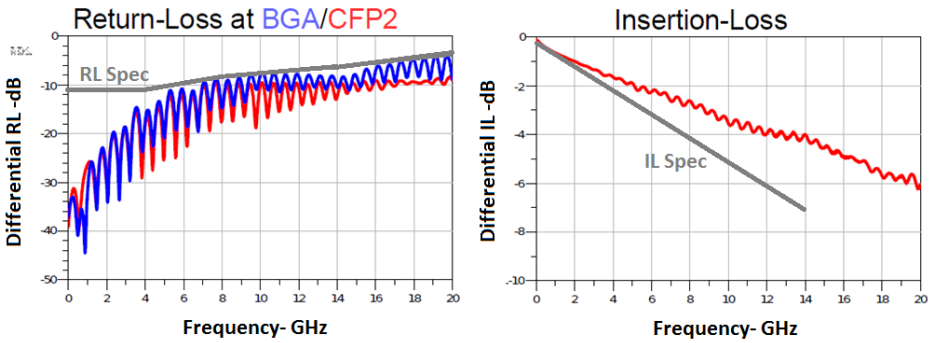

QSFP+/zSFP/QSFP28 Connector Board Layout Design Guideline

QSFP+ Module Assembly and Pinout

The Quad Small Form-factor Pluggable (QSFP) specification is based on the SFF-8665, SFF-8679 and OIF CEI v3.1 standards.

Figure 75. Simplified QSFP+ Channel

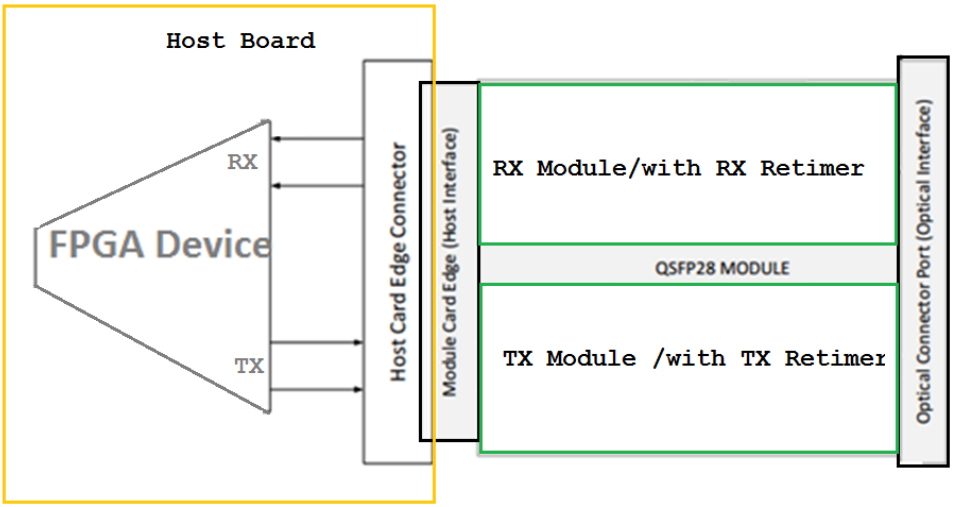

The roadmap for QSFP+ standards operating up to 28 Gbps includes:

- 10 Gbps QSFP10 SFF-8635
- 14 Gbps QSFP14 SFF-8685
- 28 Gbps QSFP28 SFF-8665

Figure 76. QSFP+ Module and Pin Map Layout

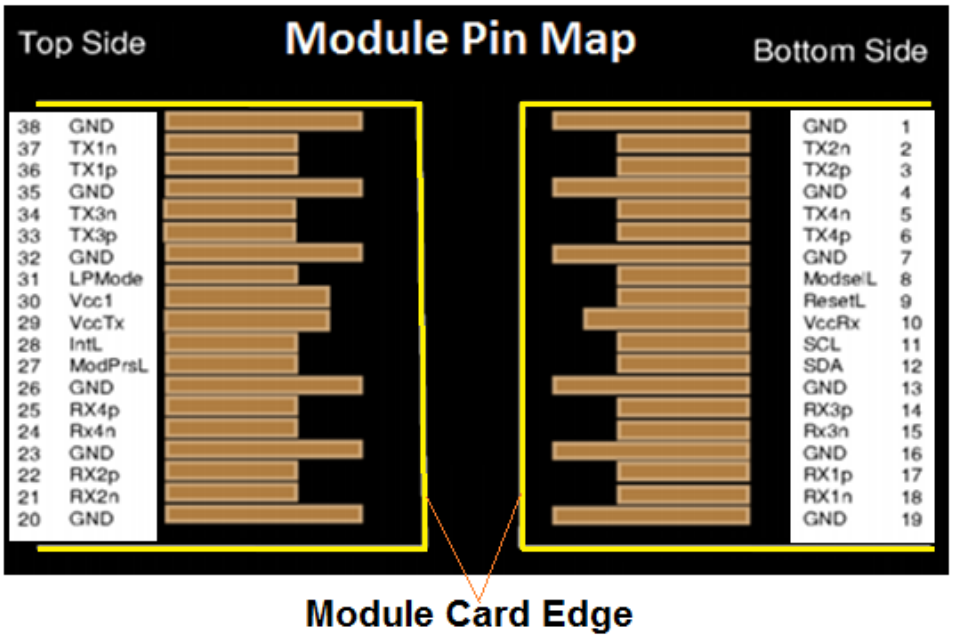

The RX(n)(p/n) and TX(n)(p/n) are module receiver data outputs and transmitter data inputs. They are all AC-coupled 100 Ω differential lines that should be terminated with 100 Ω differentially at the Host ASIC (SerDes). The AC coupling is inside the module and not required on the host board. For operation at 28 Gbps the relevant

standards (OIF-CEI-03.1 standard document) define the signal requirements on the high-speed differential lines. For operation at lower rates, refer to the appropriate standards.

Recommended PCB Design Guideline for QSFP+/zQSFP/QSFP28 Channels

Figure 77. Recommended PCB Layout at the QSFP+ Connector

Figure also applies to zQSFP and QSFP28 connectors on a PCB (host) operating up to 28 Gbps.

The yellow portions indicate GND2/GND4 cutout dimensions under the connector pads on the 1st and 2nd reference planes, respectively.

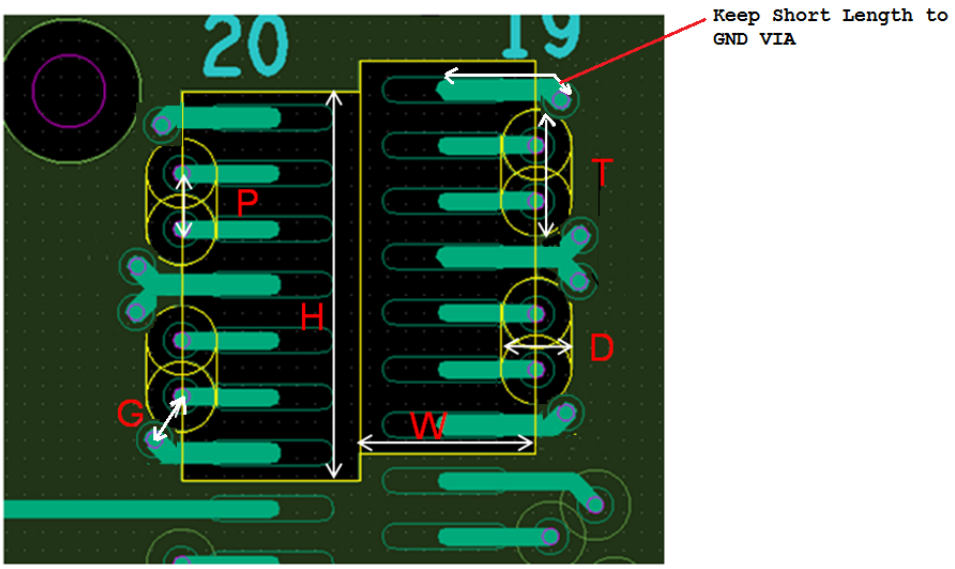

- Rectangular (W x H) cutout on the GN02 (2) layer (recommended for data rates up to 17 Gbps) under the QSFP+ pads. For data rates up to 28 Gbps, add one additional GND layer cutout (GND04). W = 100 mil, H = 210 mil.
- Signal and GND vias, finished drill diameter = 10 mil, Via pad diameter = 20 mil
  - All signal vias are back-drilled
  - Each signal via must have one single GND via
  - If the stack-up height is less than 1:12, you can also use an 8-mil finished via drill and an 18 mil via pad instead.
- P (signal-to-signal via pitch) = 40 mil
- Signal anti-pads: T = 90 mil, D = 45 mil
- G (signal-to-GND via pitch) = 30 mil
- 95 Ω differential PCB routing

Intel recommends that you have both the signal via and GND via located close enough to the connectors' signal and GND pads, respectively, to avoid cavity resonance at higher frequencies.

(2) GN02/GND04 represents the first and second GND layers below the connector.

Recommended QSFP+ Signal Routing

Figure 78. Recommended Signal Break-out Routing at the QSFP+ Connector

Different routing colors indicate different routing layers.

