MoSys Application Note

Title: High Speed Board Design Guidelines
Author: M. Baumann
Date: 2020, 04, 03

Document No: AN_GD_HS BOARD DESIGN GUIDELEINES_200406

SUMMARY

This application note provides easy to follow guidelines for designing a high-speed PCB.

It includes layout recommendations for use with the MoSys® LineSpeed™, Bandwidth Engine® (BE), and Programmable HyperSpeed Engine (PHE) high-speed serial products.

Reference board schematic and layout are available for most products.

KEY POINTS:

- Signal Integrity
- Power Integrity
- Power and simulation
- Sample Check List for PCB Layout

MoSys Device: Component overview cross-section of devices and PCB board.

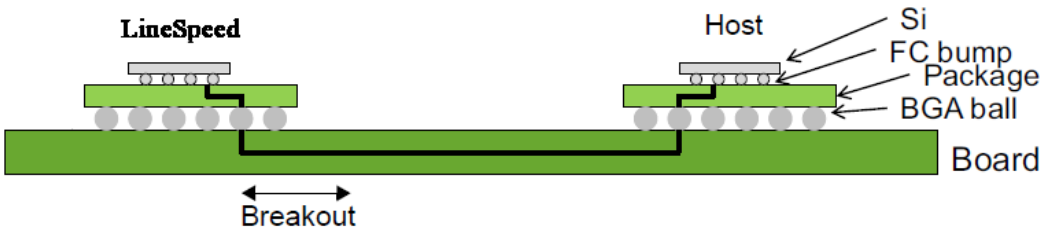

1 Introduction

This application note describes board layout recommendations for use with the MoSys® LineSpeed™, Bandwidth Engine® (BE), and Programmable HyperSpeed Engine (PHE) high-speed serial products. Reference board schematic and layout are available for most products. These can be used as starting points. There are several useful guidelines available for other products that are relevant to high-speed traces, such as Altera's layout guidelines (AN-672, 2013). In the following, several alternative guidelines are presented, as the constraints in layout sometimes make it difficult to follow a particular guideline. Using the most conservative guidelines is recommended as it will improve the odds of success as well as help reduce the effort required to run simulations and verify with measurements.

Both signal integrity and power integrity guidelines are provided. The emphasis is on the signal integrity, as it is generally easier to achieve a decent power distribution on the board with multiple planes available and room for capacitors on both top and bottom. Following these guidelines does not guarantee success on the first try. However, following the more conservative guidelines with consistent layout across all lanes will minimize the amount of simulation needed. I.e. one should be able to simulate and optimize only one or a small representative number of signals in order to validate the target impedance, the via transition, backdrill connector transitions before releasing the board to fabrication.

Table 1: Definitions of terms

Term	Description
LineSpeed	The MoSys LineSpeed 100G PHY product family for high-speed datapath signal integrity
Bandwidth Engine (BE)	The MoSys Bandwidth Engine® Serial memory product family. Uses serial high-speed GigaChip® Interface (GCI) to host
Programmable HyperSpeed Engine (PHE)	The MoSys Programmable HyperSpeed Engine family – intelligent offload. Uses serial high-speed GCI to host
Host	The ASIC/ASSP/FPGA connecting to a LineSpeed, BE, or PHE device. It could also be a connector in certain application.
Flip Chip (FC or C4)	The bump connecting the chip to the package.
BGA (Ball Grid Array)	The solder ball connecting the package to the board.
Breakout	The routing region under the BGA and between BGA to the main board route. These routes may require special considerations such as narrower traces, plane voiding, via fields, etc. to escape out.
MoSys Package Model	Will include chip pad/Flip Chip bump to BGA/PCB interface.

Contents

Figures

Tables

2 Signal Integrity

Reference clocks and high-speed signals in LineSpeed, BE, and PHE products are designed as differential pairs, which have one positive (P) and one negative (N) pin. Usually these two pins are next to each other in the ball map and have reference/GND net surrounding them. The remainder of this section addresses the following: First what type of routing to use (microstrip or stripline), then what material and stackup to use, and finally layout guidelines to minimize both return insertion loss.

2.1 Strip-Line or Microstrip Routing

To connect the differential pair to its counterpart in another device, stripline or microstrip could be used. The cross-section below illustrates the two.

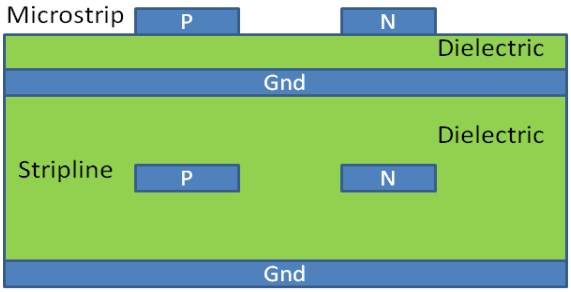

Figure 1: Cross-section illustrating microstrip and stripline differential pairs

The main advantage of a microstrip is that it can eliminate vias between the two devices if the routing is done on the top surface. In addition, using microstrip eliminates the need for backdrill if routing is constrained to top and bottom of the PCB. There are many disadvantages to using microstrip. Due to the coupling mainly down to the ground underneath, the majority of the conduction in the trace is constrained to the bottom layer due to skin effect, while that of a stripline can spread to both the top and the bottom of the traces, enabling lower conductive loss for a given pitch. In addition, special plating of outer copper layers and solder mask may cause high frequency losses. Other components on the board that may not be known or characterized at the time of the board design may affect the trace impedance too (e.g. sockets, stiffeners, mats put on surface during testing etc). An example is shown in the following table. Optimized test traces with Rogers RO4350B dielectric in a 50-Ohm stripline (L4, L6, L8) and microstrip (Bottom) traces of equal width and 2" and 4" length were measured at 12.9 GHz (25 Gbps). The effective loss per inch of the bottom microstrip trace was found to be about 2X that of the inner layers.

Table 2: Loss per inch for microstrip and stripline traces

Layer	4" trace	2" trace	Loss/inch	1 MXP
L4	**5.36 dB**	**3.94 dB**	**0.71 dB**	**1.26 dB**
L6	**5.25 dB**	**3.91 dB**	**0.67 dB**	**1.29 dB**
L8	**5.3 dB**	**4.05 dB**	**0.63 dB**	**1.4 dB**
Bottom	**8.28 dB**	**5.25 dB**	**1.38 dB**	**1.38 dB**

MoSys therefore generally recommends stripline (routing in the middle layer), compared with micro-strip (routing on the surface). As long as the customer evaluates all pros and cons based on the system requirement, and use electrical simulation to guide the selection, both stripline and microstrip routing can work when properly optimized.

2.2 Material Selection and Stackup

It is preferred to select metal/dielectric layers that have low loss and small process variation. Materials such as Nelco and Rogers have lower loss than FR4. The variation in layer thickness will be translated to signal impedance variation. Thus, it is better to control the variation at the stackup design stage. In general, larger dimensions are better. Due the wide variety of materials and system constraints, no particular recommendation is provided. However, it is generally possible to mix both low loss layers and low cost FR4 layers in one board design as shown in the following stackup.

