



XAPP1316 (v1.1) December 21, 2018

PCB Design Considerations for FPGA Accelerator Cards

Author: Ravindra Gali

Summary

This application note gives an overview of the accelerator card form factor as defined by the *PCI Express Card Electromechanical Specification, Revision 3.0* [Ref 1]. It addresses printed circuit board (PCB) design challenges, from stackup design, to dielectric material selection, to the PCB fabrication technology used in the PCB design process. [Component Placement Guidelines](#) focuses on layout best practices for high-speed memory design, high-speed SerDes channel design, and power delivery network design.

Download the reference design files for this application note from the Xilinx website. For detailed information about the design files, see [Reference Design](#).

Introduction

In addition to hyperscale (cloud) computing, there has been steady growth in high-performance computing (HPC) from government agencies, oil, financial services, and life sciences industries focusing on data mining and analysis for threat monitoring, pattern and image recognition, encryption/decryption, options evaluation, risk analysis of assets, seismic modeling and analysis, gene encoding and matching, and drug modeling and discovery. These applications are compute- and data-intensive and constantly need increased compute power and memory bandwidth. This growth has resulted in data center architects trying to come up with new server architectures to increase performance and efficiency. CPU-based systems augmented with hardware accelerators as coprocessors are emerging as an alternative to CPU-only systems. This has created opportunities for accelerators like graphics processing units (GPUs), FPGAs, and other accelerator technologies to advance HPC hyperscale systems to previously unattainable performance levels.

Any accelerator deployed in a data center must be compatible with the target workloads through its deployment lifetime. This requirement is a challenge, given the diversity of workloads and the rapid rate at which workloads change. It is thus highly desirable for accelerators incorporated into data centers to be programmable. This programmability makes an FPGA the ideal candidate for the data center market.

Data Center PCB Form Factor Requirements

The accelerator card form factor requirements are defined by the *PCI Express Card Electromechanical Specification, Revision 3.0* [Ref 1].

The accelerator card can be designed in any of the form factors in [Table 1](#), depending on the application requirements.

Table 1: Accelerator Card Form Factors

Card Type	Maximum Height (inches)	Length (inches)
Low profile/Slim PCIe® card	2.731	4.72 or 6.59
Half-length	4.381	6.6
Three-quarter length	4.381	9.45
Full length	4.381	12.28

All cards should have a thickness of 1.57 mm (0.062 in.) \pm 0.13 mm (0.005 in.) to comply with the PCIe specification. The maximum height listed in [Table 1](#) is from the bottom of the edge fingers to the top of the card.

A PCIe Gen3 x16 card edge connector is used to interface to the host server. The interface provides PCIe signals and power to the card via the 12V and 3.3V pins.

[Figure 1](#) is an example of a half-height, half-length (low profile) card.



X19322-070517

Figure 1: Example of a Half-Height, Half-Length, Low Profile Card

Figure 2 is an example of full height, half-length card.



X19323-070517

Figure 2: Example of a Full Height, Half-Length Card

Figure 3 is an example of a full height, three-quarter length card.

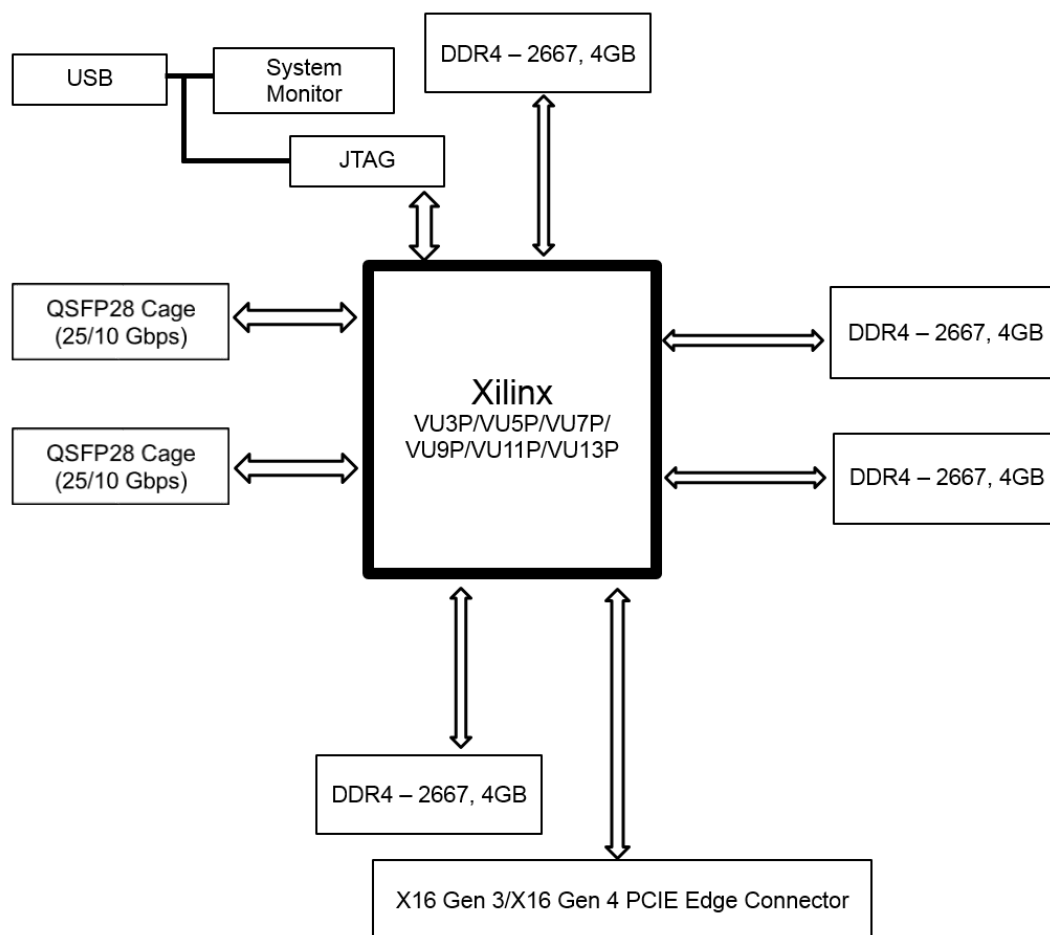


X19324-070517

Figure 3: Example of a Full Height, Three-Quarter Length Card

Architecture

Figure 4 presents a block diagram of a typical accelerator card.



X19325-082317

Figure 4: Block Diagram of a Typical Accelerator Card

Features of a typical card are summarized as follows:

- Target Device: Xilinx® Virtex® UltraScale+™ FPGA
- SDRAM: 2 to 4 x72 DDR4 interfaces operating at 2667 Mb/s in a device down configuration
- Ethernet Ports: 2 to 4 QSFP28 connectors 2 x100 GbE to 4 x100 GbE
- A Gen3/Gen4 PCIe x16 interface connected to a x16 card edge connector to interface with the host
- Quad SPI flash memory to configure the FPGA
- USB/JTAG/SYSMON for monitoring temperature and power

Typically, HPC customers such as government agencies select a full height, full length form factor card using one of the bigger Virtex FPGAs such as the VU9P, VU11P, or VU13P. These cards typically do not have a specific power limit and are constrained by reliability and thermal considerations. The FPGA power on these cards is on the order of 100W to 200W.

Hyperscale customers usually prefer a low profile card (half-height, half-length) or a half-length, full height or three-quarter length, full height card using a VU3P, VU5P, VU7P, or VU9P device, depending on the application. These cards have different power limits, ranging from 10W (x1), 25W (x4, x8), to 300W, depending on the addition of optional power connectors.

Stackup Definition and Material Selection

A significant challenge in designing an accelerator card that complies with the PCIe specification is ensuring the card thickness stays within 62 mil \pm 5 mil. The thickness restriction significantly limits the total number of layers that can fit in, considering signal/power integrity, mechanical and thermal requirements, and PCB fabrication requirements while ensuring the overall cost is kept to a minimum.

A typical PCB stackup is comprised of alternate layers of copper foils separated by a dielectric material. The PCB fabricator purchases the standalone laminate and prepreg sheets from the laminate vendors and stacks them alternately, depending on the layer count. These laminate and prepreg sheets are glued using a heat-press treatment as part of the PCB fabrication process. The laminate/core is comprised of a thin dielectric material with copper clad foils bonded to both sides. These copper clad foils form the inner layers of the PCB. The core's dielectric material is comprised of cured (hardened) fiberglass-weave material with epoxy resin that acts as an insulation layer between the copper foils. The main purpose of the woven fiberglass weave is to provide mechanical strength for the PCB in both the X and Y directions. The woven glass fabric is available in different roll widths, styles (1035, 1080, 2113, 3313, or 7628), and thicknesses.