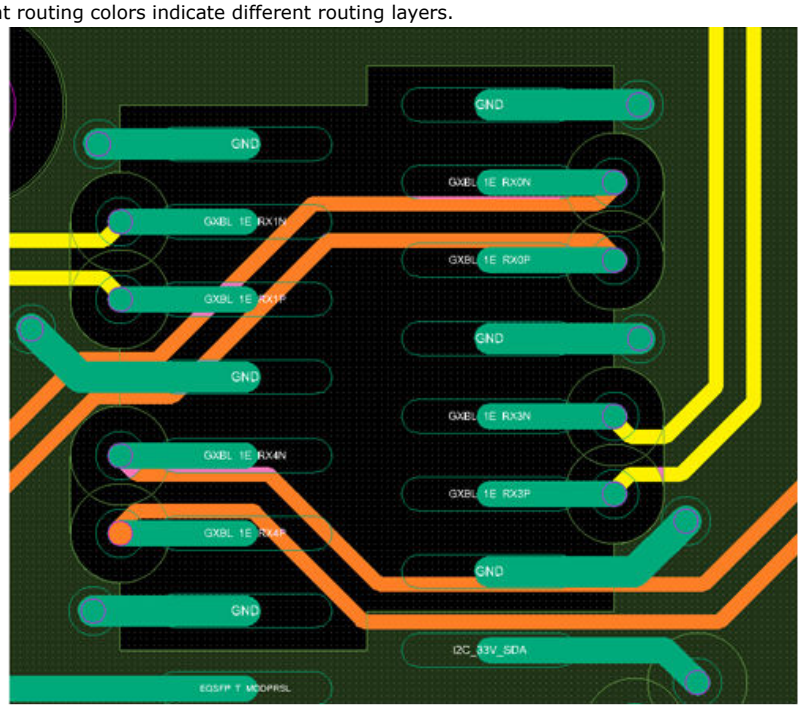

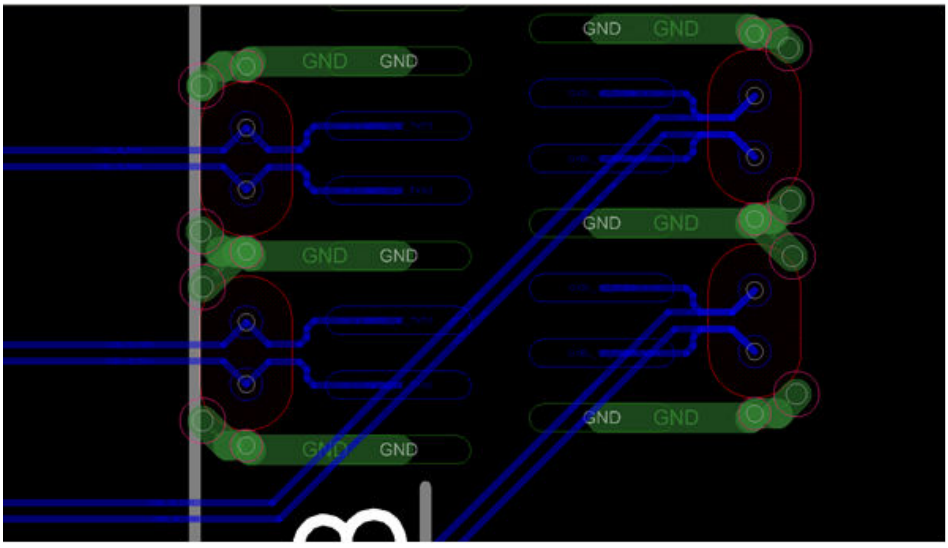

You must ensure that you have a proper GND reference plane for signal routing. The differential lanes in yellow can be routed on any signal layer. The differential lanes in orange can only be routed on any signal layer in which the GND reference layers are not GND02/GND04. This is due to crossing the cutout area below the connector.

The orange differential lanes can be routed from the opposite direction (similar to the yellow lanes) provided that there is adequate space for routing. In this case, all signal layers can be used for signal routing.

Observe these guidelines for improved QSFP+ performance at 28 Gbps on the main channel:

- Length matching for each pair (between P and N lanes) is required. Both P and N lanes must be in phase to recover the data. The skew matching in a pair is 2 ps.
- Length matching between pairs is not required unless specified by a designer.
- For optimized FPGA break-out layout design, refer to *FPGA Fan-out Region* chapter.
- Always use the minimum routing length from the FPGA to the connector to minimize insertion loss. Refer to *PCS Stackup Selection Guideline* chapter for stack-up and material selection, and *Recommendations for High Speed Signal PCB Routing* chapter for HSSI PCB routing.
- The insertion and return loss of the channel must meet specifications. Refer to *Electrical Specifications* chapter.

Related Information

- PCB Stackup Selection Guideline on page 6
- Recommendations for High Speed Signal PCB Routing on page 9
- FPGA Fan-out Region Design on page 10
- Electrical Specifications on page 87

QSFP28 Example Design and Performance Optimization

This section describes the optimization performed at the connector for the best insertion and return loss and crosstalk (isolation) performances of the channel.

Figure 79. QSFP28 Layout Design at the Connector with One Layer GND Cutout Below the Connector Pads (W(mil) x L(mil))

The performance results here are achieved by using a GND cutout at 50 mil x 65 mil, 67 mil x 75 mil, and the recommended 100 mil x 110 mil.

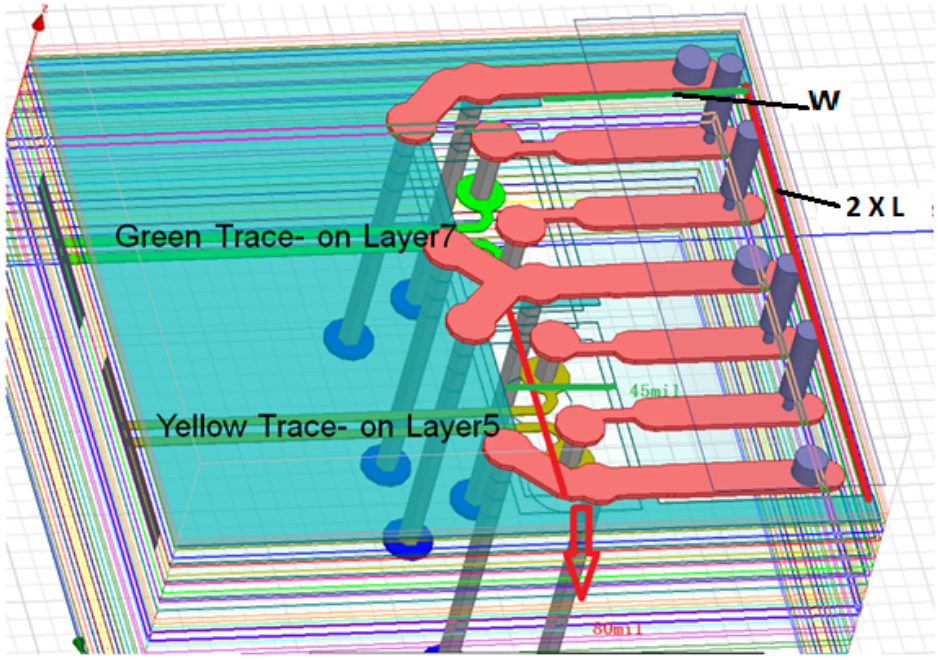

In the figure above, the signal finished drill size is 8 mil, the via pad is 18 mil, and the rectangular differential via-anti-pad is 45 mil x 80 mil. The signal via to signal via pitch in one pair is 40 mil. The signal via to GND via pitch is 27 mil. Routing is on layer 5. All signal vias are back-drilled up to layer 6.

The reference design is considered without the GND cutout below the connector pads (blue reference design in Figure 80 on page 61). The performance in Figure 79 on page 60 is observed by using various GND cutout (W) mil x (L) mil below each transceiver pair. W is swept from 50 mil to 100 mil and L from 65 mil to 110 mil. The GND cutout had been applied only on GND layer#2 as shown in Figure 79 on page 60

In the figure above, the differential return loss performance is from the stripline trace ports.

Figure 80. Differential Insertion Loss Performance

The blue reference design has no GND cutout.

The green design features a GND cutout at 50 mil x 65 mil.

The brown design features a GND cutout at 67 mil x 75 mil.

Intel recommends the red design which features a GND cutout at 100 mil x 110 mil.

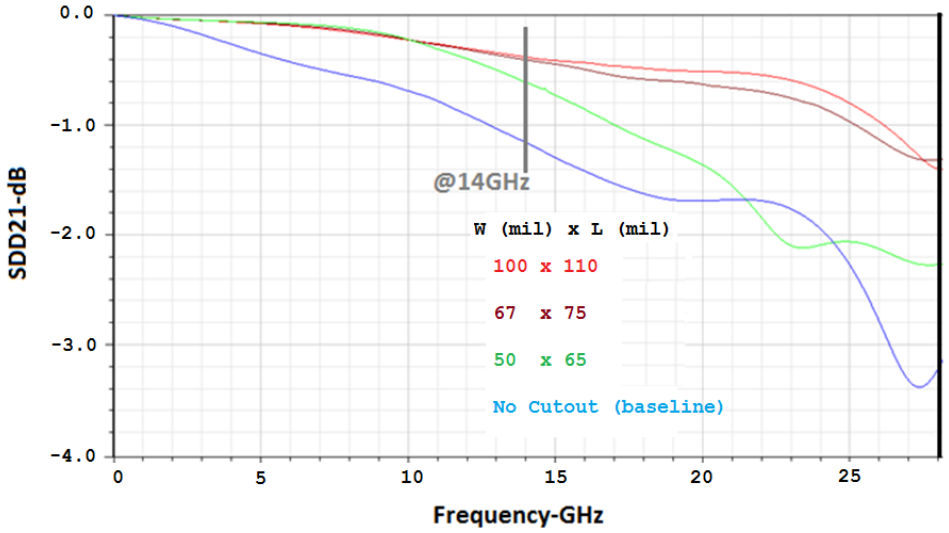

Figure 81. Differential Return Loss Performance

The blue reference design has no GND cutout.

The green design features a GND cutout at 50 mil x 65 mil.

The brown design features a GND cutout at 67 mil x 75 mil.

The red design features a GND cutout at 100 mil x 110 mil.

Intel recommends the red design which features a GND cutout at 100 mil x 110 mil.

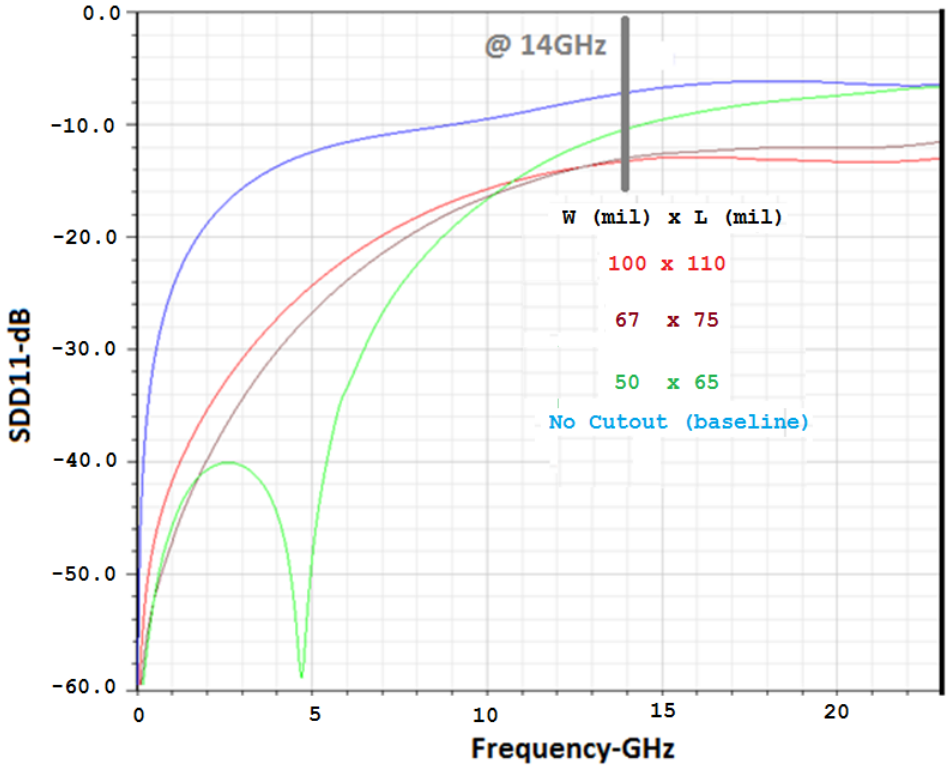

The figures above indicate that the recommended GND cutout has the best performance compared to all the others.

To further optimize the layout structure in Figure 79 on page 60, two GND reference diving boards for stripline trace routing have been added.