For low speed signals or power planes, regular low cost FR4/370HR is used. In this board, the high-speed signals at data rates of 28 Gbps are given low loss material in the middle and upper layers. To minimize the loss, the Rogers dielectric should be placed both above and below the Signal layer. A ground layer should also be placed above and below the signal trace to provide a continuous return path and ensure that cross-talk is contained. Notice that Rogers has been used between the top signal layer and the ground, as there is inevitably some routing to the components in the top layer. Even though it is a short routing, the high frequency loss of FR4 is quite high, so even a short trace can affect the overall loss. When backdrilling the vias on the high frequency traces from the bottom of the board, it essentially eliminates all FR4 coupling, so there is no signal integrity degradation of using a mix of Rogers and FR4 in this stackup.

The PCB material (dielectric, metal, and metal roughness) and trace lengths all affect the insertion loss once the traces are optimized and the return loss is minimized.

Figure 2: Stackup with Rogers dielectric for high-speed signal layers

2.3 Characteristic Impedance

Many parameters have a strong function of frequency, and it is therefore important to consider frequencies in a wide range vs. only the Nyquist rate (0.5X baud rate). The common electrical interface spec (CEI Spec, 2014) specifies the insertion and return loss, as a function of the frequency, normalized to the baud rate. An example of the return loss spec (characteristics impedance match) is shown in Figure 1. For calibration with a single number spec; at the Nyquist frequency of 0.5 the S11 spec from the graph is about -8dB. Along with the spec, the actual S11 simulation data from two revisions of a PCB board at 28 Gbps (0.5 in the graph corresponds to 14 GHz) is also shown. They both meet the CEI spec. However, the new revision has better impedance matching and will therefore provide more margin to RX, TX or other imperfections in the channel (e.g. vias or connectors).

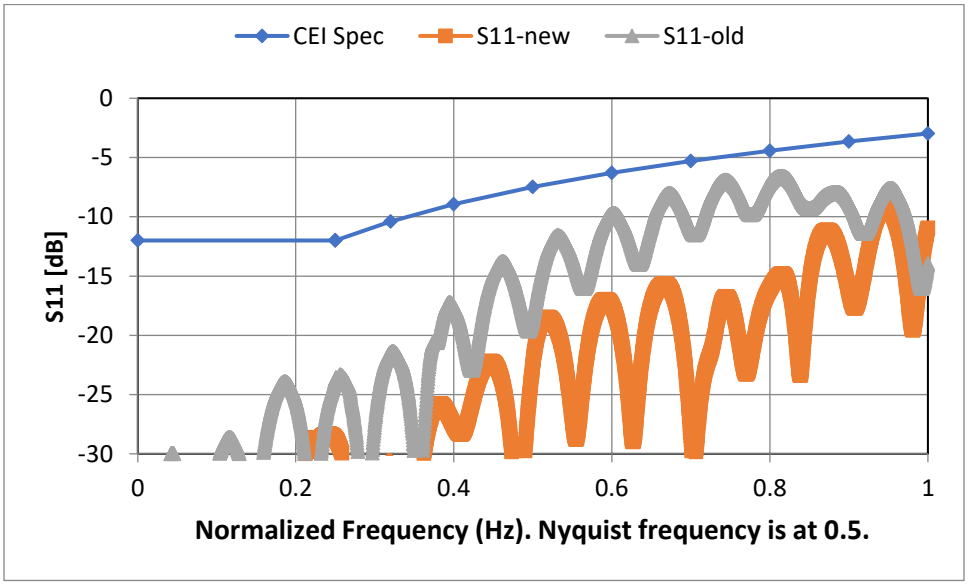

Figure 3: S11 normalized fb – the maximum baud rate in channel is 1

2.4 Impedance Control

The impedance control for a differential pair has two parts: horizontal and vertical (via) routing. The following sections provide guidelines for both.

2.4.1 Impedance Control Layout Guidelines - Horizontal Routing

Unless there are already guidelines from fab or earlier designs, start with a simple 2D simulation tune the trace geometry to meet the impedance target. If there are more than one setting that meet the target, the wider and/or thicker trace is preferred, since the metal with bigger cross-section will have relatively less conductor loss and smaller process variation. The following tool is one of many that are readily available: (Multi-Teknik Simulator, 2016). Below are some general guidelines:

- Impedance control: nominal differential impedance of 100 Ω ± 3% in simulation
- The reference plane above and beneath stripline/microstrip should be solid and should not have void along the traces.
- Minimize the number of vias and avoid discontinuity in the reference plane, such as split planes.
- In the ball-grid-array (BGA) area, narrow traces with narrow spacing can be used to escape the ball field, but should be kept as short as possible.
- Via-in-pad is recommended for escaping the ball field.
- Avoid tight bends, and use smooth curves to make turns.

2.4.2 Impedance Control Layout Guidelines - Vertical Routing (VIA's)

Even though the vertical routing through the vias of the board represents only a small proportion of the total trace length, it is often the cause of the most significant impedance discontinuity. The following guidelines need to be considered for every via along high frequency signals or clock traces.

A differential pair of vias is shown below. The critical diameters are also illustrated. This pair of vias connects a pair of chip balls at layer L1 down to inner traces of L4. Notice that the antipads of both L1 and L4 are bigger than the other layers to reduce the capacitive coupling of the landing pads. For the purposes of this engineering guideline, the recommendation is to make as large as practical antipad for the L1 and L4 to the point where this capacitance can be ignored. The main guidelines will be for the other layers. Notice the need to specify backdrill up to L4, which means that there will be a stub from Layer 4 to Layer 5 (typically 5-10 mils) as shown in the cross-section. This also requires extending the antipad guidelines to L5 and L6 in order not to add extra capacitance to the vias, i.e. to at least two layers below the routing signal layer (L4).

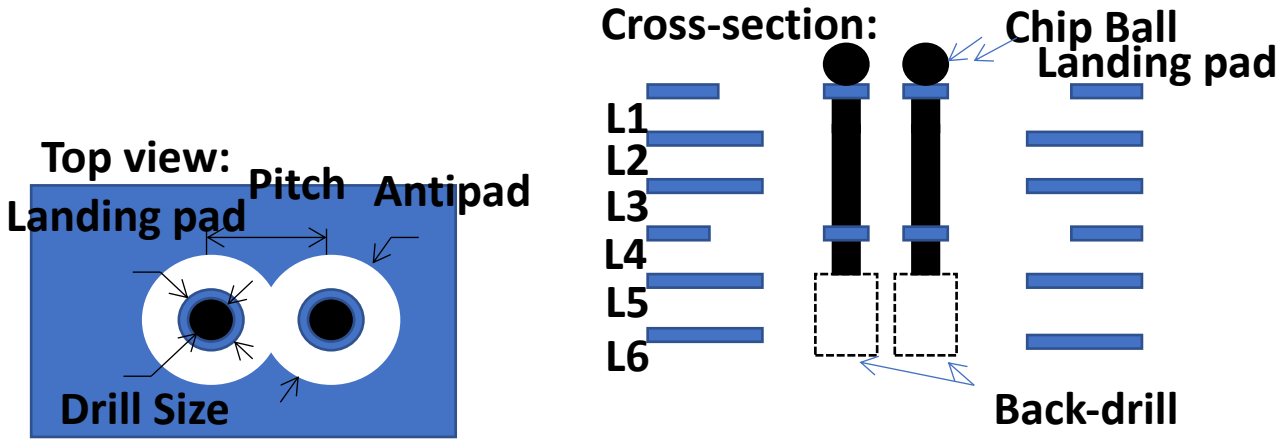

Figure 4: top and cross-section views of VIA's

Typical guidelines for a Rogers or FR4 PCB via are illustrated below. Use 8 mil drill size and an antipad of 50 mils in order to achieve approximately 100Ω characteristic impedance. Detailed simulations will be required to optimize this via transition depending on the actual stackup and material.