Various resin systems are available for specific application requirements. Some of the different resin types are epoxy-based systems (FR4), Polyimide, Teflon-based resins, Polyphenylene Oxide (PPO), Polyphenylene Ester (PPE), Cyanate Ester (CE), and Bismalamine Triazine (BT). The parameters of interest are dielectric constant D_k or ϵ_r , Loss tangent (Df) or dissipation factor, glass transition temperature requirements (Tg), ability to withstand soldering and rework, moisture absorption, and cost. The resin selected must be compatible with the glass fabric and the copper foil.

Unlike the laminate, a prepreg is a sheet of the dielectric material (fiberglass sheet saturated with resin) whose resin has not been fully cured. The prepreg acts as the insulation between core layers and is the gluing agent for the cores. The heat-press treatment cures the prepreg and binds all the layers together to form the PCB. The laminate vendors provide the laminate/core sheets in a wide variety of constructions such as different glass types for enhanced signal integrity performance, different fiberglass weaves, resin content (%), core thicknesses, copper weights (1/2 oz., 1 oz., 2 oz.), and copper foil types, among other considerations.

Table 2 lists specifications for a typical 16-layer stackup that meets PCIe thickness specifications.

Table 2: Typical PCI-Compliant Sixteen-Layer Stackup

Layer		Material Type	Material Name	Glass Style	Material Pressed ϵ_r (@ 10 GHz)	Material Pressed Thickness (mil)
1	Top	Copper				0.7
		Prepreg	I-Speed IS	1067MS ⁽¹⁾	3.27	3.0
2	PWR/GND	Copper				0.6
		Core	I-Speed IS	1067MS	3.27	3.0
3	SIG	Copper				0.6
		Prepreg	I-Speed IS	1067MS	3.27	3.5
4	PWR/GND	Copper				0.6
		Core	I-Speed IS	1067MS	3.27	3.0
5	SIG	Copper				0.6
		Prepreg	I-Speed IS	1067MS	3.27	3.5
6	PWR/GND	Copper				0.6
		Core	I-Speed IS	1067MS	3.27	3.0
7	SIG	Copper				0.6
		Prepreg	I-Speed IS	1067MS	3.27	3.5
8	PWR/GND	Copper				1.2
		Core	I-Speed IS	1067MS	3.27	3.0
9	PWR/GND	Copper				1.2
		Prepreg	I-Speed IS	1067MS	3.27	3.5
10	PWR/GND	Copper				1.2
		Core	I-Speed IS	1067MS	3.27	3.0
11	PWR/GND	Copper				1.2
		Prepreg	I-Speed IS	1067MS	3.27	3.5
12	SIG	Copper				0.6
		Core	I-Speed IS	1067MS	3.27	3.0
13	PWR/GND	Copper				0.6
		Prepreg	I-Speed IS	1067MS	3.27	3.5
14	SIG	Copper				0.6
		Core	I-Speed IS	1067MS	3.27	3.0
15	PWR/GND	Copper				0.6
		Prepreg	I-Speed IS	1067MS	3.27	3.0
16	Bottom	Copper				0.7

Table 2: Typical PCI-Compliant Sixteen-Layer Stackup (Cont'd)

Layer		Material Type	Material Name	Glass Style	Material Pressed ϵ_r (@ 10 GHz)	Material Pressed Thickness (mil)
Pressing thickness excluding plating and solder mask						59.33 $\pm 10\%$
Total thickness including plating and solder mask						62.87 $\pm 10\%$

Notes:

1. MS is *mechanically spread*.

Table 3 shows nominal trace width spacing requirements using the Table 2 stackup to achieve a 39 Ω (single-ended), 50 Ω (single-ended), and 100 Ω differential impedance.

Table 3: Typical PCI-Compliant Sixteen-Layer Stackup Impedance Table

Single-Ended (50 $\Omega \pm 10\%$)		Single-Ended (39 $\Omega \pm 10\%$)		Differential (100 $\Omega \pm 10\%$)		
Line Width (mil)	Simulated Impedance (Ω)	Line Width	Simulated Impedance (Ω)	Line Width (mil)	Spacing (mil)	Simulated Impedance (Ω)
3	~50	5	~39	3	10	~100
3	~50	5	~39	3	10	~100
3	~50	5	~39	3	10	~100
3	~50	5	~39	3	10	~100
3	~50	5	~39	3	10	~100

The choice of dielectric material plays a significant role in determining the overall layer count. The key parameters to consider when selecting a dielectric material are:

- the dielectric constant Dk/ϵ_r
- the loss tangent Df
- the glass transition temperature (T_g)
- the fiber weave characteristics
- the dielectric breakdown voltage (DBV)

Dielectric Constant

The relative dielectric constant (ϵ_r or Dk) is equal to the ratio of the capacitance of a parallel plate capacitor with the given insulator (dielectric material) compared to the capacitance of an identical capacitor in a vacuum without an insulator (dielectric material). The dielectric constant of a vacuum is 1. All other materials have dielectric constants higher than 1.

The increase in capacitance is due to the re-aligning of the dipoles in the presence of an electric field. It also reduces the speed of the electromagnetic waves in a medium by a factor of $\sqrt{\epsilon_r}$ compared to air (see Equation 1):

$$\sqrt{\epsilon_r} = C/V \quad \text{Equation 1}$$

where:

- C is the speed of light in free space (0.0118 in./ps)
- V is the measured propagation velocity in the presence of a given dielectric material

$$\text{Propagation delay} = (\sqrt{\epsilon_r})/C \quad \text{Equation 2}$$

The ϵ_r of a dielectric material that is specified in the material data sheet is the effective ϵ_r value based on a given fiberglass/resin composition because the glass and resin material have their own ϵ_r values based on the type of glass and resin material chosen for creating the dielectric material. The effective Dk or ϵ_r is used to calculate the impedance of the transmission lines. The Dk or ϵ_r varies with frequency and is inversely related to frequency, that is, the dielectric constant of a material decreases with increasing frequency.

[Table 4](#) lists dielectric materials that are commonly used for PCB design. Most of the values listed below are specified at 10 GHz or 12.5 GHz based on the manufacturer data sheet. Dielectric materials like Tachyon 100G or Megtron 7 are relatively new and are especially suitable for very high-speed design although at a higher cost.

Table 4: Typical Dielectric Materials Used for PCB Design

Material Name	Relative Dielectric Constant	Loss Tangent	Highest Data Sheet Frequency (GHz)
Generic FR4	4.50	0.015	N/A
Megtron 4	3.80	0.005	1
TU-872LK	3.80	0.009	10
FR408HRIS	3.37	0.0092	10
EM-888	3.80	0.008	10
EM-888K	3.20	0.006	10
Nelco N4000-13	3.70	0.008	10
Nelco N4000-13 Si	3.30	0.007	10
Megtron 6	3.63	0.004	12
Rogers 4003C	3.38	0.0027	10
Rogers 4350B	3.48	0.0037	10
Taychon 100G	3.02	0.0021	10
Megtron 7 (Low Dk Glass)	3.35	0.002	12
I-Speed	3.63	0.0071	10
I-Speed IS	3.27	0.0064	10

The advantage of picking a dielectric material with a lower ϵ_r compared to one with a higher ϵ_r is that you can attain a higher trace impedance for a given trace geometry keeping everything else the same. You need to either plan for a significantly lower trace geometry or increase the dielectric thickness to achieve the same impedance when using a higher ϵ_r dielectric material.

[Table 5](#), [Table 6](#), and [Table 7](#) are some examples of the trace geometries needed to achieve the single-ended/differential impedances needed for implementing a DDR4 QSFP+ interface. These trace geometries are arrived at based on the stackup details in [Table 2](#).

Figure 5 was used as the reference stackup for the single-ended impedance trace width calculations using FR408HRIS from Isola shown in Table 5 and Table 6, and Figure 6 was used for arriving at the trace width and spacing requirements shown in Table 7 to meet the 100- Ω differential impedance requirements. The calculations show the dielectric thickness must be increased by 1.4 mil for a given inner routing layer using a generic FR4 substrate to achieve the same impedance for a given trace width. This results in a significant board thickness increase of approximately 25% compared to the above stackup, thereby failing to meet the PCIe thickness specification. The additional thickness can also result in additional cost due to the need for an advanced PCB fabrication process if it exceeds the 10:1 via aspect ratio that is fairly common with the standard PCB fabrication process.