Figure 82. QSFP28 Layout Design at the Connector with Added GND/Reference Diving Boards

Reference diving board structures are on layers 4 and 6. The diving board width is twice the width of the outer edge-to-edge of the differential lanes. The length of the diving board is extended up towards the signal via.

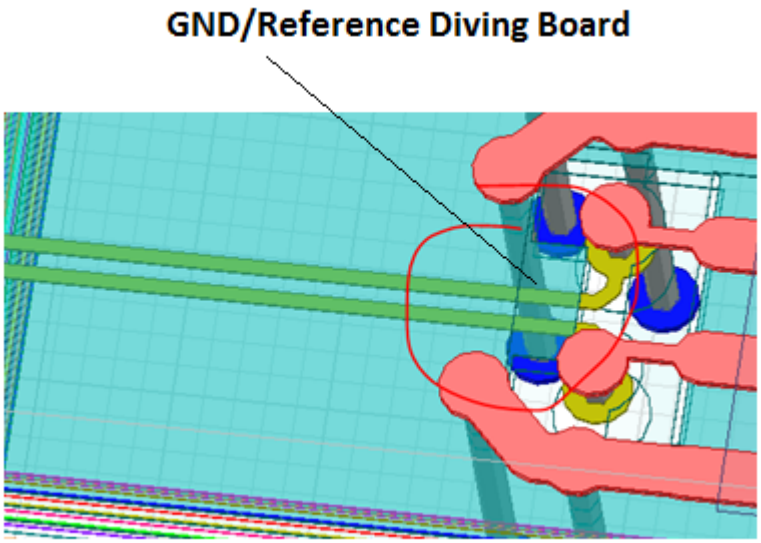

Figure 83. Differential Insertion Loss Performance

GND cutout below the connector is fixed at 67 mil x 75 mil.

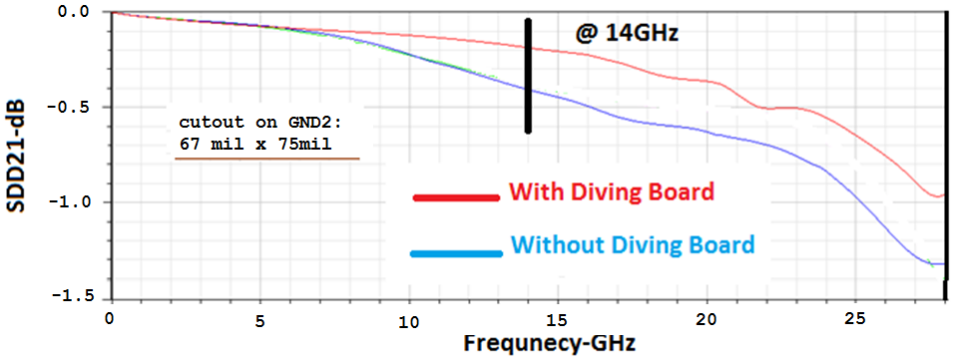

Figure 84. Return Loss Performance

GND cutout below the connector is fixed at 67 mil x 75 mil.

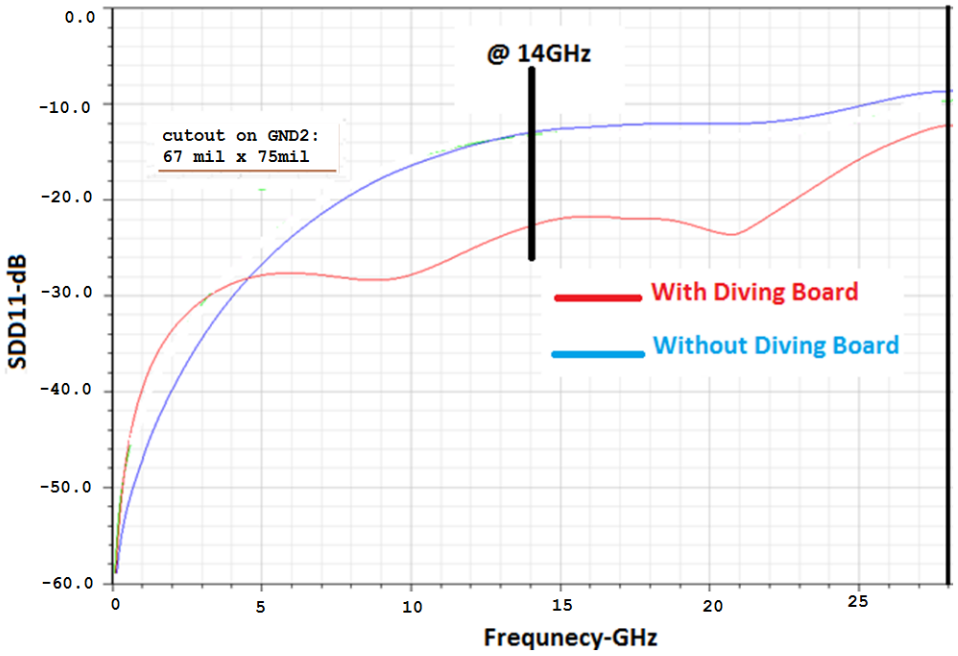

The addition of reference diving boards enhances both insertion and return loss performance.

To improve isolation between the two pairs, you can insert an additional GND via between the signal pairs. Most of the crosstalk will occur through vertical coupling between signal vias of two different signal pairs. Therefore, adding one GND via will enhance the isolation.

Figure 85. QSFP28 Layout Design with Added GND Vias

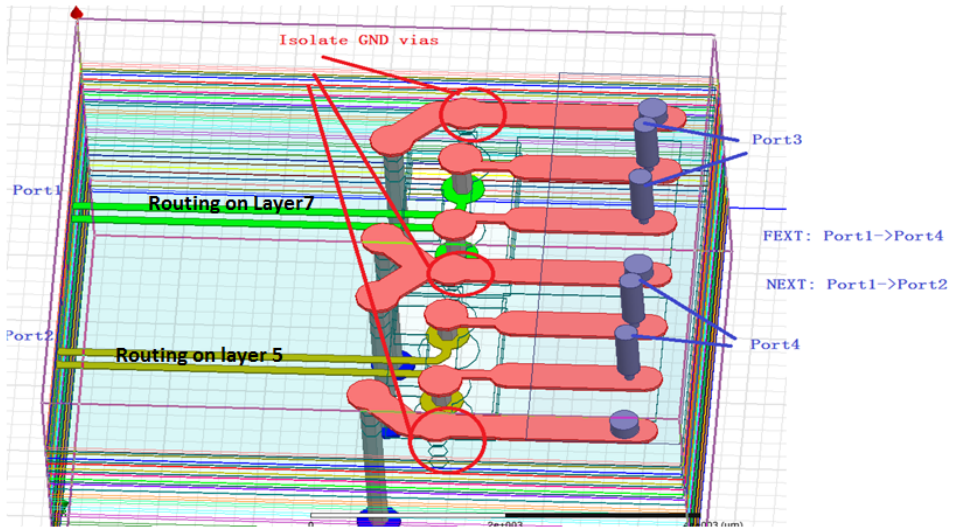

Figure 86. FEXT and NEXT Crosstalk Performance

GND cutout below the connector is fixed at 67 mil x 75 mil.

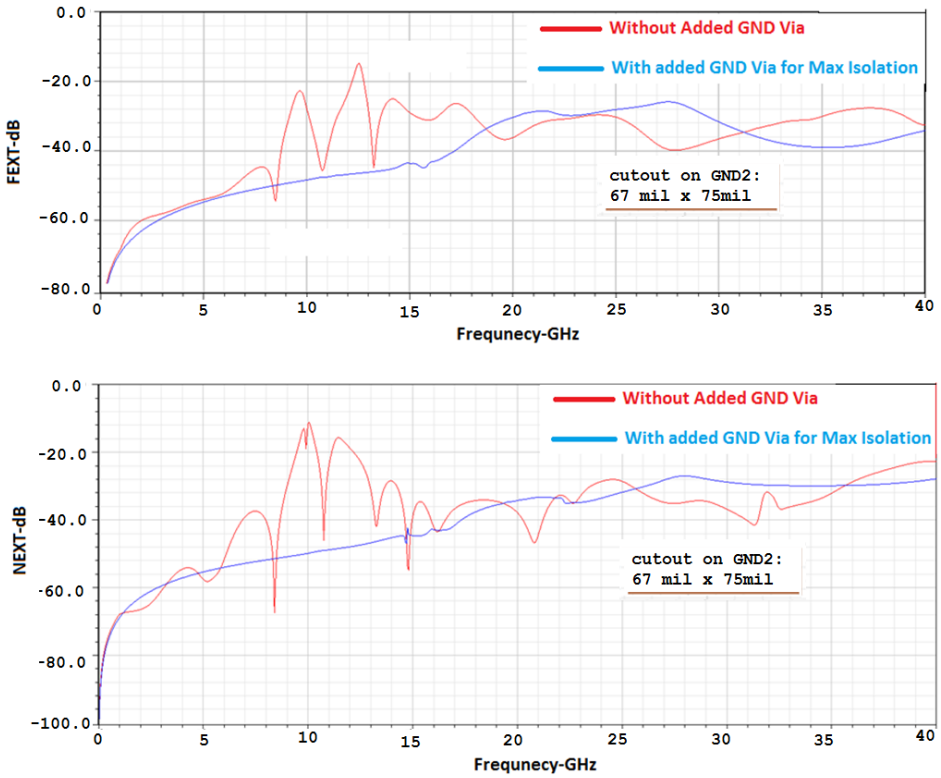

SMA 2.4-mm Layout Design Guideline

SMA 2.4 mm Molex® Connector Assembly

Figures below show the structure of SMA/Interface 2.4 mm (MOLEX_73387) for two different versions (A and B). The operating frequency from DC is up to 50 GHz, which is within the range of 28 Gbps applications.