Table 3: Typical guidelines for 100-Ohm differential vias

Pitch	Drill Size	Antipad
40 mil (1mm)	8 mil	50 mil

Similarly to the adjacent ground shields in the horizontal routing, the adjacent ground vias are important for signal vias. This is shown in the following Figure. It is important to have a continuous ground return path, which is provided by the ground vias.

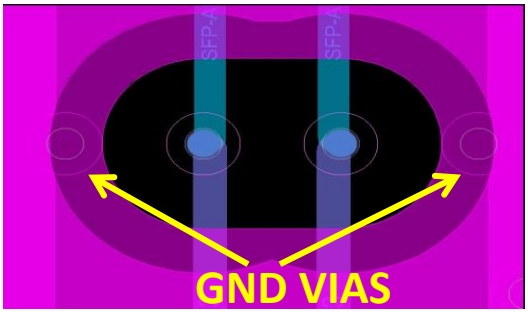

Figure 5: Top view of differential vias with adjacent ground VIAS

Summarizing the key points for the vertical/via routing:

- Antipad design is critical near the devices/balls. Make them all consistent so that simulation of one is representative of all.
- To reduce the impact of via stubs, it is recommended to back-drill vias to have the minimum residual stub. Consult your PCB vendor for this, and use accurate residual stub in your simulations.
- Back-drilling might be performed from both top and bottom of the PCB. Clearly define the drill symbol in the PCB design file.
- It is recommended to use stitching vias along with the signal vias.

2.5 Cross-Talk

The common electrical interface spec cross talk noise specification is a function of insertion loss (CEI Spec, 2014) as shown in the following Figure. To help illustrate this, consider the following: The TX output swing is generally in the order of 1V differential, and the worst case is generally two adjacent lanes that couple to a given lane. This means that the cross-talk spec for lane-to-lane at low insertion loss is in the order of 1% or -40dB [peak-to-peak to rms: (divide by $\sqrt{2}$) and from 1 to 2 lanes: (multiply by $\sqrt{2}$) so it cancels out. One percent of 1V is 10mV]. As the channel loss is increased, the cross-talk noise threshold is reduced. For <5dB insertion loss, the tolerable cross-talk from a lane to an adjacent lane is therefore about -40dB. For 20dB of insertion loss, the tolerable cross-talk from a lane close to source (e.g. TX) to an adjacent lane (RX) is 0.3% or -50dB.

The Figure below illustrates that the worst case cross-talk will be for TX to RX or a different signal (clock or digital signal) to RX. The aggressor RX for the RX-RX cross-talk would generally also have loss along the line, so the amplitude of noise injection is reduced proportionally to the loss of the line. Hence, a common guideline is to spec RX-RX at -40dB, while any other signal to TX at -50dB at the Nyquist frequency. Notice that other signals such as GPIO, MDIO, I2C, SPI and CLK signals may have much higher swing than 1V. The cross-talk from these signals to high-speed signals are therefore even stricter. The most conservative guidelines are recommended to eliminate the need to simulate this cross-talk.

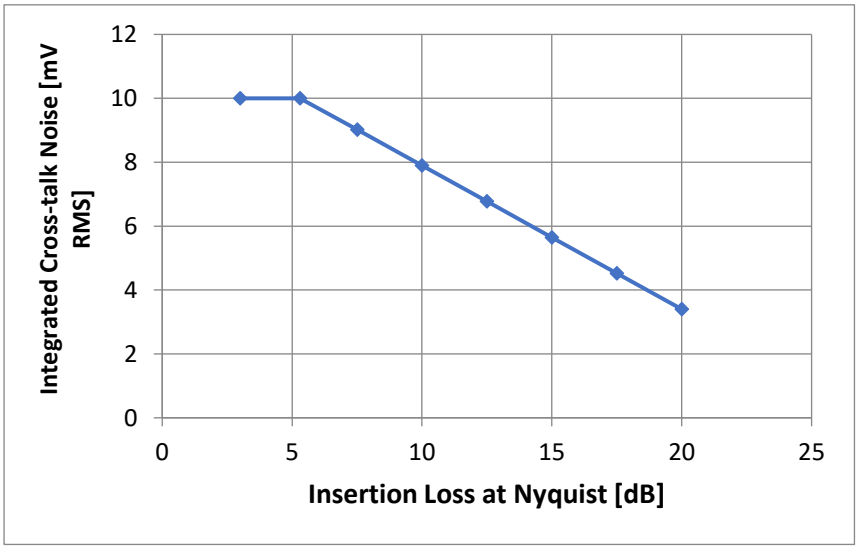

Figure 6: Integrated cross-talk Noise as a function of IL at Nyquist

2.5.1 Ideal Routing

These guidelines are presented to help achieve a given cross-talk spec. Generally, best practice is the well shielded approach that in many cases reduces or eliminates the need for doing cross-talk analysis.

Below, a cross-section is illustrated for an ideal routing of a differential pair. There is a continuous ground plane above and below the pair, and there is plenty of space to the adjacent grounds. The adjacent grounds are wide and stitched with vias as illustrated. In this case, the simple 2D calculation gives a fairly accurate estimate of the impedance, because the coupling to the adjacent grounds is small. The advantage of this approach is that cross-talk may not need to simulated/analyzed much or simulated at all in some cases.

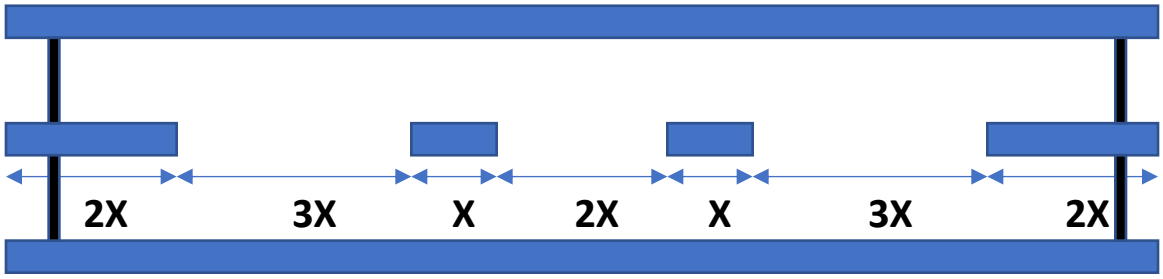

Figure 7: Cross-section of ideal routing; ground shielded differential pair

A top view layout from the PHE characterization board follows this guideline. Notice the wide ground between the two differential pairs and the smooth turns that result in low loss. Notice also that this layout is still in progress: need to add ground via stitching at the top of this layout.