Table 5: Example Stackup to Achieve 39 Ω Single-Ended Impedance


Dielectric Material	ϵ_r	Trace Thickness (T)	W (mil)	D (H1+H2+T) (mil)	Single-Ended Trace Impedance (Ω)
Generic FR4	4.50	0.6	5	8.6	~39
I-Speed IS	3.27	0.6	5	7.1	~39

Table 6: Example Stackup to Achieve 50 Ω Single-Ended Impedance

Dielectric Material	ϵ_r	Trace Thickness (T)	W (mil)	D (H1+H2+T) (mil)	Single-Ended Trace Impedance (Ω)
Generic FR4	4.50	0.6	3	8.6	~50
I-Speed IS	3.27	0.6	3	7.1	~50


Table 7: Example Stackup to Achieve 100 Ω Differential Trace Impedance

Material	ϵ_r	Trace Thickness (T)	W (mil)	Spacing (S)	D (H1+H2+T) (mil)	Differential Trace Impedance (Ω)
Generic FR4	4.5	0.6	3	10	8.6	~100
I-Speed IS	3.27	0.6	3	10	7.1	~100



X19900-122018

Figure 5: Single-Ended Stripline



X19516-092717

Figure 6: Differential Edge-Coupled Centered Stripline

Loss Tangent (Df)

This parameter is a measure of the amount of RF energy absorbed by the laminate. The resin system has a separate loss tangent and so does the fiberglass cloth. The loss tangent specified in the material data sheet is a combination of the two. Similar to Dk, the loss tangent $\tan(\delta)$ or Df also varies with frequency and is dependent on the composition of the glass-to-resin ratio.

The first order approximation for a dielectric loss is governed by [Equation 3](#):

$$\text{Attenuation [dB/in]} = 2.3 \times f \times Df \times \sqrt{\epsilon_r} \quad \text{Equation 3}$$

where:

- f is the sine wave frequency in GHz, equivalent to the Nyquist
- Df is the loss tangent
- ϵ_r is the relative permittivity

See *Loss in a channel: Rule of Thumb #9* [\[Ref 2\]](#).

Many papers published in the industry quantify the impact of dielectric and conductor losses across a wide range of frequencies. It has been estimated that the ratio of conductor loss to the total loss was about 13% for an FR4-based substrate with $Dk = 4.4$ and $Df = 0.02$ to 20 GHz [\[Ref 3\]](#). However, the conductor loss as a percentage of the overall loss increased to around 30% when using a low-loss dielectric with a $Dk = 3.7$ and $Df = 0.002$. These findings clearly illustrate that the dielectric loss is the dominant loss mechanism to be concerned with when designing a PCB for high-speed applications.

Note: Picking a dielectric material with low Dk and low Df is recommended for the best signal integrity.

Glass Transition Temperature

The fiberglass transition temperature (T_g) is the temperature at which a resin's coefficient of thermal expansion changes to a much larger value than it has at lower temperatures. This is a critical parameter for PCB manufacturability because the z-axis expansion can cause plated

through-holes to fracture due to excessive stress if the PCB soldering temperature exceeds the T_g of the dielectric material.

Fiber Weave Characteristics

The main purpose of the woven glass fabric as mentioned earlier is to provide mechanical strength in both X and Y directions for a PCB. The fiberglass weave manifests itself as variations in impedance and velocity along a trace if the distribution of the glass in the dielectric is not uniform. In a differential pair, this results in intra-pair skew between the P and N legs. This is more pronounced for data rates beyond 10 Gb/ps because the amount of skew due to this effect can often be more than 10 to 15% of UI depending on the data rate and length of the channel. There has been lot of industry research in this area to quantify the impact of fiber weave on signal integrity performance and ways to mitigate its effect [Ref 4]. A simple and effective solution to mitigate this effect during the stackup definition process is to pick a laminate/prepreg material using a mechanically spread glass. The laminate manufacturer goes through the extra step of applying energy to enlarge and flatten the glass in both X/Y directions to make it more uniform [Ref 5]. This application of energy has the effect of reducing variations in dielectric constant seen by the transmission lines as it traverses a PCB. Some of the examples of this type of glass weave style are 1067MS, 1086MS, 1078MS, and 3313MS among others. The *MS* next to the glass style indicates *mechanical spread*.

Note: Pick a laminate/prepreg material using a mechanically spread glass such as 3313MS to minimize the fiber weave effect.

Dielectric Breakdown Voltage

Apart from Dk and Df characteristics, the other key requirement when selecting a laminate and a prepreg material is the dielectric breakdown voltage requirements of the design. The typical specification for most of the products catering to the telecommunications market is 1500V. Most of the commercially available laminates have breakdown voltages around 1000V per mil of thickness (25 μm). Hence, using a laminate or prepreg less than 2 mil is not feasible when creating a stackup definition.

Copper Foil Considerations

In addition to the dielectric material considerations that have been discussed, carefully consider the copper foil characteristics needed to achieve robust signal integrity performance for multi-gigabit channels.

There are two basic methods for manufacturing copper foils. These are rolled copper foils and electro-deposited (ED) copper foils. The ED copper foil is the dominant method for manufacturing copper foils used in multilayer PCBs. The copper used must be able to achieve good peel strengths so that the copper does not pull away from the laminate material during the PCB fabrication process. The copper is typically designated by weight (1/2 oz., 1 oz., 2 oz.) and foil type. Some of the different foil types include the standard ED copper foils, reverse-treated foil (RTF), high temp elongation (HTE) shiny copper, double-treated (DT) copper, and very low profile (VLP/e, VLP/H, VLP). There are specific advantages and disadvantages associated with each foil type. The standard ED/HTE copper is not the most optimum foil to achieve robust impedance control and signal integrity performance for

supporting multi-gigabit data rates, unlike the VLP foils, because of additional loss due to copper surface roughness. However, the downside with VLP foils is the relatively high cost and the low peel strength as opposed to other copper foils.

The first-order approximation for a conductor loss is governed by [Equation 4](#):

$$\text{atten[dB/in]} \sim 1/w[\text{mils}] \cdot \sqrt{f[\text{GHz}]}$$

Equation 4

where:

- f is the sine wave frequency in GHz, equivalent to the Nyquist
- w is the line width in mils

See *Loss in a channel: Rule of Thumb #9* [\[Ref 2\]](#).

The conductor loss can be further divided into scattering loss caused by surface roughness and the skin effect loss ([Equation 5](#)).

$$\text{Total Conductor Loss} = \text{Skin Effect Loss} + \text{Scattering Loss (Surface Roughness)}$$

Equation 5

At high frequencies, the majority of the current distribution is pushed to the outer surface or "skin" of the trace due to skin effect. The amount of penetration into the metal, known as skin depth, is usually represented by the symbol δ . The skin depth is inversely related to the square root of frequency ([Equation 6](#)):

$$\delta = \sqrt{2/(\omega\mu\sigma)}$$

Equation 6

where:

- ω is angular frequency
- μ is permeability
- σ is the conductivity of copper

The skin depth is approximately 2 μm (80 $\mu\text{in.}$) at 1 GHz. Hence, specifying 1/2 oz. copper (700 $\mu\text{in.}$) for signal traces as part of stackup design is more than sufficient for all high-speed links because the bulk of the current flows at the surface or bottom of the trace. The skin effect manifests itself primarily as an increase in resistance proportional to the square root of frequency. This is called conductor loss ([Equation 7](#)):

$$R_{ac} \sim \text{Width (W)} \times \text{skin depth}$$

Equation 7


One way of minimizing the skin effect loss is to choose wider trace geometries to increase the surface area. However, the down side to that approach is a thicker PCB, higher crosstalk, and more cost. The height to the nearest reference planes has to be increased to make the traces wider for a given impedance. Moving the traces further away from the reference planes also has the potential to cause higher crosstalk if adequate trace-to-trace spacing cannot be maintained due to routing real estate limitations.

At very high frequencies (>1 GHz), the skin depth is on the same order as the typical surface roughness of the copper foils. This results in additional conductor loss because the copper roughness further impedes the current flow.

The roughness of the copper is dependent on the construction of the copper foils. The industry uses multiple terms (Ra, Rz, and RSAR) to specify copper roughness in μm [Ref 6]. Ra is defined as the average surface roughness, RSAR stands for *roughness surface area ratio*, and Rz is typically referred to as a ten-point height. The Rz value for a typical copper foil used in a PCB is around 6 to 10 μm . This value is significantly higher than the skin depth at 1 GHz, making it imperative for you to account for conductor losses due to copper surface roughness especially for PCB designs with transceiver channels operating at multi-gigabit data rates.