Figure 87. Version-A Molex 2.4 mm SMA connector

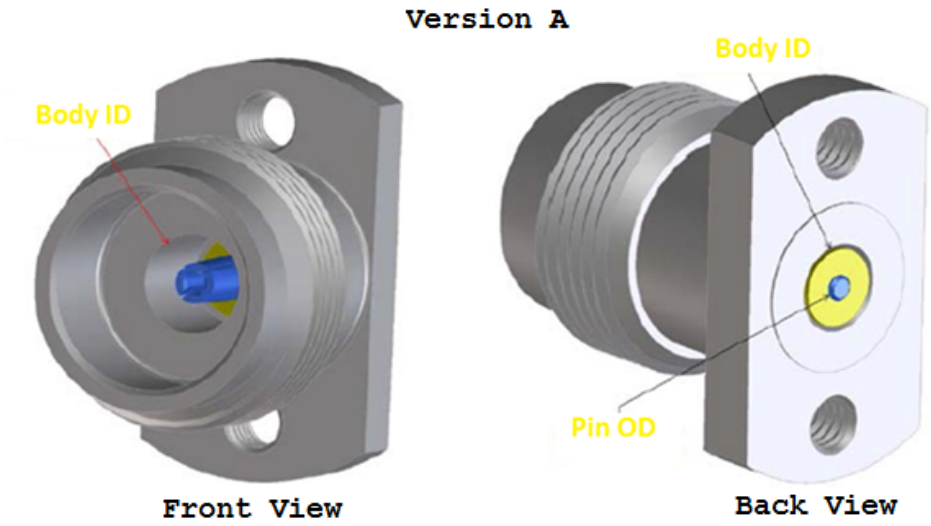

Figure 88. Version-B Molex 2.4 mm SMA connector

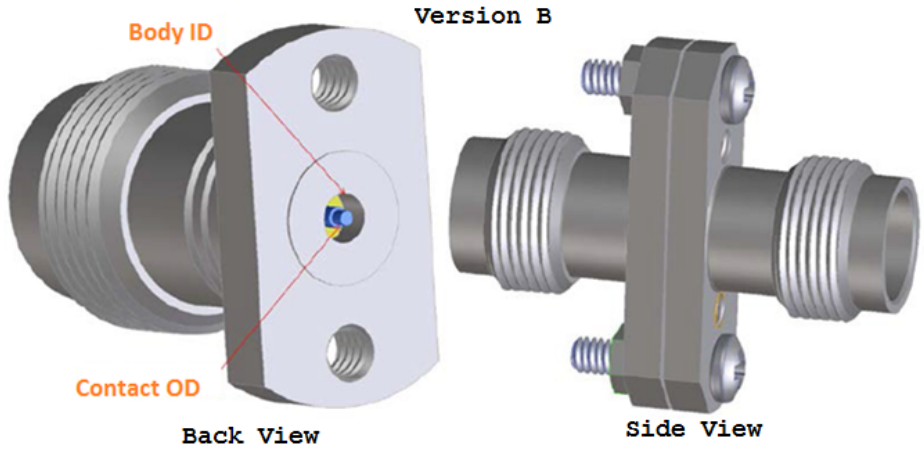

PCB Design Guidelines for Channels Using the 2.4 mm Connector

Observe the following guidelines when designing for channels using the 2.4 mm connector:

- Refer to *PCB Stackup Selection Guideline* chapter for the selection of stackup and routing layers.
- Intel recommends a routing trace impedance of 95Ω loosely differential, or 47.5 Ω single-ended. Refer to *FPGA Fan-out Region* chapter for break-out routing at the FPGA.
- Use the minimum routing length possible to minimize insertion loss and crosstalk.
- Refer to the AC coupling layout design guide in *AC Coupling Capacitor Layout and Optimization Guidelines* chapter, because all RX paths require AC capacitors.

- Match the length (less than 2 ps) for all TX and RX paths if required. Refer to *Recommendations for High Speed Signal PCB Routing* chapter for the length matching strategies at the FPGA.
- Use a back-drill for all transceiver signal vias.
- The Molex connector and cutout are standard recommendations made by Molex. This is a surface-mounted connector, and there is always a back-drill for the signal vias for transferring signals from the top layer to the inner layers.

Related Information

- PCB Stackup Selection Guideline on page 6
- Recommendations for High Speed Signal PCB Routing on page 9
- FPGA Fan-out Region Design on page 10
- AC Coupling Capacitor Layout and Optimization Guidelines on page 30

Recommended PCB layout Design for Version-A SMA 2.4 mm Connector

Figure 89. Recommended layout for Molex Version-A 2.4 mm SMA connector

The SMA is surface-mounted on the top layer.

The diameter of the signal via drill hole is 15 mil.

The diameter of the outer GND circle on the top layer is approximately 180 mil.

D1: signal anti-pad = 60 mil diameter.

D2: signal anti-pad = 26 mil diameter.

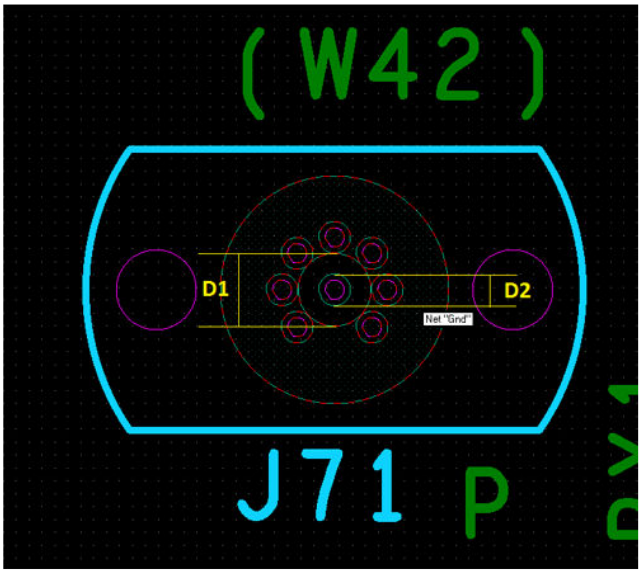

Top Layer

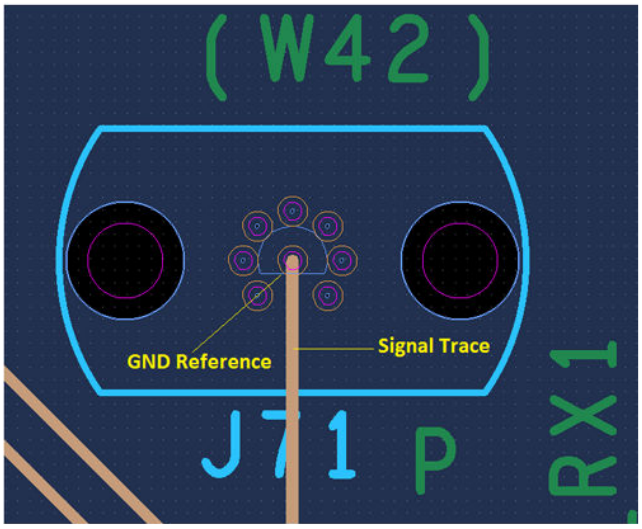

GND Reference

The GND reference planes for SMA signal routing are extended to the signal via pad to prevent mismatching in this area.

The insertion loss of the connector only is calculated as:

```
Connector RF Insertion Loss (max) = 0.03 x √f (GHz) dB
```

f is the maximum operating frequency for the channel. For example, for a channel operating at 28 Gbps, the maximum RF insertion loss of an individual SMA connector is 0.12 dB.

Recommended PCB layout Design for Version-B SMA 2.4 mm Connector

An optimized layout suppresses the natural cavity mode within the via GND ring as well as additional cavity coupling to other structures on the PCB.

Figure 90. Version-B 2.4 mm Connector Launch

The diameter of the signal via drill hole is 10 mil.

D1: diameter of the signal anti-pad on the top layer = 60 mil.

D2: diameter of the signal pad = 20 mil.

D3: diameter of the signal anti-pad on the inner layers = 40 mil.

D4: diameter selected = 118 mil.

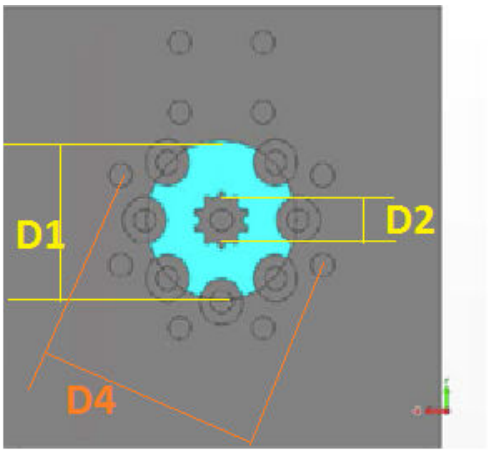

Top layer of the launch structure showing the via top

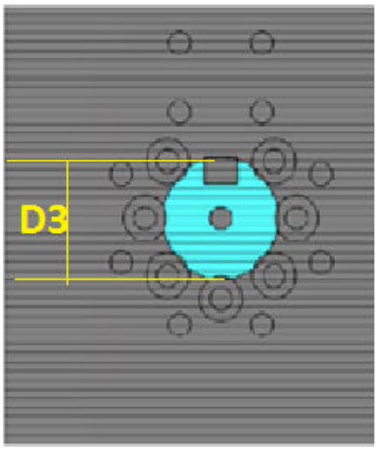

Plane layer above and below the trace

The width of the reference diving board must be at least twice the width of the signal trace.

Figure 91. Magnified Signal Trace and Reference Plane Areas at the Connector

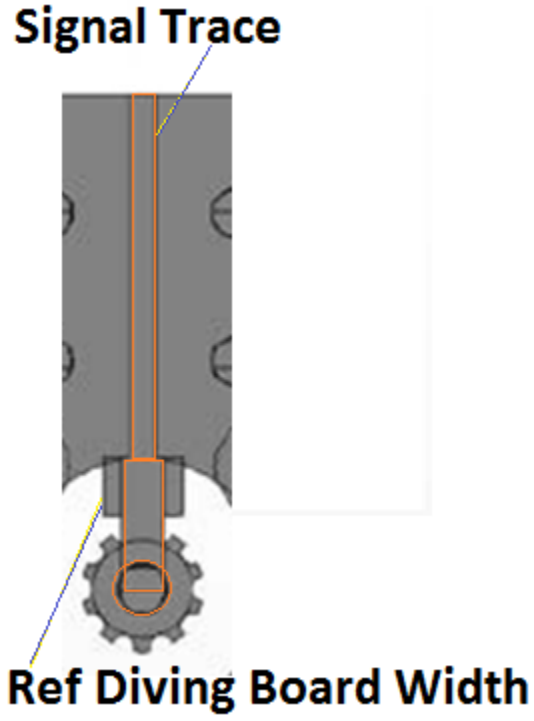

Performance Comparison between Option1 and Option2 layout

A version-B SMA connector model in figure *Version-B Molex 2.4 mm SMA connector* has been implemented on both layout design recommendations Option1 and Option2 in figure *Recommended layout for Molex Version-A 2.4 mm SMA connector* and figure *Version-B 2.4 mm Connector Launch* respectively. Only the Connector model and a small portion of Host PCB has been included into this simulation and performance comparison. Both simulations are using similar stack up materials and the main PCB routing is on layer 7. There is slight difference of about 6 mil in signal via height (Signal via height in Option 1 layout is about 6 mil more than that in Option 2). From Insertion loss point of via, there is about 0.18 dB difference at 14 GHz as shown in figure *Single ended Insertion loss comparison for Version-B SMA connector on both recommended Option1 and Option2 layout design. (Improved by 0.18 dB at 14 GHz)* .

Figure 92. Single ended Insertion loss comparison for Version-B SMA connector on both recommended Option1 and Option2 layout design. (Improved by 0.18 dB at 14 GHz)

Option 1 = Recommended PCB layout Design for Version-A SMA 2.4 mm Connector.

Option 2 = Recommended PCB layout Design for Version-B SMA 2.4 mm Connector.