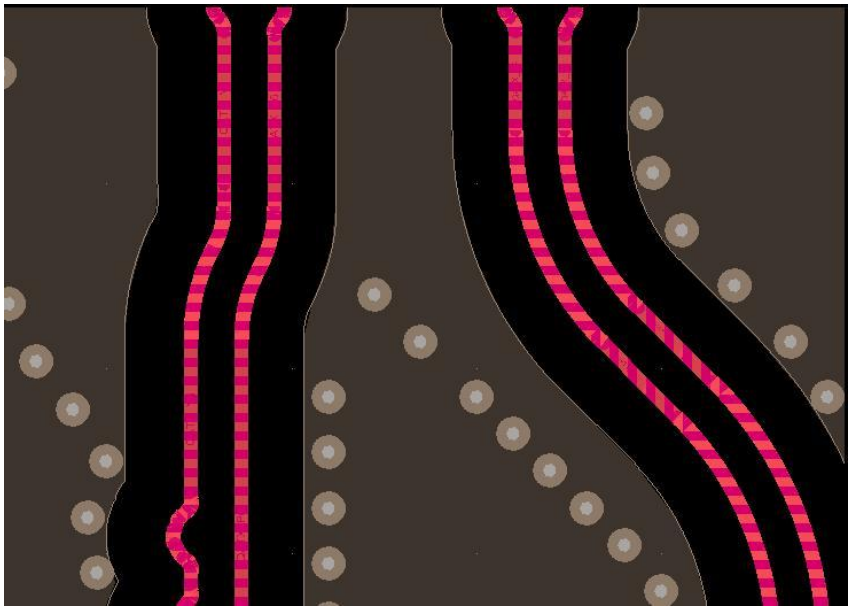

Figure 8: Top view of layout of differential pair

2.5.2 Non-Ideal Routing – Use with Care and Simulate Impedance in Detail

In some cases it may be difficult to adhere to the general rules above due to congestion of signals. In that case, all the dimensions can be reduced to still maintain a finite ground trace between adjacent pairs as illustrated in the following.

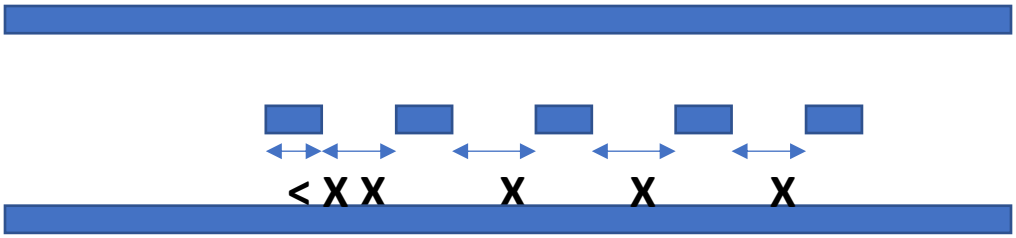

Figure 9: Cross-section of routing with small distance to adjacent GND

In the following layout example, notice the thin sliver of ground that is placed in between the differential pairs in the break-out region in order to limit the cross-talk. As soon as they come out from the breakout region, this ground as well as spacing from the traces to the grounds are widened back to the ideal routing. The breakout region needs to be extracted and simulated as smaller spacings reduce the characteristic impedance from the ideal 100Ω, and the cross-talk is increased.

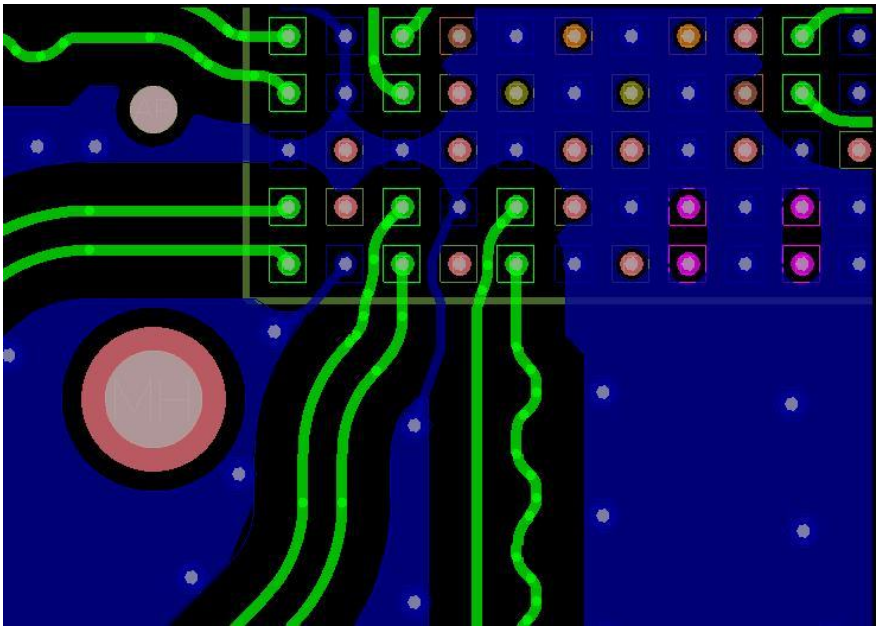

Figure 10: Top view of layout with limited distance to adjacent ground shield

2.5.3 Risky Routing

In some cases, it may be difficult to fit in a trace in between adjacent differential pairs. In this case, the cross-talk will be significant, and needs to be analyzed in detail (Bogatin, 2013). A 3X spacing is shown below, but that may not be sufficient to reduce cross-talk. This would be classified as risky routing: you will generally save both time and money to use a more conservative approach that includes better shielding vs. trying to optimize this type of layout.

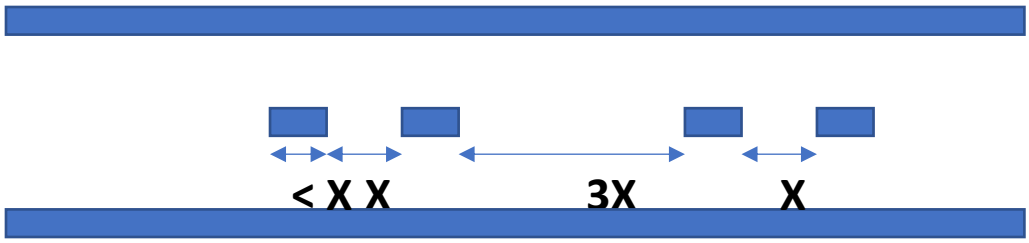

Figure 11*:* Cross-section of routing without ground shield between pairs

It is generally not recommended, but it is used a lot for cases where cross-talk can be tolerated. An example is shown for a board with a short reach requirement of 12.5 Gbps. Notice that the signal trace widths are reduced closer to the be2 break-out in order to reduce the cross-talk. This increases the loss, so it is just done for the short distance in the vicinity of the ball breakout.