To address this issue, the laminate vendors provide laminate sheets with different copper foil constructions to cater to specific applications, although at a cost premium. RTF is the standard finish used by almost all laminate manufacturers. The Rz, Ra, and RSAR specifications for these foils are significantly different, resulting in differing amounts of conductor loss. It has been shown that you can expect at least a 1.5 dB improvement in conductor loss at 20 GHz moving from an RTF to a type of VLP foil. This change results in approximately a 17% improvement for overall conductor loss [Ref 3].

Figure 7 through Figure 10 show the insertion loss profiles for the various dielectric materials on a per-inch basis. Figure 7 is the insertion loss for a 3-mil wide trace assuming a smooth copper trace ignoring the effect of surface roughness. Figure 8 shows the insertion loss for the same 3-mil wide trace assuming an average surface roughness of 0.6 μm (Ra).



X19327-070517


Figure 7: Insertion Loss Per Inch for a 3-mil Wide Trace (No Copper Surface Roughness)



X19328-070517


Figure 8: Insertion Loss Per Inch for a 3-mil Wide Trace (With Surface Roughness $R_a = 0.6\mu\text{m}$)

Figure 9 shows insertion loss for a 4-mil wide trace assuming a smooth copper trace ignoring the effect of surface roughness. Figure 10 shows insertion loss for the same 4-mil wide trace assuming an average surface roughness of 0.6 μm (Ra).



X19329-070517

Figure 9: Insertion Loss Per Inch for a 4-mil Wide Trace (With No Surface Roughness)



X19330-070517

Figure 10: Insertion Loss Per Inch for a 4-mil Wide Trace (With Surface Roughness Ra = 0.6 μ m)

Table 8 and Table 9 summarize the average loss per inch of trace across some key Nyquist frequencies for some select dielectric materials.

Table 8: Frequency Loss per Inch of Trace (4 mil with Cu Surface Roughness Ra = 0.6 μ m)

Material (Trace Width W = 4 mil) with Ra = 0.6 μ m	4 GHz (PCIe Gen3)	7.5 GHz (HMC 15G)	8 GHz (PCIe Gen4)	14 GHz (QSFP 28)
Generic FR4	0.611	1.04	1.10	1.79
Megtron 4	0.500	0.838	0.884	1.40
TU-872LK	0.482	0.804	0.848	1.34
FR408HRIS	0.480	0.800	0.844	1.33
EM-888	0.464	0.770	0.811	1.27
Megtron 6	0.393	0.637	0.669	1.02
Taychon 100G	0.360	0.574	0.603	0.91
Megtron 7 (Low Dk Glass)	0.357	0.569	0.597	0.90

Table 9: Frequency Loss per Inch of Trace (3 mil with Cu Surface Roughness Ra=0.6 μm)

Material (Trace Width W = 3 mil) with Ra = 0.6 μm	4 GHz (PCIe Gen3)	7.5 GHz (HMC 15G)	8 GHz (PCIe Gen4)	14 GHz (QSFP 28)
Generic FR4	0.686	1.16	1.23	1.97
Megtron 4	0.576	0.958	1.01	1.58
TU-872LK	0.558	0.923	0.973	1.52
FR408HRIS	0.558	0.921	0.970	1.51
EM-888	0.540	0.889	0.936	1.46
Megtron 6	0.467	0.753	0.791	1.20
Taychon 100G	0.439	0.698	0.732	1.10
Megtron 7 (Low Dk Glass)	0.434	0.690	0.724	1.08


A brief summary of findings based on these tables shows:

- The insertion loss is strongly dependent on the Dk and Df of the dielectric material along with the trace width for given surface roughness.
- The Ra parameter plays a significant role in the overall insertion loss as shown in [Figure 7](#) through [Figure 10](#). Your laminate vendor and PCB fabricator should be able to provide you with the Rz, Ra, RSAR, and Rq values for the surface roughness of the copper depending on the copper foil you choose for your PCB stackup.
- You can minimize conductor loss by selecting a wider trace width to achieve the characteristic impedance. However, the penalty for doing so is an increase in board thickness. Going with a 4-mil trace is a reasonable compromise between minimizing loss and increasing board thickness for an FPGA accelerator card because the trace lengths are relatively small; unlike a backplane where a 40-inch trace is the norm.
- Considering the limited form factor of the accelerator card, choose a dielectric material like the Nelco N4000 -13/-13SI, TU-872LK, FR408HRIS, I-Speed, I-Speed IS, or EM-888 with Df = 0.0065 - 0.009 assuming a 4-mil wide trace in the 5-inch to 7-inch range from the FPGA to the module because the CAUI-4 specification for 100G Ethernet is roughly 11dB @ 14 GHz. Similarly, the PCIe specification for Gen4 and Gen3 standards is upwards of 15 dB, which should be relatively easy to achieve given the short trace lengths expected on this type of card.

PCB Fabrication Technology

In addition to the material selection, the choice of PCB fabrication technology plays a key role in the number of signal layers needed for routing the accelerator card while meeting PCB thickness requirements. If your goal is to minimize the PCB layer count at the expense of cost, you can use advanced fabrication techniques like micro vias, blind vias, and buried vias in addition to thinner trace widths to achieve this goal. A brief description of various industry terms along with approximate cost adders on top of the standard PCB fabrication cost that can

be expected for high-volume manufacturing is presented in the following section for reference. Figure 11 shows the various via types.



X19331-092817

Figure 11: Types of Vias

Via Aspect Ratio

The via aspect ratio is the ratio of PCB thickness to the smallest unplated via drill hole diameter. This is used as a guide to ensure that the PCB fabricator does not exceed the mechanical capabilities of the drilling equipment. A via aspect ratio of 10:1 is fairly common with standard PCB fabrication. The via aspect ratio can be increased to 20:1 using advanced PCB fabrication while maintaining design for manufacturing (DFM) rules.

Back-drilled Vias

A back-drilled via is a through-hole via that has a portion of its length drilled out so that it is no longer conductive. This improves signal integrity because it removes an unneeded stub from the route. The typical cost adder for back drilling vias ranges from 5% to 10% of the total PCB fabrication cost.

Via-in-Pad

A via-in-pad is a via drilled directly beneath a pad. This removes the need for a separate metal trace (stringer) to be drawn in order to drop down a via. This can help with breakout routing and improved signal integrity at the expense of a higher board fabrication cost. The cost adder varies from 10% to 15% of the PCB fabrication cost and is dependent on the via aspect ratio.

Blind and Buried Vias

A buried via is located entirely inside the PCB and does not touch the top or bottom layers. A blind via either travels from the top or bottom layer to an inner signal layer. This frees room above or below for other routing, unlike a through-hole via which travels all the way from the top to the bottom layer. The cost adder for a blind or a buried via depends on the number of

different types of blind or buried vias that exist on the PCB. Each type of blind or buried via requires a separate lamination cycle, resulting in extra cost. For example, a PCB with three different types of blind or buried vias (L1 to L4, L16 to L12, L4 to L8) on a 16-layer PCB results in a +30% cost adder for each type of blind or buried via. The use of blind or buried vias is often called *buildup construction* or *high-density interconnect* (HDI).

The primary limitation of blind vias is the difficulty of plating copper in them that is thick enough and properly bonded to the inner layer to which they connect [Ref 5]. The reliability of the plating becomes an issue when the depth of the blind via exceeds the diameter of the via (via aspect ratio limitation). Hence, most PCB fabrication shops recommend that the via diameter be at least 1.5 times the depth or greater to achieve a reliable plating. This limits blind vias to connecting only the top two layers of most PCBs. The blind vias must be stacked for connections below layer 2. This requires additional steps as part of the PCB fabrication process like "button" plating, which involves filling the blind vias with copper and using a sanding process to remove the excess copper that protrudes from some of these blind vias. This results in increased PCB fabrication cost.

Micro Vias

A micro via is a via whose diameter is less than 8 mil (0.2 μm). They are usually laser-drilled and typically cannot penetrate more than one or two layers at a time. The cost adder is approximately 15% for each type of via.


Note: Blind vias are often erroneously called micro vias.

Component Placement Guidelines

This section provides component placement guidelines to help you design a best-in-class FPGA-based accelerator card. A typical block diagram for an accelerator card as shown earlier in [Figure 4](#) comprises multiple DDR4 x64/x72 interfaces, multiple QSFP28 ports, and a PCIe X16 interface.