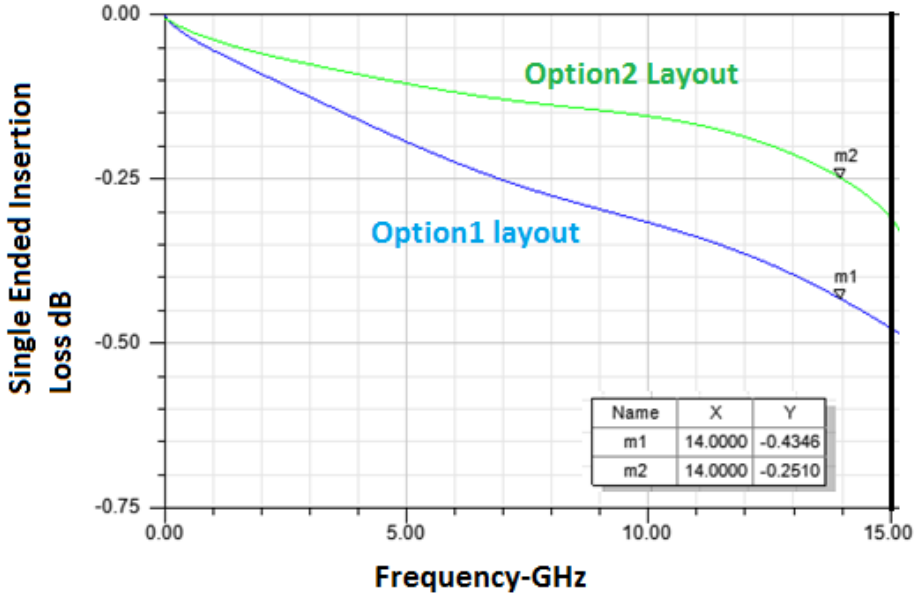

Insertion loss from the point of the via improves by approximately 0.18 dB at 14 GHz.

Figure 93. Single Ended Return loss (from connector port) comparison for Version-B SMA connector on both recommended Option1 and Option2 layout design. (Improved by 4.5 dB at 14 GHz)

Option 1 = Recommended PCB layout Design for Version-A SMA 2.4 mm Connector.

Option 2 = Recommended PCB layout Design for Version-B SMA 2.4 mm Connector.

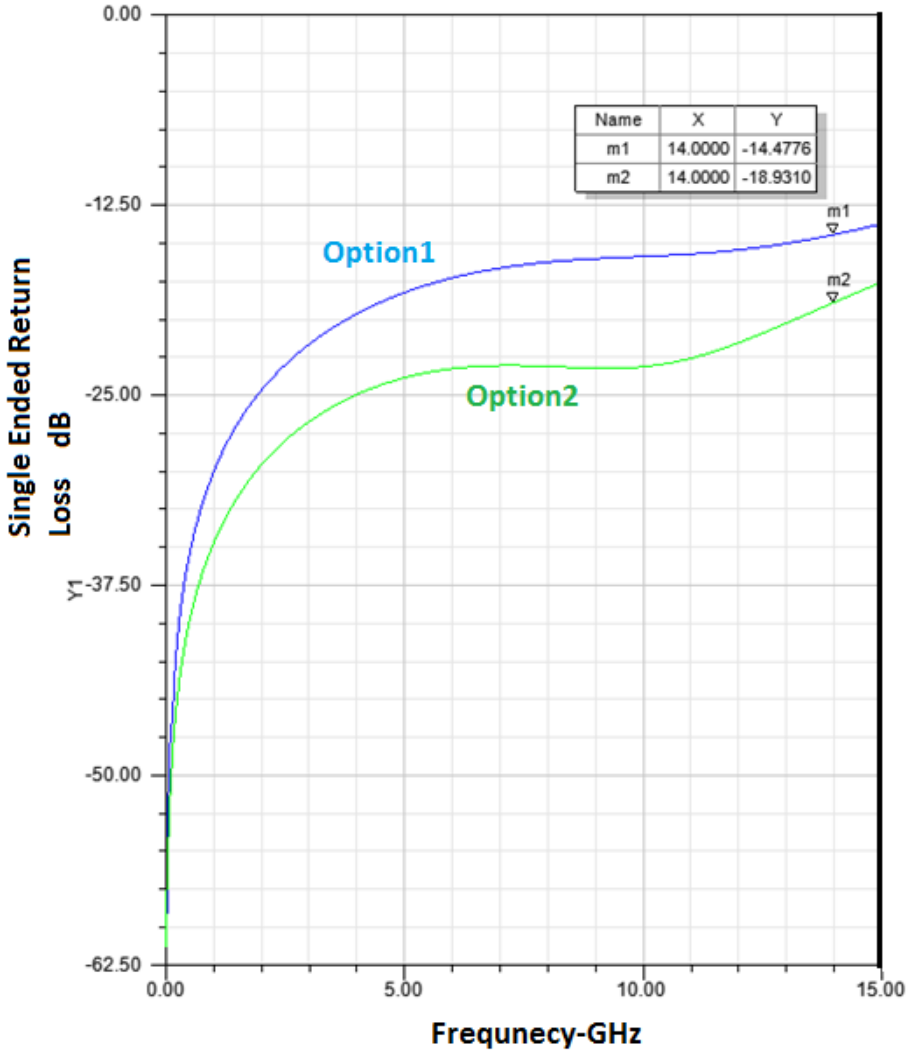

Return loss from the point of the connector port improves by 4.5 dB at 14 GHz.

Other 2.4 mm Connectors

Make sure you follow your connector vendor's recommended layout on the PCB at the connector. Intel recommends using a large signal anti-pad on the reference GND plane. Additionally, you should ensure that you have a circle of GND vias and a GND reference diving board.

Figure 94. GND Cutout and Reference Diving Board

Figure shows the first GND layer below the microstrip routing.

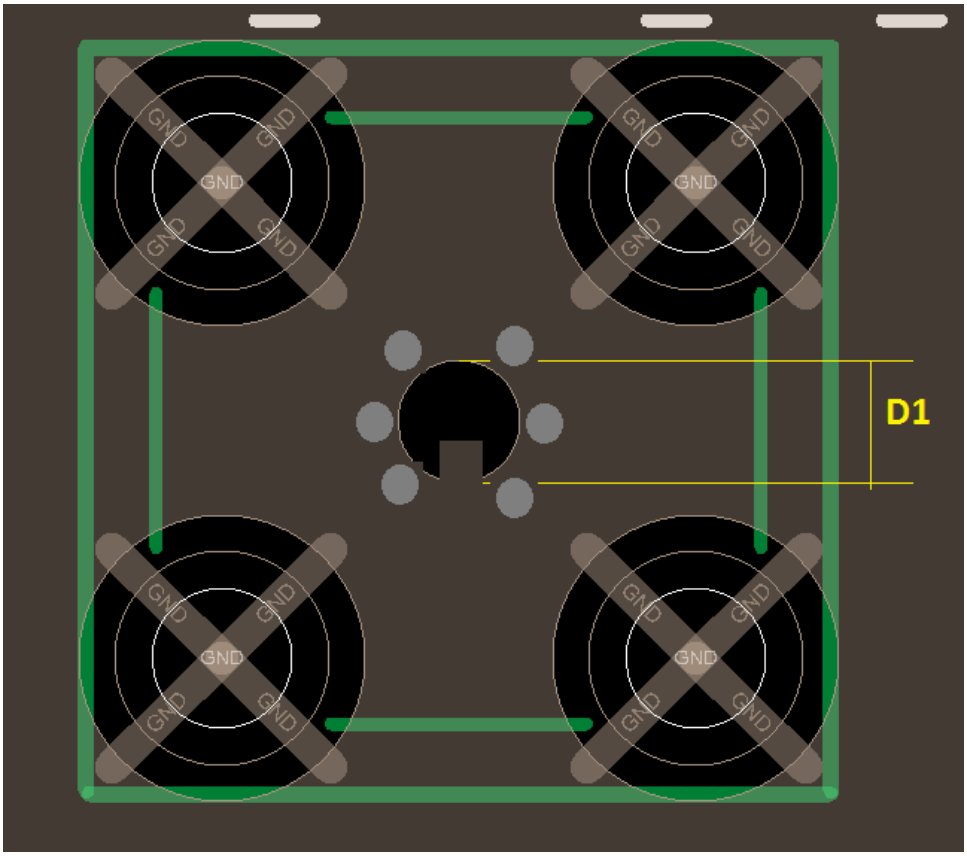

2.4 mm Channel Specifications

There is no defined interface or specification available for 2.4 mm channels up to 28 Gbps. The layout design strategy above at the connector is based on the minimum reflection obtained from all discontinuities in the channel. The material selection of stackup and total length of signal routing will define the least insertion loss achievable at 14 GHz.

2.4 mm Example Design Performance

The layout design strategy in the previous section showed the implementation at the connector area.

Figure 95. Arria10 device SI Development Kit Channel Layout for the TX0 Lane

The total trace length is approximately 4.5 inch excluding the connector.

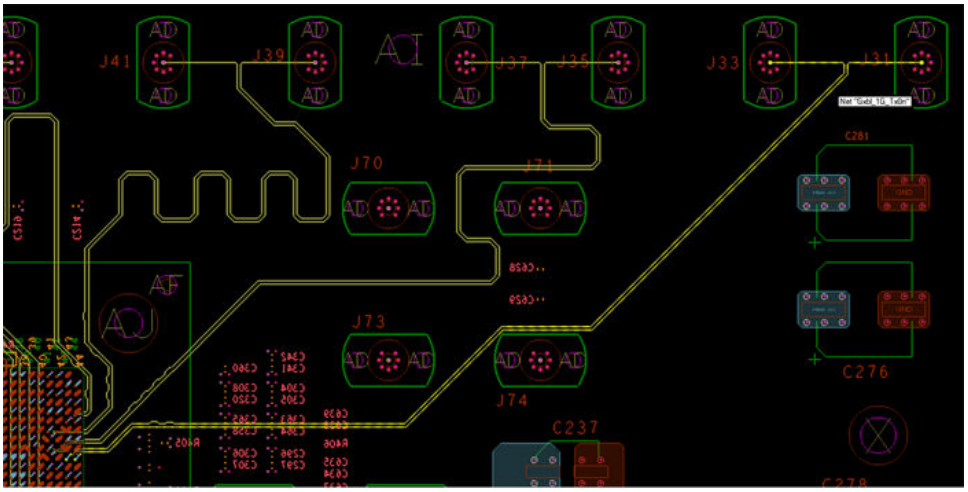

Figure 96. Magnified Single-ended TDR measurement from the Connector

The TDR rise time used for below measurement < 17 ps.