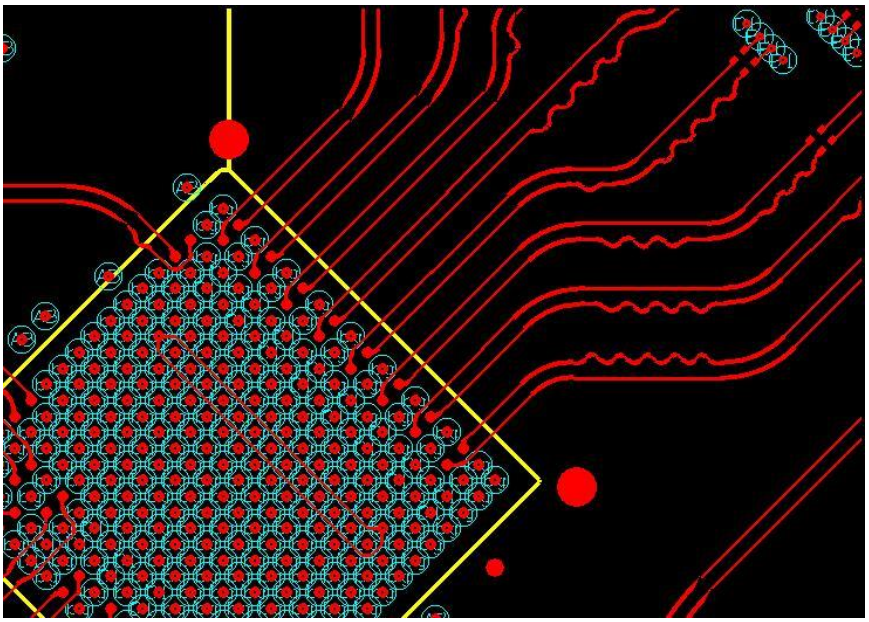

Figure 12: Top view of routing of relatively low speed differential pairs

2.6 Skew Between P and N

For a differential pair, positive (P) and negative (N) lines are ideally exactly the same length and matched/mirror routing all the way from TX to RX. In reality, there will be a difference between them, physically as well as electrically. A skew between P and N causes a reduced eye margin and a finite common mode signal (differential to common mode conversion). The spec of this conversion is a function of frequency (CEI Spec, 2014). At Nyquist frequency (1/2 bit rate), the maximum allowed differential to common mode conversion is less than -15dB. Even with a perfectly match length between P and N, the non-uniform material properties of a PCB could affect the signal delay (Ritchey, 2016). It is therefore critical to match the lengths accurately in the layout.

For data rates of 10-28 Gbps, NRZ, matching lengths within ± 1 mil is recommended. In addition, it is recommended to match the length close to the location of the original mismatch. Several traces are illustrated below. The trace illustrated with note 1 has a length mismatch due to the ballout at the device. The wiggle pattern is used to increase the length of the shortest trace as close to the original mismatch as possible (Intel Addendum, 2010). Due the constraints of minimizing cross-talk, it was not possible to do this for all traces in this test board: e.g. the trace to the right of #1 has the length equalization mid-way along the trace instead of at the bottom.

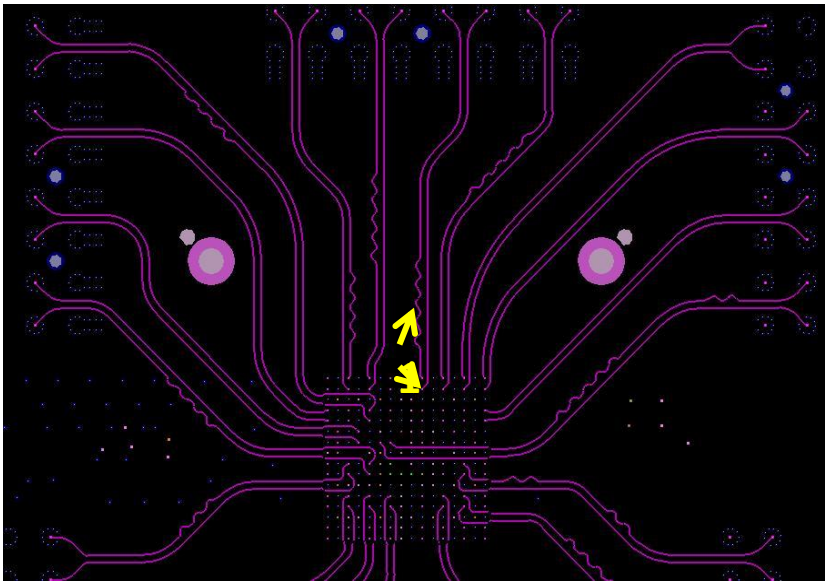

Figure 13: Illustrations of delay matching with a wiggle pattern (12.5 Gbps).

2.7 Reverse Polarity to Optimize Layout

In most MoSys products, polarity of the differential pair can be reversed for both transmitter and receiver. I.e. a TXP from the host can go to RXN to the MoSys product (and TXN to RXP), as the design can internally correct for the polarity. This feature gives freedom to a layout engineer to reduce length, simplyf routing, and in some cases more accurately match length without. In the figure below, yellow lines are the correcting traces that reverse the polarity of the original differential pair. Please note that the change in layout needs to be back-annotate to PCB schematic and software team needs to be notified for the polarity reverse.

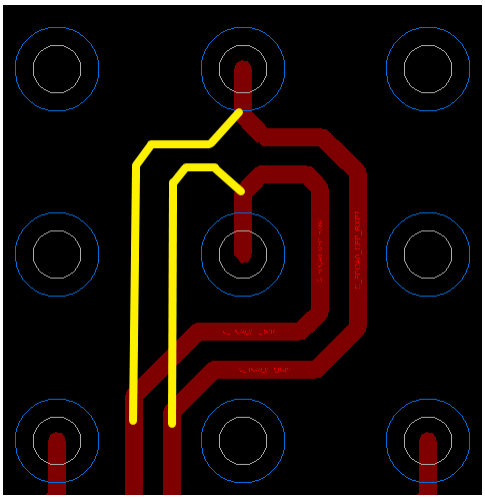

Figure 14: Allow polarity reversal to reduce extra turns in routing

2.8 AC Coupling Capacitor

AC coupling capacitors are sometimes required for high-speed signals. The best solution is to choose a low loss, broad band AC coupling capacitor that has a size that is about the same width as the traces in order make the characteristic impedance of the components match as closely as possible to the signal impedance. One case is illustrated below where 0201 size 0.1□F capacitors are used. They are slightly wider than the traces, and would therefore normally cause a capacitive discontinuity. To compensate for this, the ground plane underneath the component is partially cut out, as shown in the Figure below. This is not ideal, as it makes the ground return not uniform. Ideally even smaller AC coupling capacitors should be used with a continuous ground plane underneath.

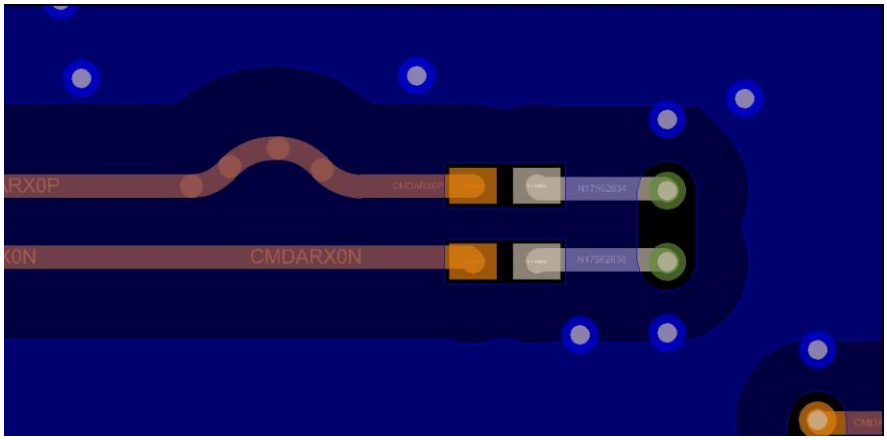

Figure 15: Illustration of 0201 size AC coupling capacitors

2.8.1 Example Layout Including Horizontal and Vertical Routing

In the following, details of a trace between a flip chip BGA and a QSFP28 connector is shown. The layout of the differential trace is shown using a 3D viewer (HFSS EM Field Solver, 2016). The signal traces are on the inner layers (as seen from the connection to the vias) and both BGA and QSFP28 connectors are at the top of the board.