Careful upfront planning is needed in terms of component placement to successfully route all the high-speed interfaces, considering electrical, thermal, and mechanical requirements. At times, the keepout areas for heat sink attachments and back-brace requirements can be at odds with signal integrity routing requirements, so these factors must be considered early, during the planning phase and before the actual layout. All Kintex and Virtex UltraScale and UltraScale+ device architecture facilitates easy PCB breakout for high-speed transceivers and DDR4 interfaces. [Figure 12](#) shows a pinout for the VU9P device in a B2104 package. This pinout, like that of other Xilinx packages, is organized so that the high-speed transceivers are bonded out on the east/west side of the device with the SelectIO™ BGA balls on the north/south of the

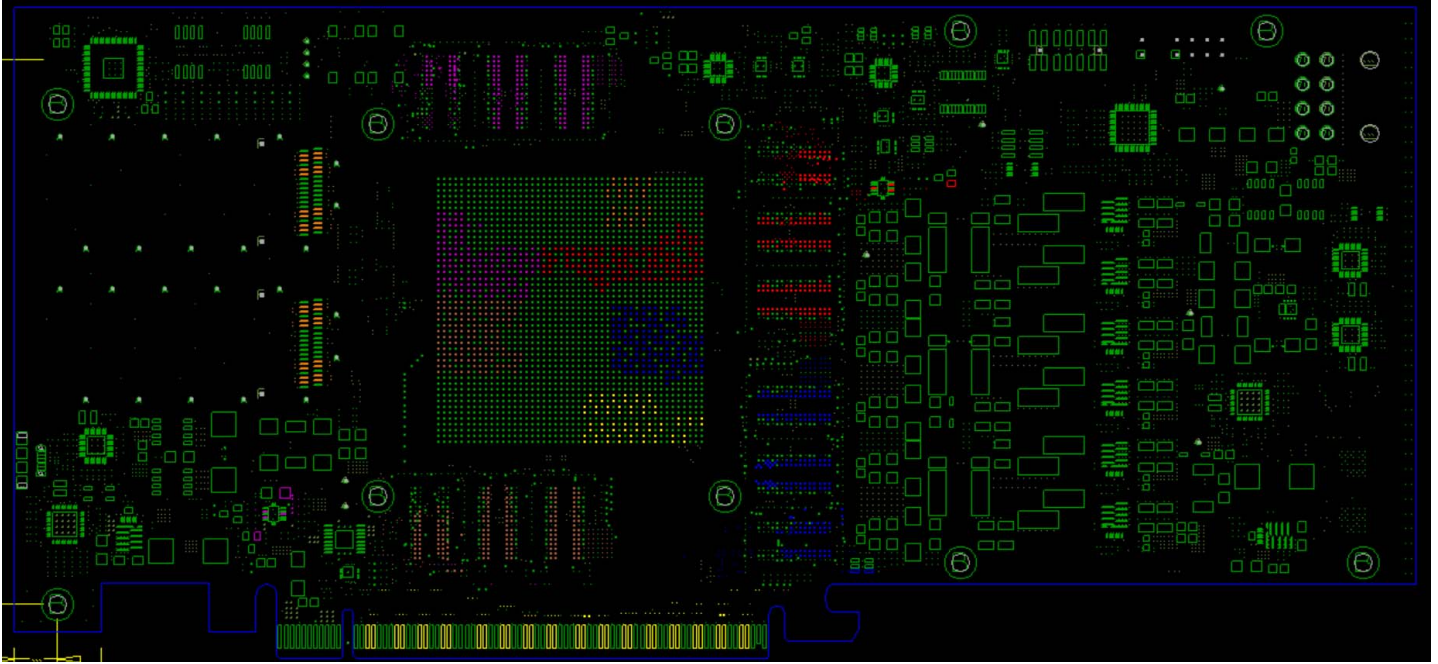
device. The V_{CCINT} BGA balls that power the FPGA fabric along with other digital rails are bonded out in the middle of the device.



X19333-100417

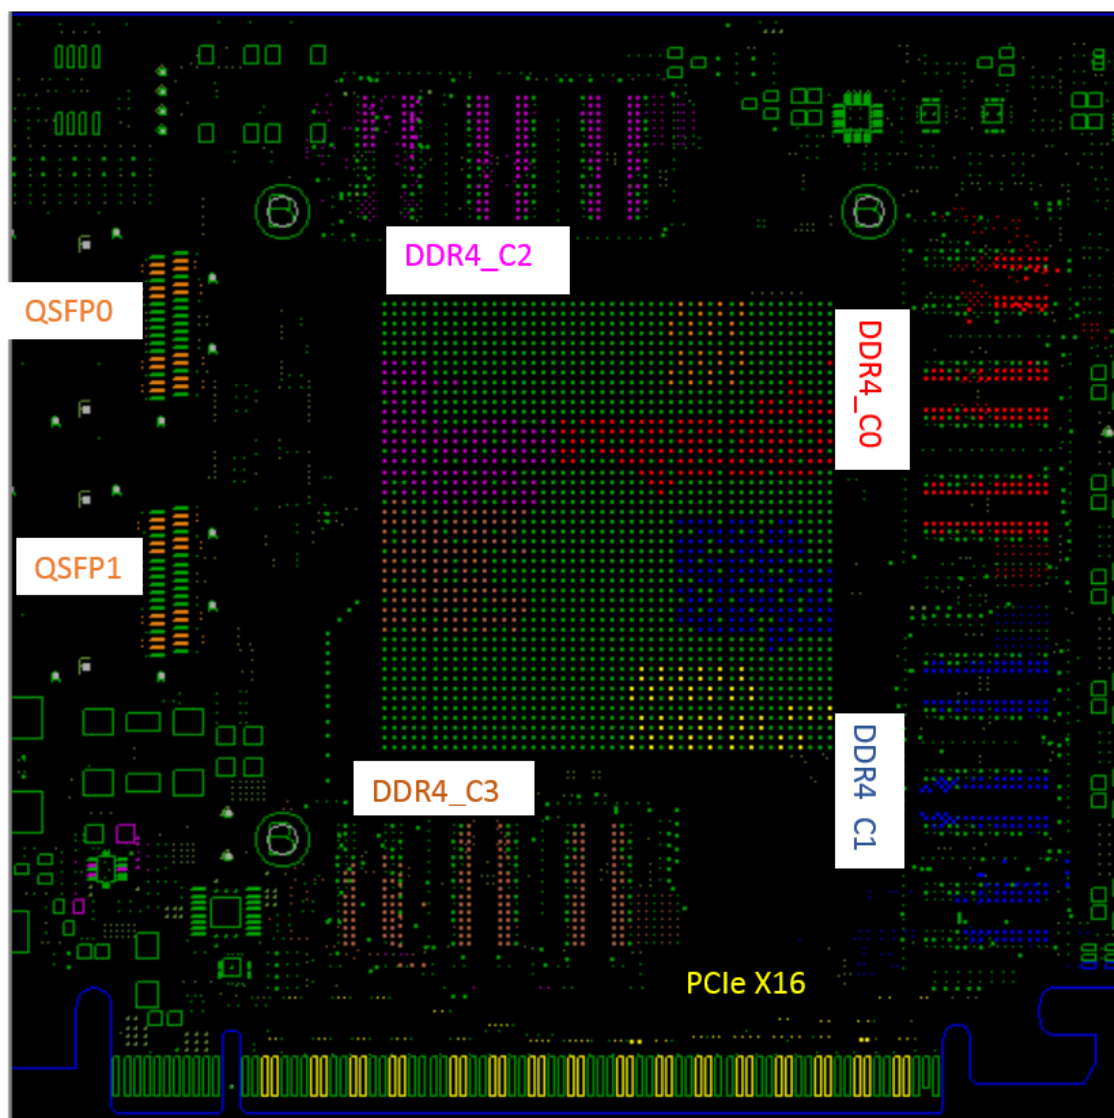
Figure 12: BGA Ball Locations for DDR4 and QSFP28, PCIe

Figure 13 shows the top view from an actual layout for a three-quarter FPGA accelerator card using the above device. Figure 14 is a closeup view for the FPGA device highlighting the various high-speed interfaces on the accelerator card. This PCB has four x72 DDR4 channels operating at 2400 Mb/s along with 2 x100G Ethernet ports. It has a x16 PCIe interface to interface to the host.



X19334-070517

Figure 13: Sample PCB Layout, Top View, Three-Quarter Length Accelerator Card



X19335-070517

Figure 14: Close-up of FPGA to Highlight BGA Ball Locations

The FPGA device on this board is oriented so that the GTY transceiver BGA balls are in the north/south direction and the I/O BGA balls are in the east/west direction with BGA ball A1 at the bottom left. Respective FPGA BGA balls for a given high-speed interface are highlighted in various colors.