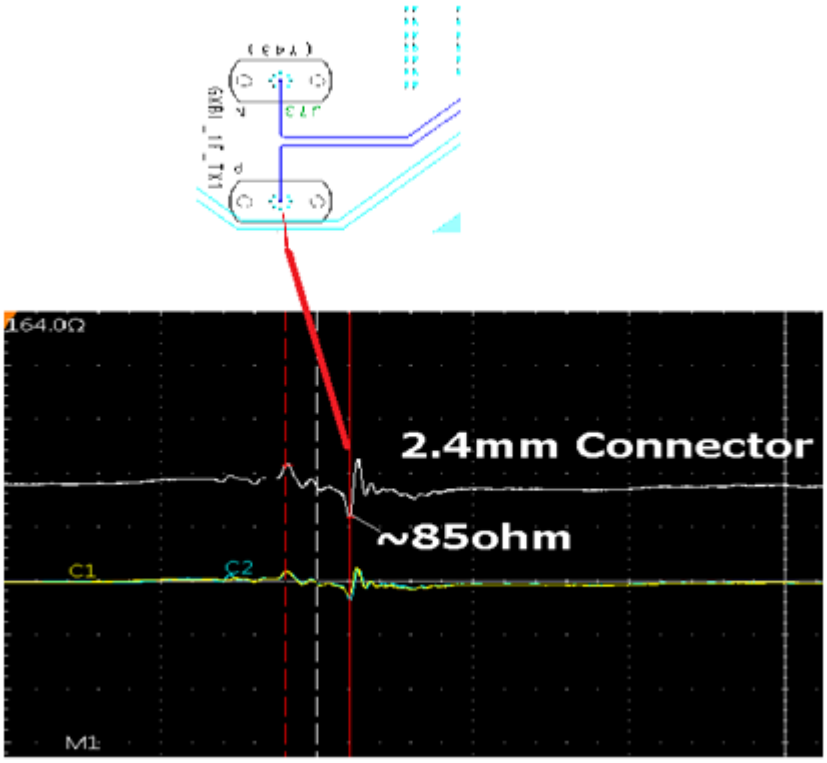

The signal via impedance reaches 45 Ω at the minimum and to 55 Ω by transition from signal via to the main inner layer routing.

Figure 97. Single-ended Scatter Parameters S21 and S22 (from the Connector)

The single-ended insertion loss is approximately 5 dB and the return loss is approximately -12 dB at 14 GHz.

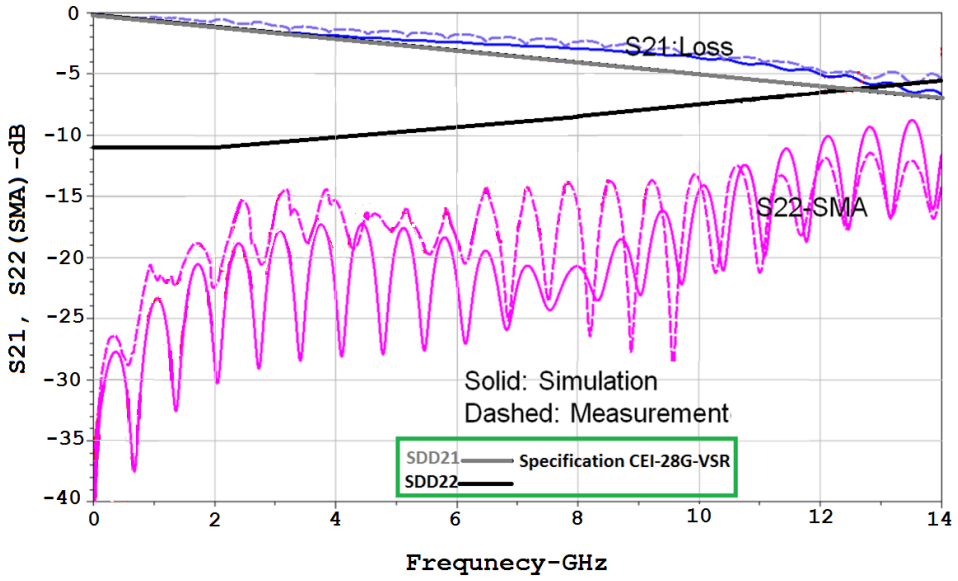

Tyco/Amphenol Interlaken Connector Design Guideline

Connector Signal and Pin Assignment

Intel recommends the TE Connectivity® 3.9 mm STRADA Whisper connector (R/A receptacle) or IMPACT connector for the 25+ Interlaken interface. This connector is ideal for cabled connections. The part number is TYCO 2187194-1. This connector can handle both TX and RX signals. The cable assembly crosses the signal to provide connectivity from TX to RX and from RX to TX. A single connector supports up to 24 channels (24 TX and 24 RX). The advantage of this connector is that it can operate up to 57.8 Gbps with no skew in pair.

Figure 98. Interlaken Single or Dual Connector Setups

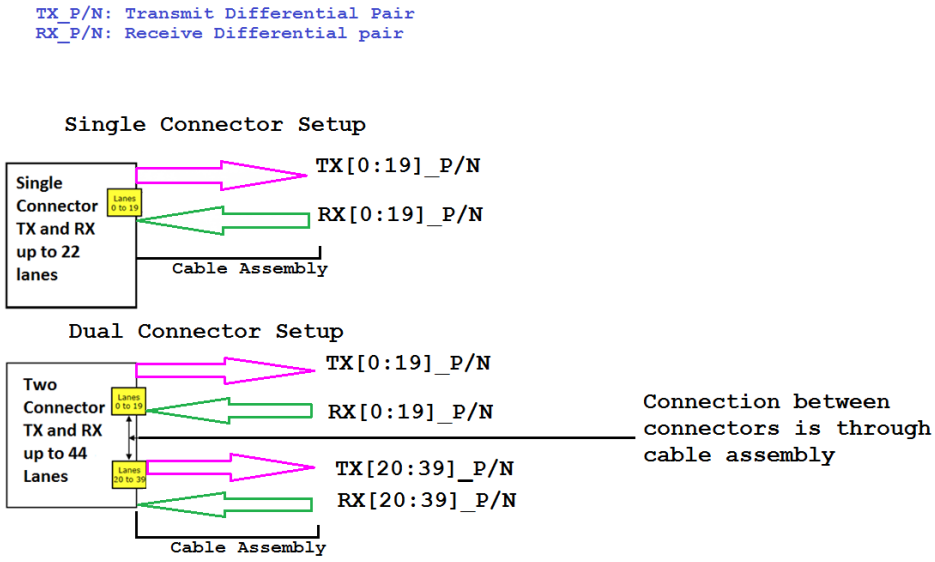

The single connector allows up to 22 lanes of high-speed differential connectivity to another board for interoperability. The Interlaken Alliance currently specifies only 20 lanes are needed for interoperability.

The dual connector setup allows up to 44 lanes of high-speed differential connectivity to another board. In this case, the Interlaken Alliance currently specifies the need for only 40 lanes for interoperability. The connector and cable assembly are rated to support data rates up to 28.9 Gbps with non-return to zero (NRZ) encoding.

Note: TX[20:21] and RX[20:21] are spare lanes that you can use for other cases where extra lanes may be required.

Figure 99. Connector Signal Pin Assignment and Magnified View of Signal Pins

Pin assignments for the high-speed differential pairs. Signals labeled TX are outputs from the board, while RX signals are inputs to the board.

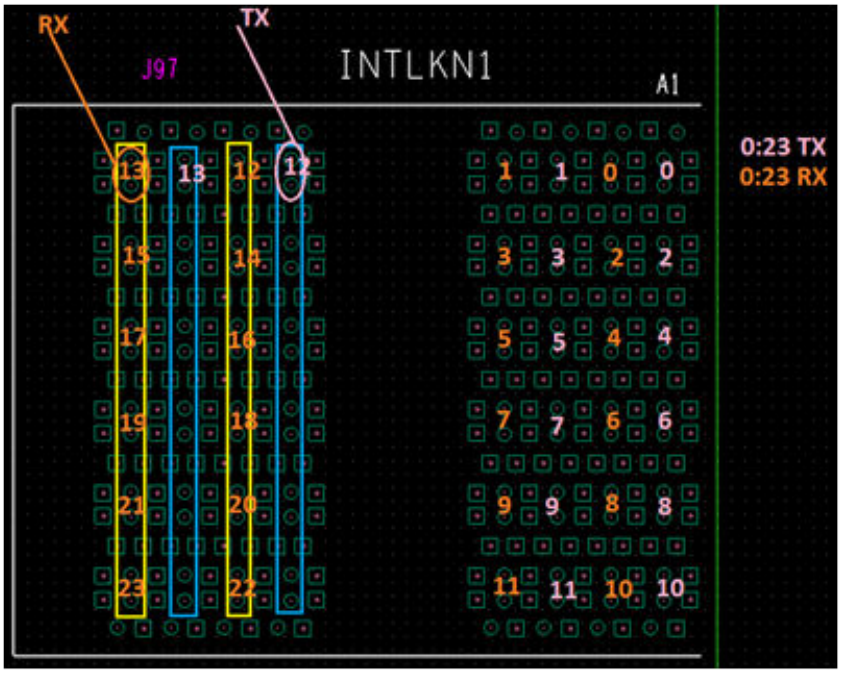

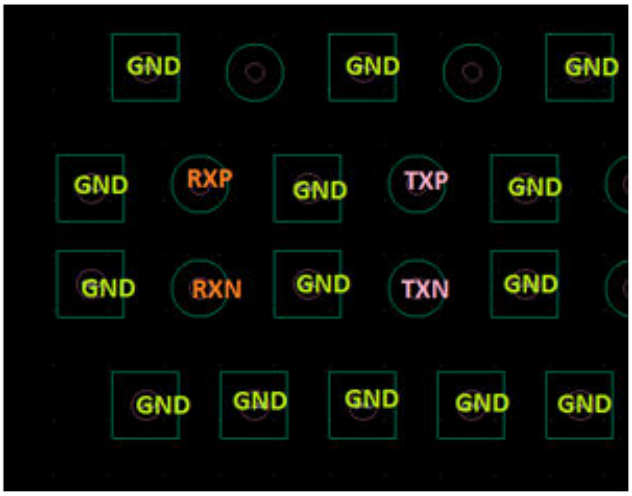

The cable assembly connecting each PCB's on-board connector(s) (STRADA Whisper connectors to R/A receptacles) takes care of swapping TX to RX. Therefore, the connector pinouts are identical on either inter operating board.

Note: TX_REFCK_P/N and RX_REFCK_PN are not provided for high-speed physical interoperability.

The insertion loss of the connector is less than 1 dB and is linear up to 20 GHz. For an 85 Ω channel, Intel recommends that you use the 85 Ω version of the connector. Likewise, for a 100 Ω channel, Intel recommends that you use the 100 Ω version of the connector. There is little to no skew through the high-speed differential pairs within the connector.