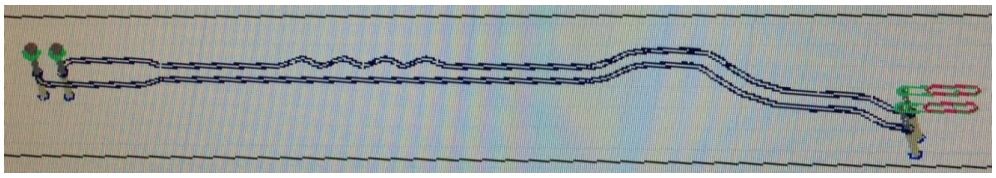

Figure 16: Differential trace between BGA (left) and QSFP28 (right)

The top two layers are shown below. There is a large cutout around the QSFP pads in the L1 top layer (in red) and a much smaller cutout in the L2 layer below the pads (in green).

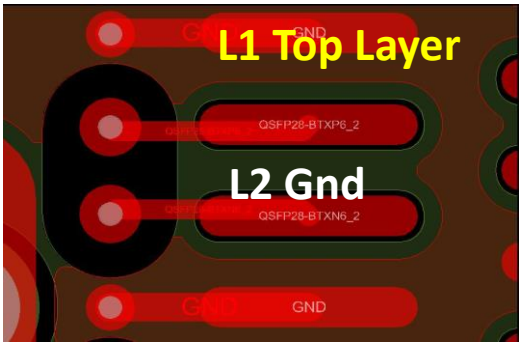

Figure 17: Top view of QSFP pads and L2 ground underneath pads

From measurements of test structures it was discovered that the impedance of the QSFP connector appeared to have lower than the target 100□ characteristic impedance. Upon a a detailed analysis of the original layout (HFSS EM Field Solver, 2016) it was found that the initial layout had in fact too low characteristic impedance near the QSFP pads. After optimization, the improved layout shows a much better match as shown below.

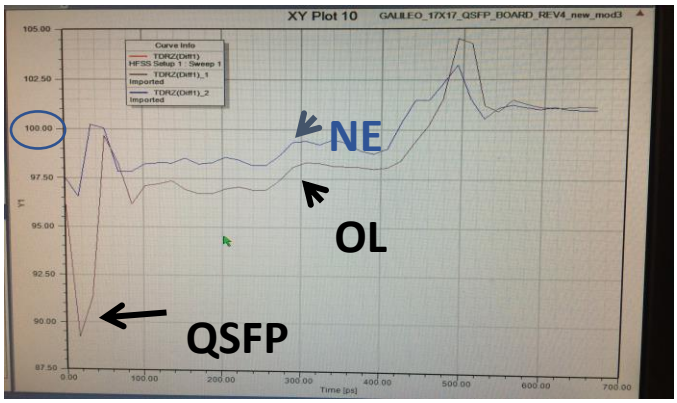

Figure 18: TDR simulations of initial and revised layout

The actual change was mainly to reduce the coupling to L2 ground, and make the L2 similar to L1 ground as shown below. The main ground shield for the signal trace was instead a continuous L3 (not shown). I.e. the capacitance was reduced. Instead of trying to accommodate a different trace width from the vias to the wide pad, the trace width was set equal to the pad width. This made it a clean continuous return path in L3 instead of the return path that jumps between L3 and L2 in the old layout.

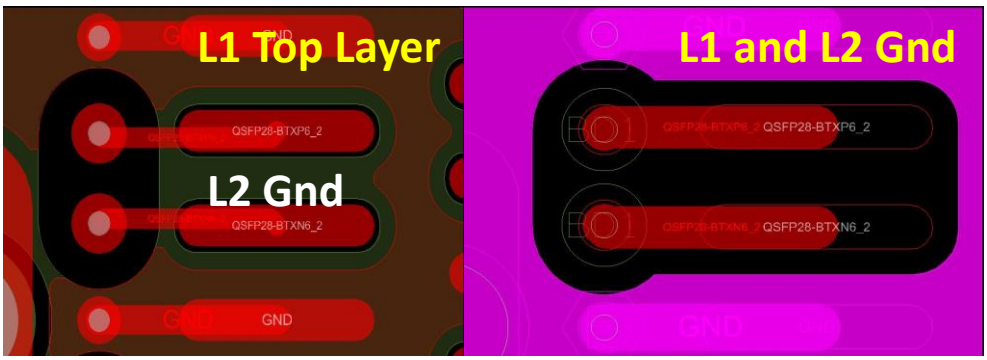

Figure 19: Initial (left) and revised (right) layout of L1 and L2 Gnd

The change in layout is also illustrated in 3D in the following. The vias are back drilled; shown with a different color in the layers below the signal routing below. There was a stackup change from the old to the new layout which was not due to this signal integrity problem, so the signal trace is now closer to the center of the board. This is not related to the impedance problem and fix near the QSFP connector. The difference in routing between the QSFP pads and the vias are highlighted with an arrow.

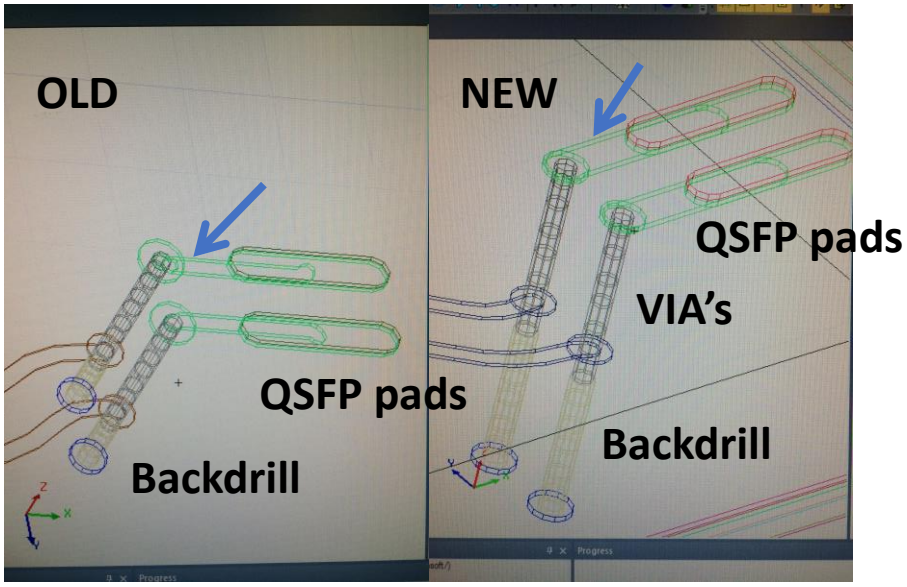

Figure 20: Initial and Updated differential signal layout near QSFP pads

The final S-parameters of the initial and updated layout are shown along with the CEI spec in the figure. Although both pass the CEI spec, the new layout will give better margin to account for imperfections in the TX/RX as well as for process variations of the package and board.

3 Power Integrity

To reduce supply ripple voltage, it is generally recommended to use power planes in the PCB with alternating VDD/GND layers, and to put capacitors as close to the device as possible. There are several documents available that generally recommend the use of multiple, small, low inductance capacitors (Intersil Application Note 1325, 2011), (Ripple Voltage and ESR, 2012). What is not always mentioned is the inductance in the PCB traces that connect to the capacitor. The inductance in those traces is just as important. It is recommended to treat the supply, the GND as well as parallel path of decoupling capacitors as a continuous plane: the current needs to flow through all of them. Simlarly to the importance of vertical routing of signal, consider the vertical routing of power as well: place VDD and VSS vias adjacent as pairs to minimize the inductance.