A quick way to estimate the number of routing layers required to fully break out the I/O BGA balls and transceivers from the FPGA is to use [Equation 8](#). The signal pins can be assumed to be approximately 40% to 45% of the total number of BGA balls for the Virtex UltraScale+ FPGAs, with the remaining BGA balls designated for power and ground.

$$\text{Layers} = \text{Signal Pins (I/Os, MGTs)} / \text{Routing Channels} \bullet \text{Routes per Channel}$$

Equation 8

Routing channels are the total number of available routing paths out of the FPGA, for example, ((number of BGA balls on one side – 1) x 4 sides). Routes per channel are either one or two,

depending on the number of traces that can be routed between the BGA pads on the top/bottom layer. The routes per channel on the inner layers depend on the spacing between the vias, taking into account the drill to copper specifications. A quick calculation shows that at least six routing layers are needed to break out the various I/O BGA balls and transceivers on this device. A 14- or a 16-layer card should suffice to meet the above signal layer requirement while ensuring the overall PCB thickness stays within 1.57 mm (0.062 in.) \pm 0.13 mm (0.005 in) to comply with the specification.

The next section describes specific recommendations for laying out a DDR4 QSFP interface and optimizing the power delivery network design.


DDR4 Layout Guidelines

The majority of the routing real estate on an FPGA accelerator card is needed to route the DDR4 interfaces. Based on the application requirements, DDR4 memories are connected to the FPGA as either a set of discrete SDRAMs or as a DIMM module. Regardless of the topology, successful operation of the DDR4 interface at the highest possible data rate depends on ensuring robust signal and power integrity. This means minimizing crosstalk, signal reflections due to impedance discontinuities, and skew due to package flight time differences, and optimizing decap placement for V_{TT} and V_{DDQ} rails. Xilinx has done extensive signal integrity analysis using a design of experiments approach by sweeping the various critical parameters like driver slew rate, driver impedance, trace width, trace length, spacing, load capacitance, and fly-by termination resistor values that impact the eye opening with a goal to optimize the eye opening for both read and write scenarios across process, voltage, and temperature (PVT) corners. The memory guidelines are listed as part of the *UltraScale Architecture PCB Design User Guide* (UG583) [Ref 7] and cover both the device down and DIMM scenarios in detail. The memory layout guidelines cover critical aspects of PCB layout such as breakout routing requirements, maximum via count requirements, and trace length and spacing requirements for various signal groups that make up the memory interface. The guidelines provide a way to route the DDR4 interface using trace lengths as long as 6 inches on the DQ lines along with up to 13 inches on the Address and Command signals (ADDR/CMD) by spacing out the signals to minimize crosstalk.

On a typical accelerator card, the FPGA and the DRAM components are placed very close to each other, causing some of the guidelines to not entirely apply. The typical trace length needed to route the DDR4 interface is not more than three inches for the DQ lines under this scenario. However, some resulting specific challenges include:

- Less space for signal breakout, increasing the potential for signal coupling between the traces
- Layout spread across multiple layers, increasing scope for via coupling
- Difficulty with impedance matching due to trace neck downs in the FPGA/DRAM region

Figure 15 shows routing on one of the inner layers of the 16-layer accelerator card for reference.

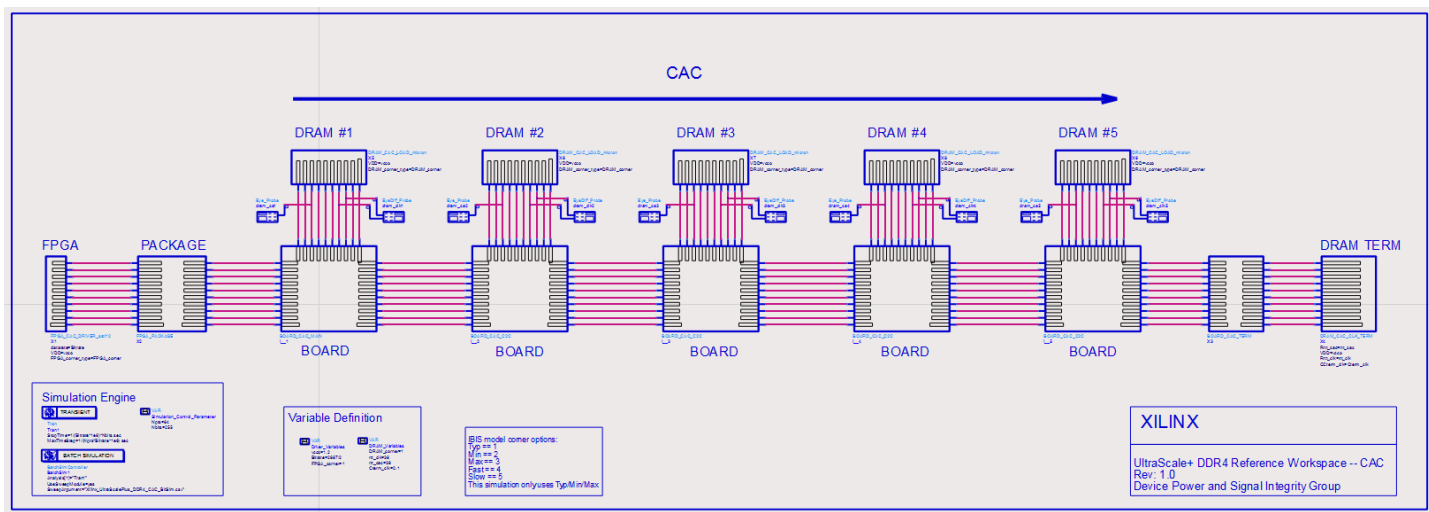


X19336-070517

Figure 15: Inner Layer Routing on a Typical 16-Layer Accelerator Card

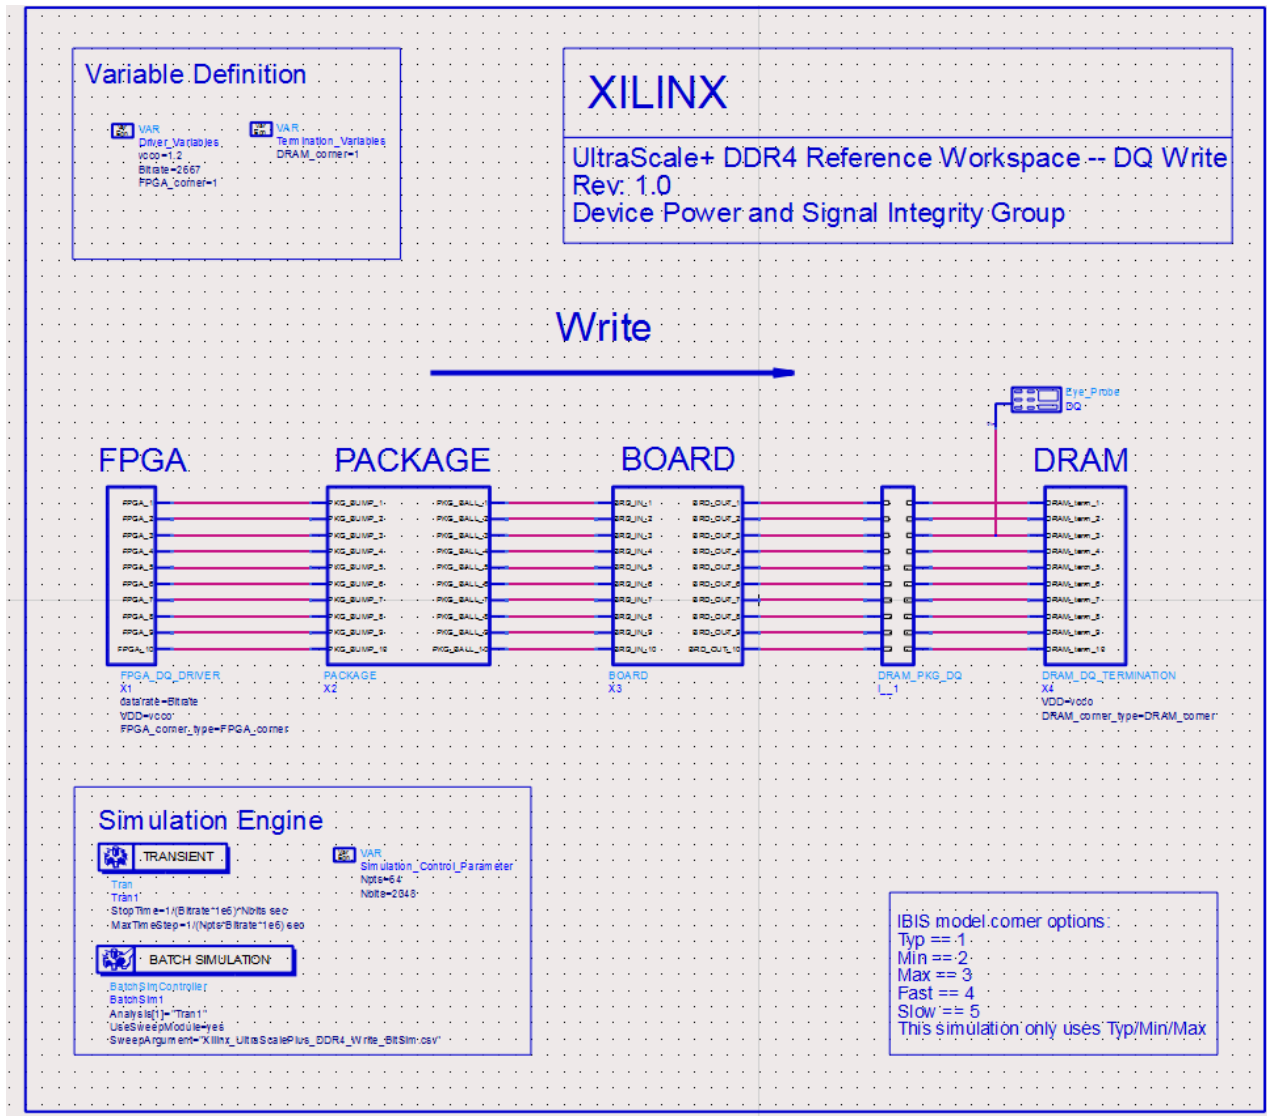
To help you with the DDR4 layout for the above scenario, Xilinx has developed an easy-to-use simulation methodology across a broad suite of EDA tools. An ADS workspace is readily available for the UltraScale and UltraScale+ families with two built-in topologies:

- Simulation workspace to simulate a group of ADDR signals with regard to the CK signal (Write). See [Figure 16](#).
- Simulation workspace to simulate a DQ byte group with regard to the DQS strobe (Write/Read). See [Figure 17](#).



X19337-070517

Figure 16: Simulation Workspace for Command Address Signals

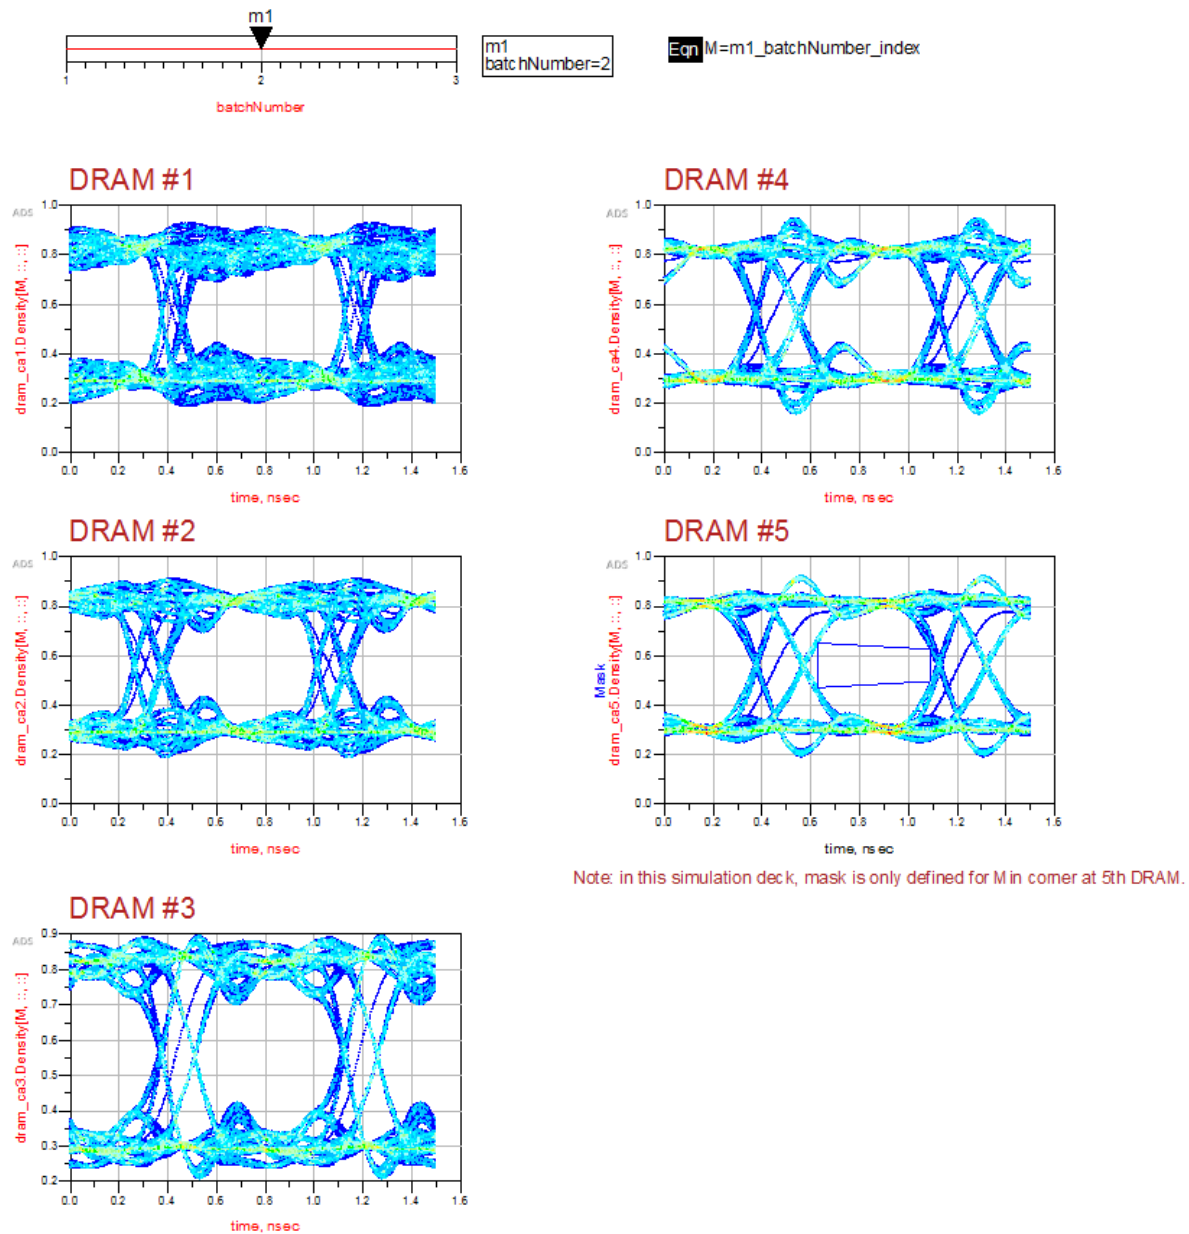


X19338-070517

Figure 17: Simulation Workspace for Data (DQ) Signals


The purpose of both these workspaces is to provide a reference to evaluate a given PCB layout against the routing guidelines listed in *UltraScale Architecture PCB Design User Guide* (UG583) [Ref 7]. The simulation methodology is fairly straightforward, where you run the reference simulation as is to generate the baseline eye diagrams with the default buffer settings, PCB model, and data pattern that is provided as part of the workspace. The second part of the simulation is to replace the default PCB model with the coupled S-parameter model extracted from your actual PCB layout with the byte group and ADDR signals of interest, and rerun the analysis to compare the eye opening with the baseline eye diagram. An eye opening equal to or greater than the baseline eye diagram means that the layout is satisfactory. The ADS workspace for both the DQ and ADDR assume coupled S-parameter package models for the FPGA and DRAM, along with a coupled PCB model to account for crosstalk, and is set up to sweep across the PVT corners. The workspace assumes ideal power and accounts for power delivery network (PDN) noise due to simultaneous switching outputs (SSOs) or simultaneous switching noise (SSN) as part of the eye mask requirements.

Figure 18 and Figure 19 show the baseline eye opening for the ADDR and DQ (write) workspaces, respectively.



X19339-070517

Figure 18: Baseline Eye Opening for Command/Control/Address Signals



X19340-070517

Figure 19: Baseline Eye Opening for Data Signals (Write)

Download these workspaces from the lounge:

- [UltraScale Signal and Power Integrity Lounge](#)
- [UltraScale+ Signal and Power Integrity Lounge](#)

In addition to Keysight's ADS simulation workspace, Xilinx has developed a HyperLynx DDRx wizard timing model for the DDR4 interface on the UltraScale+ family that is available upon request. These parameters can be used in conjunction with the PCB signal integrity analysis using LineSim/BoardSim tools from Mentor Graphics to determine the overall margin in your DDR4 interface.