PCB Design Guidelines for Channels with 25 Gbps + Interlaken Interface Connector Recommendations

Figure 100. TYCO Interlaken Connector GND Plane Cutout on a Host PCB

H = 72.5 mil

P = 47.5 mil

L = 82.5 mil

W = 52.5 mil

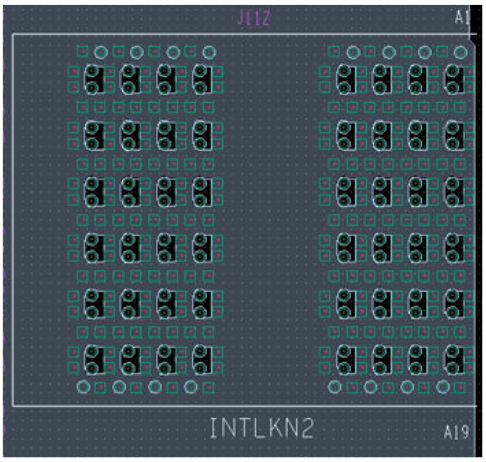

High-level View

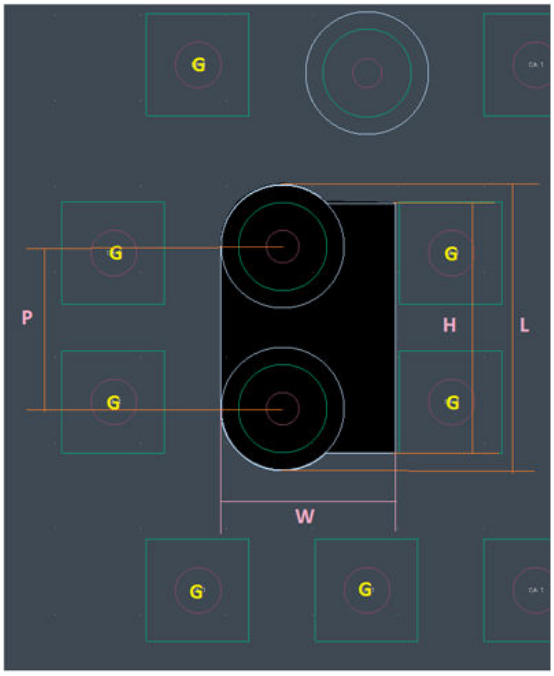

Magnified View

The maximum intra-pair skew (between P/N) is within 2 ps.

The GND reference cutout must be applied to all GND and reference planes. The dimensions of the GND cutout are 52.5 mil (W) x 72.5 mil (H).

The standard signal via used in Figure 100 on page 79 has the following characteristics:

- 10 mil drill size diameter
- 26 mil pad diameter
- 36 mil anti-pad diameter

You must remove all non-functional pads.

Signal routing is differential on the main PCB before it reaches the Interlaken connector where it is turned into a single-ended routing.

Figure 101. Option 1: TYCO Interlaken Connector Routing on a Host PCB with Skew Matching

50 Ω single-ended routing up to the connector edge.

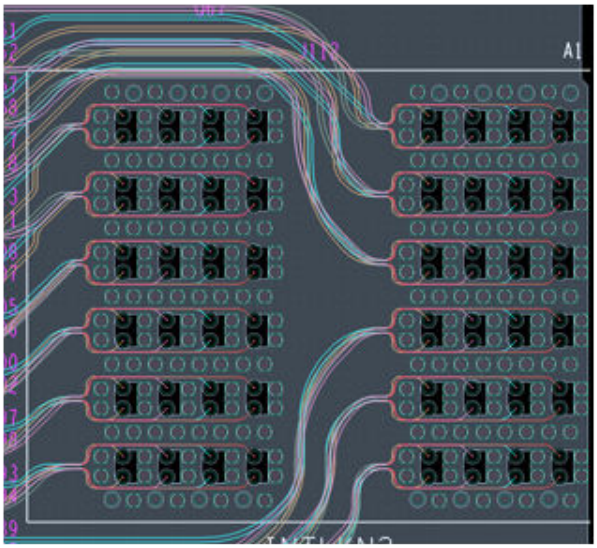

Overview

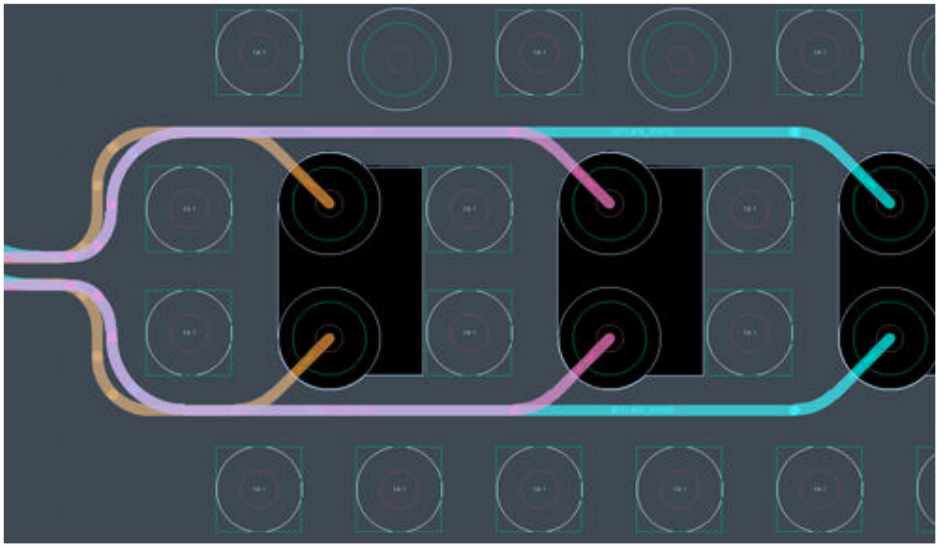

Magnified View

Figure 102. Option 2: Alternate TYCO Interlaken Connector Routing on a Host PCB with Skew Matching

Below are two figures. The first one describes the use of 4 layers for routing. The second figure describes the use of two layers for routing (one layer for TX and one layer for RX for maximum isolation).

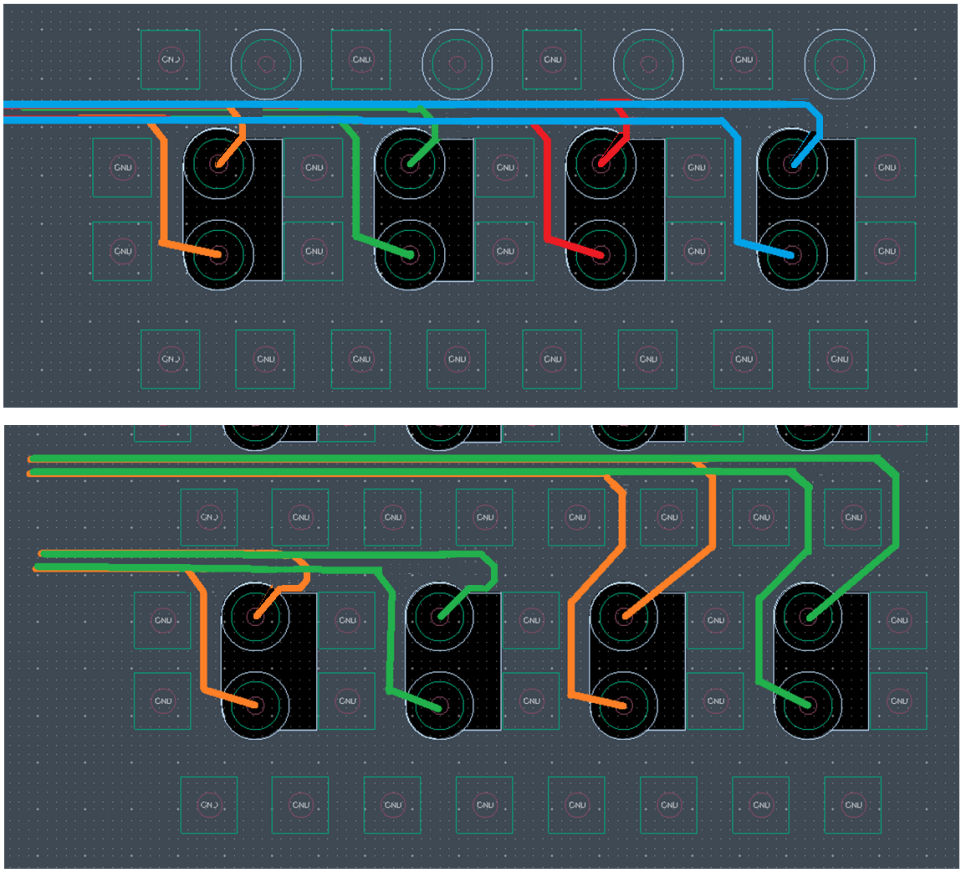

Figure 103. Interlaken Interface Connector Configurations

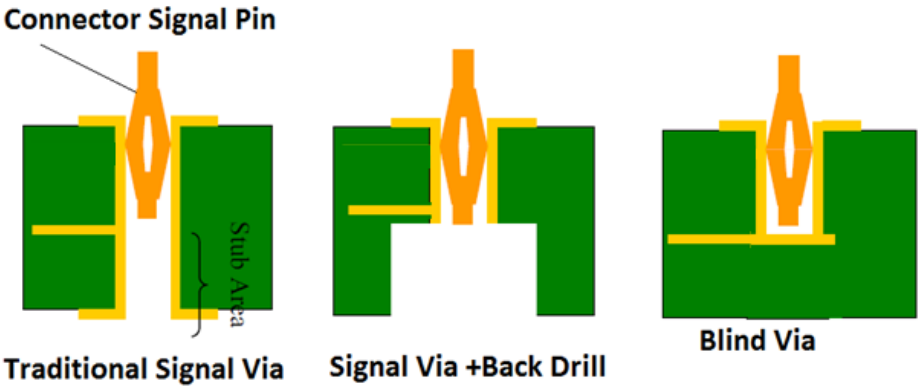

You can use either the signal via with back-drill, or a blind via at the Interlaken interface connector.

Figure 104. Recommended Routing Layer Due to the Connector Pin Length

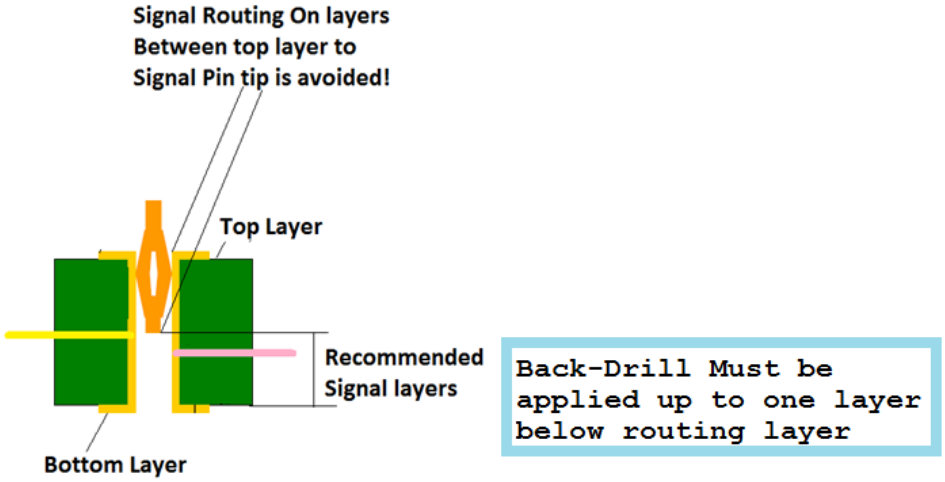

Signal routing is located on the bottom stack-up layer to minimize the impact of stubs in the channel. Observe the following recommendations for stack-up PCB design:

- Select the appropriate stack-up. Refer to *PCB Stackup Selection Guideline* chapter .
- Select the PCB routing layers:
  - Separate TX and RX layers for maximum isolation
  - The Interlaken interface connector pins have the following lengths:
    - Connector GND pin = 1.88 mm
    - Connector signal pin = 1.3 mm
  - Use layers below the signal pin tip for all signal routing. This helps to avoid the stub from the connector pin if the upper layers are used for routing. Use either the signal through via with back-drill or the blind via.
- Intel recommends using a 95 Ω routing trace impedance because it aligns with the following guidelines:
  - For FPGA break-out, refer to *FPGA Fan-out Region* chapter.
  - Use a 100 Ω loosely differential routing on the main host PCB if you are using option 1 in Figure 101 on page 81 at the connector.
  - Use a 100 Ω tightly differential routing on the main host PCB up to the connector pins if you are using option 2 in Figure 102 on page 82 at the connector.
- Use the smallest routing length possible to minimize insertion loss and crosstalk.