The backside of the PHE characterization board illustrates this: almost the entire backside of the PCB is filled up with decoupling capacitors. The package is supporting this by having supply balls adjacent to GND balls, enabling a simple placement of supply capacitors between power balls. The cross in the middle does not have decoupling capacitors to allow the socket stiffener to connect at the backside.

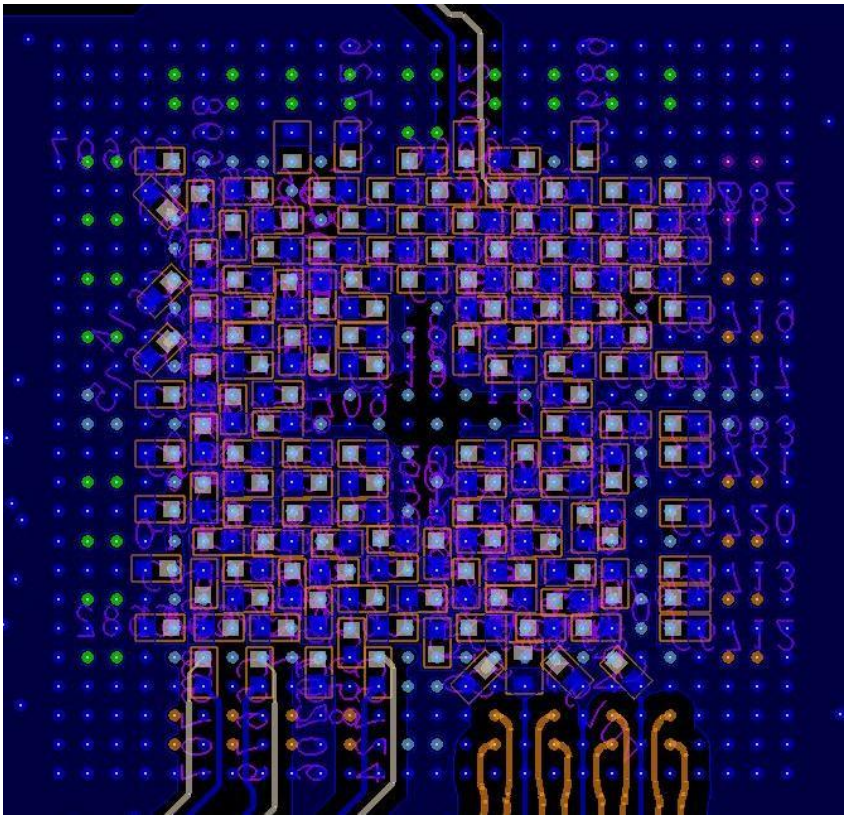

Figure 21: Backside of PCB, illustrating good power decoupling.

In the design below, there are not as many decoupling capacitors. The power supply di/dt current surges in this product are about an order of magnitude lower than those of the PHE, so not as many decoupling capacitors are required. However, a similar approach is used: all supply balls in the package are placed next to GND balls, allowing the capacitors to be placed between supply/GND vias at the backside of the board, avoiding any high-inductance supply traces to the capacitors. In addition, there is a ground plane in the bottom layer to reduce the inductance.

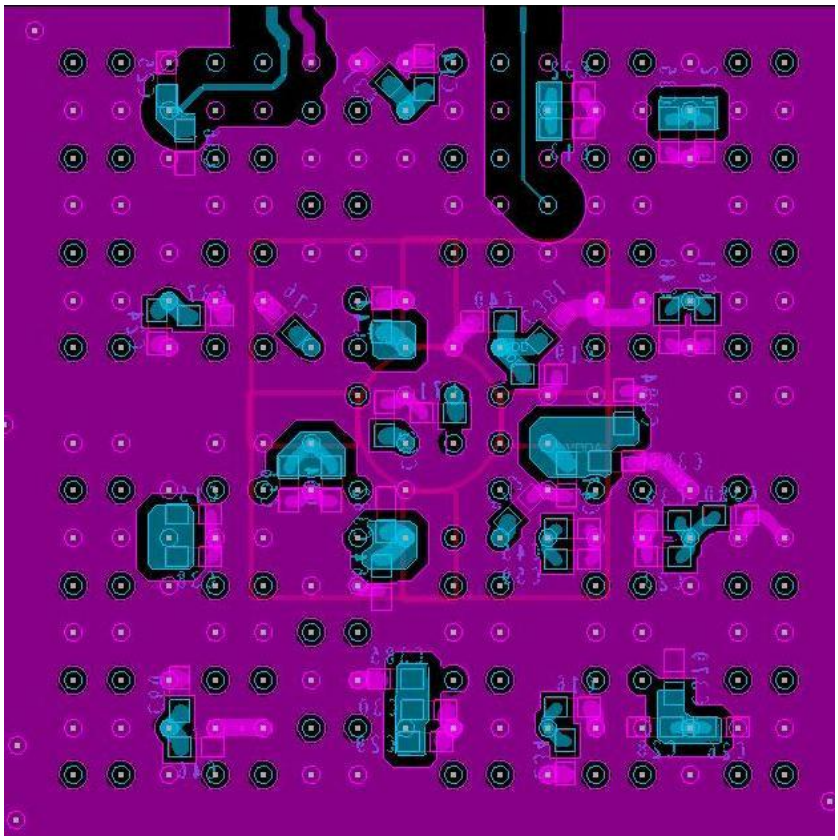

Figure 22: Backside of PCB, illustrating OK power decoupling capacitors

The most accurate way to determine decoupling requirement is to obtain di/dt information for each supply and simulate the board regulators, layout and the current to obtain a given voltage ripple/droop spec. All this information is not always available. It is therefore generally recommended to use a range of decoupling capacitors to reduce the ESR over a wide frequency range – e.g. capacitors at 100, 10, 1, and 0.1microF.

In some products, there are already many low inductance capacitors on the package, and it is then recommended to use larger capacitors on the board. For example, the PHE product has 38 low inductance 2.2microF capacitors on the package and it does not help much to add equal or lower value capacitors on the PCB board. For the PHE, 22microF capacitors on the backside are recommended in as small footprint as possible.

The recommendation of ‘as small a footprint as possible’ is not always valid. Smaller footprint than 0201 may in fact result in larger effective inductance when considering the traces connecting to the capacitors.

Please check following aspects during power supply design of board in accordance to specifications listed in the MoSys data sheet:

- Bulk Decaps - Check if quantity and values are sufficient for filtering
- Decap-Mid and High freq - Check if quantity and values are sufficient for filtering
- Tantalum Caps Polarity visible and properly connected
- Check if Caps have sufficient Voltage rating: generally recommend >2.5X rating
- Check if Ferrite bead used when a power supply is shared between analog and digital domain
- Check if Ferrite bead DC resistance is low such that the power supply voltage level is within tolerance
- Check if Ferrite bead chosen have sufficient current rating
- Make sure Sense line connected very close to the load
- Check if Power supply sequencing follows the data sheet

Once above considerations and design choices are made, use appropriate tool to make sure routing of power supplies meet the IR drop target specified for both DC and AC tolerance limits.

4 Measurement and Simulation

MoSys supplies IBIS-AMI model for I/O. It is recommended that these models are incorporate into system level simulation with PCB and ASIC/FPGA models to verify the channel performance.

The channel layout on the PCB must be optimized in order to meet the strict insertion and return loss masks defined by corresponding CEI-28G-VSR/SR/MR/LR and/or IEEE specifications depending on the application.