- Ensure that all RX paths have an AC capacitor for AC coupling. Refer to the AC coupling layout design guideline in *AC Coupling Capacitor Layout and Optimization Guidelines* chapter.
- Ensure that you have length matching (less than 2 ps) for all TX and RX paths if this is a requirement. Refer to *Recommendations for High Speed Signal PCB Routing* chapter for length matching strategies at the FPGA.
- Use a back-drill for all transceiver signal vias if a through via is used.

Related Information

- PCB Stackup Selection Guideline on page 6
- Recommendations for High Speed Signal PCB Routing on page 9
- FPGA Fan-out Region Design on page 10
- AC Coupling Capacitor Layout and Optimization Guidelines on page 30

Interlaken Channel Interface Performance Example

The channel in this example is designed for the Interlaken interface using a TYCO Interlaken connector. The board layout recommendations provided above are used in this channel design. This example design is implemented on the Intel Arria 10 device development kit.

A TX channel has been selected for these 3D HFSS simulations. The TX channel is routed on layer 26 using a back-drill up to layer 27. The total PCB routing is approximately 3.94 inch using stripline routing with rounded corners.

Figure 105. Stackup Layer and Material Data for the Example Interlaken Channel

These are the specifications for the example channel:

- 30 layers
- Copper Foil - HVLP
- Surface roughness = 2 µm
- Back-drill
- Material = Megtron6
- Total thickness = 153.3 mil

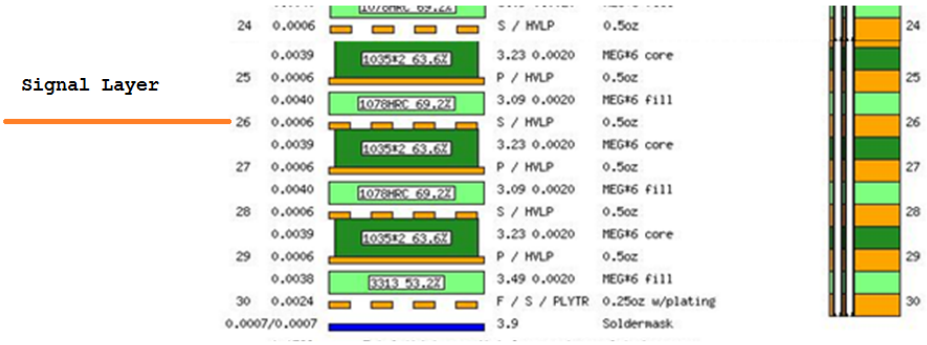

Figure 106. Interlaken TX Channel for Simulation

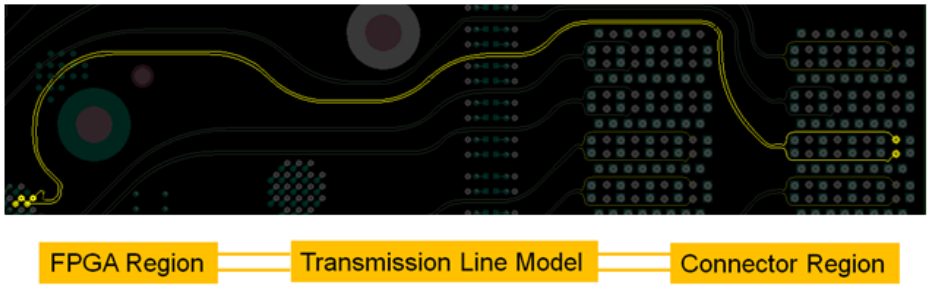

The following figures show the host PCB TX channel performance from the FPGA BGA/ball to the Interlaken connector signal pads on the top layer.

Figure 107. Differential Insertion Loss on the Host PCB Only

Results exclude the Interlaken connector.

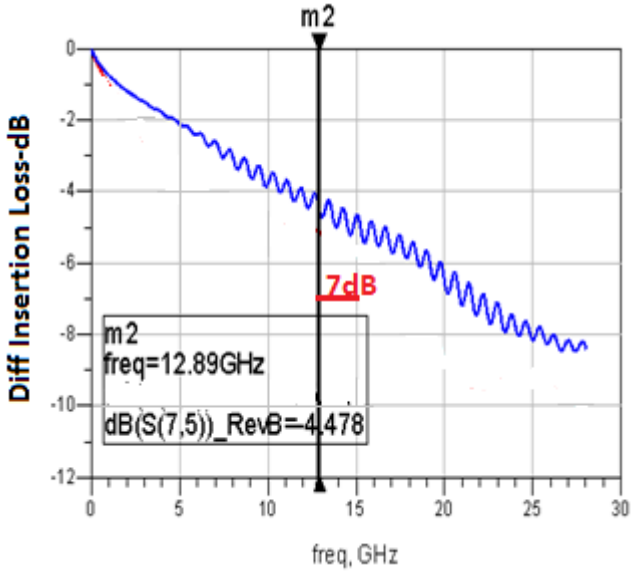

Figure 108. Differential Return Loss from the Interlaken Connector Pads on the PCB

Results exclude the Interlaken connector.

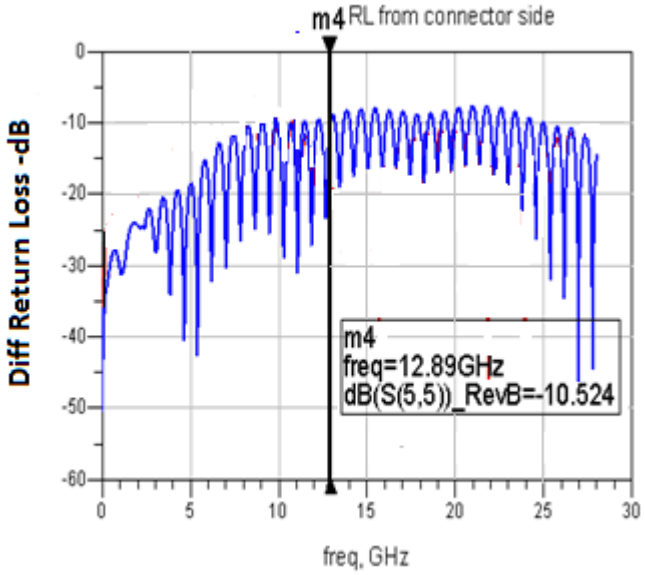

The insertion loss is below the mentioned specifications in *Electrical Specifications* chapter, which specifies less than 7 dB insertion loss for the host PCB.

Figure 109. Differential Measured TDR Performance for Various TX Interlaken Channels on an Arria 10 device Development Kit

Results show performance focused only on the PCB trace impedance and connector fan-out areas for various TX channels

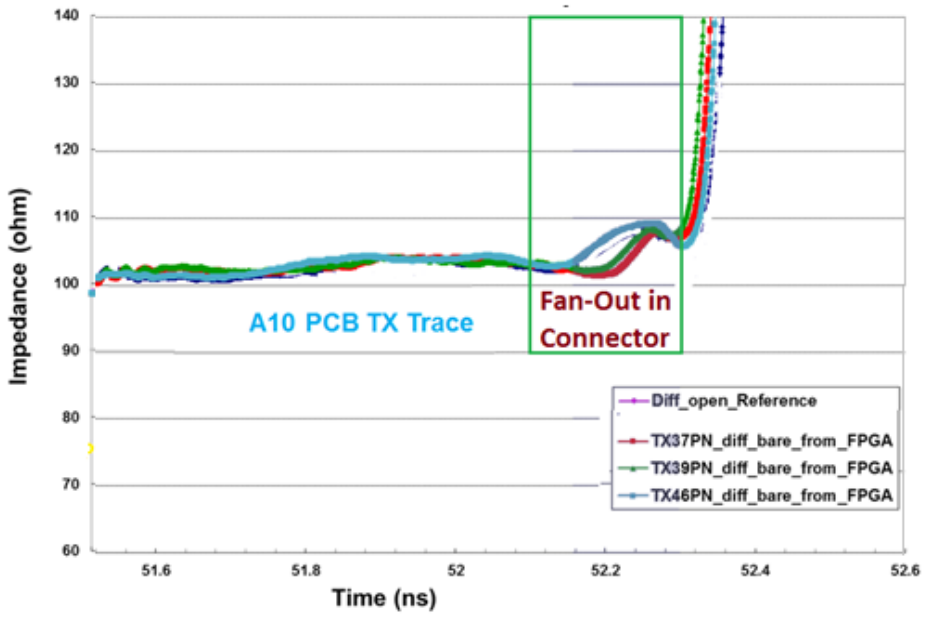

Related Information

Electrical Specifications on page 87

Electrical Specifications

CEI 28 Very Short Reach (VSR) Specifications for CFP2/CFP4/QSFP+/QSFP28/zQSFP/SMA 2.4 mm

Very short reach (VSR), short reach (SR) and long reach (LR) electrical specifications can be found in details in *Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps and 25G+ bps I/O* document.

Please refer to chapter 13 of *Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps and 25G+ bps I/O* for all electrical specifications for CEI-28G-VSR Very Short Reach channel/Interface. You can find the reference model and test points, return loss (both differential and common mode) specifications for all reference points, full channel reference model and its maximum insertion loss specifications in *Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps and 25G+ bps I/O*.

Related Information

Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps and 25G+ bps I/O

Interlaken Interface Specification

Interlaken interface electrical specifications can be found in *Interlaken Interoperability Recommendations*.

This document focuses on Interlaken interface recommendations for typical applications up to 400 Gbps packet transfer. It also shows the connector and cable assembly channel and contributed differential insertion loss. The end to end interoperability loss budget can also be found in this document as the required specification in channel design process.

Related Information

Interlaken Interoperability Recommendations

Document Revision History for AN 766: Intel Stratix 10 Devices, High Speed Signal Interface Layout Design Guideline

Document Version	Changes
2019.03.12	Updated maximum transceiver data rates. NRZ was 30 Gbps, is 28.9 Gbps, PAM4 was 56 Gbps, is 57.8 Gbps.
2018.08.14	Global editorial changes only.
2017.05.08	Updated the capacitor from 0404 to 0402 in the "AC Coupling Capacitor Layout and Optimization Guidelines" topic.
2016.11.11	Initial release.