PCB model can be obtained through either extraction or measurement. To be specific

- Four-port S-parameter is recommended to check insertion and return loss.
- Eight-port S-parameter is recommended to check crosstalk between two differential pairs.

To obtain transmission line model from PCB layout, user needs to feed the layout file to extraction tool, define port at the terminal of the traces, specify frequency sweep and allow the numerical engine to crunch numbers to produce S-parameter. It is not uncommon that hybrid method is used to obtain the final S-parameter, in which via model is extracted using 3-D electromagnetic software, while traces are extracted using 2-D software. These models are consolidated into a single S-parameter by cascading them sequentially and running a frequency sweep simulation. It is worth mentioning that user needs to pay attention to the following:

- Whether the stackup of the layout file contains the correct material property. If not, correct it in the extraction software
- Whether the backdrill is applied in the 3-D model. If not, certain stub has to be shortened or removed.

PCB channel model can also be obtained through measurement. Vector-network-analyzer is a commonly used equipment in the lab for obtaining such model. The measurement technique involves calibration, probe selection etc., which is beyond the scope of this document. User can consult local equipment vendor for more detail. Besides measuring the data in frequency domain, time domain measurement such as TDR is also another way to characterize the PCB channel. Again, equipment vendor is a good resource to gather information and details for that.

For simulation purposes, the following table is a general guideline; refer to product for specific/variations (especially for insertion loss S21).The specs are often specified at the Nyquist frequency (0.5X maximum baud rate). That is a minimum for these parameter. Ideally, these specs are achieved up to 2X Nyquist frequency. E.g. For 28 Gbps, these parameters are should be attained in simulations at 28 GHz.

Table 4: Sample maximum specifications for high-speed signal traces.

S11	**S21**	**Xtalk, Rx-Rx**	**Xtalk, any→Rx**	**Skew**
-12 dB	-5 to -30dB	-40dB	-50dB	+/- 1 mils

5 Sample Checklist for PCB Layout

It is recommended to review the following consolidated list of items prior to releasing the board to fab.

Signal and Power Integrity	Comment or Waiver
Check whether test coupon/structure, if any, covers enough variety of the PCB, such as single-ended vs. differential, top layer vs. middle layer, etc.	Covered high-speed traces.
High-speed trace via stub lengths are less than 10 mils. Back drilling is used for vias with stub lengths higher than 10 mils.	
Return loss is < -12dB for all high-speed diff. pairs	at 1.5X Nyquist frequency
Planes used for power delivery of all rails. If a plane is not possible, sufficiently wide traces are used (width > 100 mils)	
Capacitor pad is connected to power and ground plane with larger vias to minimize loop inductance	
Wide - short traces are used between the vias and capacitor pads or vias are placed adjacent to capacitor pads	
Power supply sense lines are connected very close to the load	
High frequency decoupling capacitors are placed very close to device	
Routing and Placement	**Comment or Waiver**
There are no 90-degree corners on traces	Only applies to high-speed SerDes traces
High-speed traces routed on impedance-controlled layers	With Rogers
Multiple vias on high-speed traces and clock lines are avoided	
High-speed differential pairs are not routed close to clock lines	
GND vias are placed close to single ended and differential signal vias	
Stitching vias are used to tie all GND planes.	
Stitching via diameter roughly equal to trace width	
Each GND pin or via are connected to plane individually	
High-speed traces are not routed near or across discontinuities in the reference plane such as splits or voids	
High-speed traces are not routed over an antipad	
Signal via pads on unused internal layers removed	
Traces on unshielded neighboring signal layer runs perpendicular to minimize crosstalk	

Tight bends are avoided	
Differential Pair Design	
Skew between P and N signals of a differential pair are matched	+/- 1 mil for HS traces, +/- 5 mil for clocks
Skew between P and N signals are matched on a per layer basis	+/- 1 mil for HS traces, +/- 5 mil for clocks
Differential vias are placed as a pair in a symmetrical fashion	
Spacing between differential pairs is 3X higher than spacing between P & N traces to reduce cross talk	Include ground shielding
P and N traces of a differential pair are routed on the same layer	
Width and Spacing of the differential pair P and N traces are as per fab vendor specification	
High current power rails are not routed close to TX or RX differential pairs	Planes are separated by gnd shields
Impedance variation is less than +/- 10%	HFSS simulations
AC Coupling Capacitors	
GND Plane under the AC coupling capacitor pads removed to improve return loss	Instead use 01005 capacitors
AC coupling capacitors placed far way from TX	
Signal and Power Integrity	
Check whether test coupon/structure, if any, covers enough variety of the PCB, such as single-ended vs. differential, top layer vs. middle layer, etc.	Covered high-speed traces.
High-speed trace via stub lengths are less than 10 mils. Back drilling is used for vias with stub lengths higher than 10 mils.	
Return loss is lower for all high-speed diff. pairs	
Planes used for power delivery of all rails. If a plane is not possible, sufficiently wide traces are used (width > 100 mils)	
Capacitor pad is connected to power and ground plane with larger vias to minimize loop inductance	
Wide - short traces are used between the vias and capacitor pads or vias are placed adjacent to capacitor pads	
Power supply sense lines are connected very close to the load	
High frequency decoupling capacitors are placed very close to device	
Mechanical	
Connectors/ switches/sockets have sufficient clearance space around them	
Adequate number of mounting holes are present	

If a socket is used, sufficient clearance space is provided around DUT	
If a socket is used, socket mounting hole locations are verified to be correct	
If SMA connectors are used, there is sufficient clearance space between SMAs for handling	
A heatsink if required can be attached without any issues	

6 References

AN-672. (2013, February 15). *Transceiver Link Design Guidelines for High-Gbps*. Retrieved from https://www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/an/an_672.pdf

Bogatin, E. (2013, June 26). *Robust crosstalk design rules*. Retrieved from EDN Network: http://www.edn.com/design/test-and-measurement/4417648/2/Robust-crosstalk-design-rules

CEI Spec. (2014). *http://www.oiforum.com/public/documents/OIF_CEI_03.1.pdf*.

HFSS EM Field Solver. (2016). *http://www.ansys.com/products/Electronics/ANSYS-HFSS*.

Intel Addendum. (2010, March). *Intel Atom processor D400/D500 storage platform*. Retrieved from http://download.intel.com/design/processor/specupdt/323311.pdf

Intersil Application Note 1325. (2011, October 10). Retrieved from http://www.intersil.com/content/dam/Intersil/documents/an13/an1325.pdf

Multi-Teknik Simulator. (2016). *http://www.multek.se/engelska/engineering/pcb-structures-2/differential-stripline-impedance-calculator-2*.

Ripple Voltage and ESR. (2012, March). Retrieved from http://www.niccomp.com/help/ESR-RippleVoltage-032012.pdf

Ritchey, L. (2016, October 4). *Solving signal integrity problems at very high data rates*. Retrieved from http://goo.gl/9gJZRh

7 Version History

Date	Version	Changes
January 2014	**0.9**	**Initial release**
June 2014	**1.0**	**Editorial changes.**
November 2014	**1.1**	**added guidelines for mating Ardent cable assemblies**
October 2016	**2.0**	**Added detailed guidelines for layout**
October 2016	**2.1**	**Corrections to diagram and verbiage in section 2.1**

MoSys, Inc.

2309 Bering Drive

San Jose, CA 95131

USA

Phone 408-418-7500

Fax 408-418-7501

Web http://www.mosys.com