Xilinx also provides hardware-correlated power-aware Input/Output Buffer Information Specification (IBIS) 5.0 SelectIO interface buffer models for UltraScale and UltraScale+ families to enable signal integrity analysis in a multitude of EDA tools capable of supporting the power-aware IBIS 5.0 models. These models can be requested from the [UltraScale+ SelectIO IBIS Models Lounge](#).


Some key recommendations to ensure a robust DDR4 layout from a signal and power integrity perspective include:

- Use the top signal routing layers for the DQ signals to minimize crosstalk due to via coupling.
 - Our analysis has shown at least a 30 ps channel timing degradation due to far end crosstalk (FEXT) for DQ nets routed on deep signal layers (~60 mil effective via length) compared to upper layer routing [\[Ref 8\]](#).
 - Give precedence to DQ - Data signals compared to ADDR signals, because ADDR signals operate at one-half the data rate of the DQ nets.
- Ensure that a given byte group along with the corresponding strobe is routed on the same layer to minimize skew.
- Breakout one DQ signal instead of two DQ signals between the BGA pads if you want to maintain a longer breakout length (approximately 1.5 inch) to minimize coupling in the via field region under the FPGA.

- Using larger antipads creates a slotted PWR/GND plane underneath the FPGA and the DRAM device.




RECOMMENDED: Do not route signals across the slotted ground planes as shown in [Figure 20](#) to avoid crosstalk issues. Reduce the size of the antipads to avoid this issue as shown in [Figure 20](#).



X19341-070517


Figure 20: Routing over Slotted Ground Plane

- Ensure adequate GND return vias to minimize crosstalk, especially in areas where there are fewer ground pins, such as near the address pins of the DRAM devices. [Figure 21](#) shows bad routing practice. [Figure 22](#) shows good routing practice.



X19736-122018

Figure 21: Bad Routing Practice



X19737-122018

Figure 22: Good Routing Practice

- For address/command/control V_{TT} termination, every four termination resistors should be accompanied by one 0.1 μ F capacitor, physically interleaving among resistors as shown in Figure 23.



X19342-070517

Figure 23: Recommended V_{TT} Termination Topology


SerDes Channel Guidelines

Some of the key recommendations for high-speed differential channels to ensure a robust layout from a signal and power integrity perspective include:

- Length-match the P and N legs for a given differential pair when doing the layout for QSFP, PCIe interfaces.
- It is a good design practice to ensure that the stripline layers chosen for QSFP, PCIe interfaces are surrounded by GND planes on either side to minimize coupling.
- Route RX and TX on separate stripline layers, giving higher priority to RX channels over TX channels from a layer assignment perspective.
 - Route the RX differential pairs on one of the top signal layers assuming the FPGA is on the top layer. It is recommended to place the AC coupling capacitors on the top layer if they are needed for the RX channels. This ensures a smaller via transition from the AC

coupling capacitor to the RX channel thereby minimizing the impedance discontinuity on the RX channel helping with the return loss. It is also recommended to back-drill the vias from the bottom layer under this scenario. The same strategy should be implemented in a vice versa fashion if the FPGA is placed on the bottom layer.

- Route the high-speed differential channels in true differential fashion as shown in [Figure 24](#) using neck down traces with back jog for skew compensation when breaking out of the FPGA. It is also recommended to use via-in-pad technology when breaking out of the FPGA for high-speed QSFP channels operating at 28 Gb/s. The via-in-pad technology allows more space for breaking out the high-speed channels under the FPGA.



X19343-070517

Figure 24: Recommended Differential Channel Breakout Routing

- For a given AC coupling capacitor, it is recommended to use a smaller profile capacitor (0201 over a 0402 capacitor) to minimize the return loss on the RX channel. The smaller capacitor needs a smaller pad, thereby minimizing the pad capacitance resulting in a smaller impedance discontinuity with respect to the trace compared to a bigger capacitor. It is also recommended to route the via transitions to the AC coupling capacitors in true differential fashion on the P and N traces rather than in a single-ended fashion.
- It is a good design practice to remove all non-functional via pads for both the signal and GND vias to minimize the via capacitance. The via inductance can be reduced by minimizing the via stubs through back drilling. It is recommended to back drill the signal vias for all high-speed channels. The via impedance can be further optimized by adjusting the via antipad size through 3D simulation analysis along with ensuring a good return path by having the ground vias close to the signal vias. By optimizing the via impedance to closely match the trace impedance by reducing the via capacitance, inductance results in a significant improvement in the insertion and return loss performance across a wide range of frequencies.

Similar to the SelectIO interface pins, Xilinx also provides hardware-correlated GTH/GTY IBIS algorithmic modeling interface (AMI) buffer models for UltraScale and UltraScale+ families to enable high-speed channel design in various EDA tools capable of supporting IBIS AMI models. These models can be requested from the access-controlled [UltraScale+ IBIS AMI Lounge](#).


Power Delivery Design Considerations:

Ensuring robust power integrity is as important as ensuring good signal integrity for a successful board design. The number of PWR/GND layers that need to be allocated as part of the stackup design is dictated by the following considerations:

- Number of unique power rails that need to be decoupled on the FPGA
- DC IR drop considerations for a given rail
- Ripple noise requirements

A typical power delivery network (PDN) on the PCB is comprised of four main components as shown in [Figure 25 \[Ref 9\]](#).

- Voltage regulator module (VRM)
- PCB
- Package
- FPGA (silicon)




X19344-092717

Figure 25: Components of a Typical PDN System

When the logic circuitry on the FPGA switches, it creates a transient current that has to come from the external power supply through all the PDN components in [Figure 25](#). Each component of the PDN has a certain non-zero impedance associated with it. That impedance causes voltage

variations (voltage noise) as the transient current passes through the elements of the power supply. A power delivery network can be represented as a chain of equivalent lumped RLC circuits that correspond to impedances of the PDN components as shown in Figure 26.




X19345-092717

Figure 26: Equivalent Circuit Representation of a PDN

Each component of the PDN contributes to the on-chip voltage noise in different frequency bands as shown in Figure 27. A typical $Z(f)$ impedance profile for a given power rail as seen by the FPGA based on the interaction of the various PDN components is as follows:

- 1st peak - Resonance between the on-die capacitance and the inductance of the package
- 2nd peak - Resonance between the on-package decoupling capacitors and the PCB inductance
- 3rd peak - PCB capacitors and the combined inductance of board PDN components and the VRM




X19346-092717

Figure 27: Example of a PDN Impedance Profile

Ideally, you want the impedance profile to be flat and stay below the target impedance across the wide range of frequencies to ensure the voltage noise does not get greatly amplified around the PDN resonance frequencies. Your goal is to ensure the PCB is designed to maintain a flat impedance profile by keeping the magnitude of the 2nd and 3rd peak to the minimum, because the magnitude and the frequency of the 1st peak is controlled by the package design and the on-die parasitics (R_{die} , C_{die}) of the FPGA. The design decisions on the PCB impact the resonance frequency and magnitude of the 2nd and 3rd peak, with minimal impact on the 1st peak.

Figure 28 shows the key parameters to focus on when designing the power delivery network on the PCB.




X19347-092817

Figure 28: Key Design Parameters for PDN Design

The stackup design and layer assignment for PWR/GND layers play a key role in optimizing the various parameters. The total inductance a given decoupling capacitor sees can be divided into the following:

- Capacitor mounting inductance
- Plane spreading inductance
- BGA via inductance under the FPGA (dependent on the PWR/GND BGA balls for a given rail)

In addition to the inherent equivalent series inductance (ESL) listed in the capacitor data sheet, a decoupling capacitor experiences additional inductance depending on how it is mounted on the PCB as shown in [Figure 29](#).



X19901-122018

Figure 29: Capacitor Mounting Considerations on a PCB

A connecting trace length has a large impact on the capacitor mounting inductance and if used, should be as short and wide as possible. When possible, a connecting trace should not be used and the via should butt up against the land pad. Placing the vias to the side of the capacitor land pads or doubling the number of vias further reduces the capacitor mounting inductance.

The second major component of the inductance is the plane spreading inductance associated with the PCB PWR/GND planes. The PWR/GND planes work in pairs, with their inductance coexisting dependently with each other. The spacing between the power and ground planes determine the pair's spreading inductance. The closer the spacing, the thinner the dielectric, and the lower the spreading inductance. The spreading inductance is specified in units of picohenries (pH) per square. The square is a dimensionless quantity. If the PWR/GND plane is in the shape of a square, the length equals the width, and the ratio is always 1, independent of the length of a square, so it results in the same spreading inductance. The shape of a section of the plane, not the size, determines the amount of inductance. The approximate values of spreading inductance for different dielectric thickness are shown in [Table 10](#).

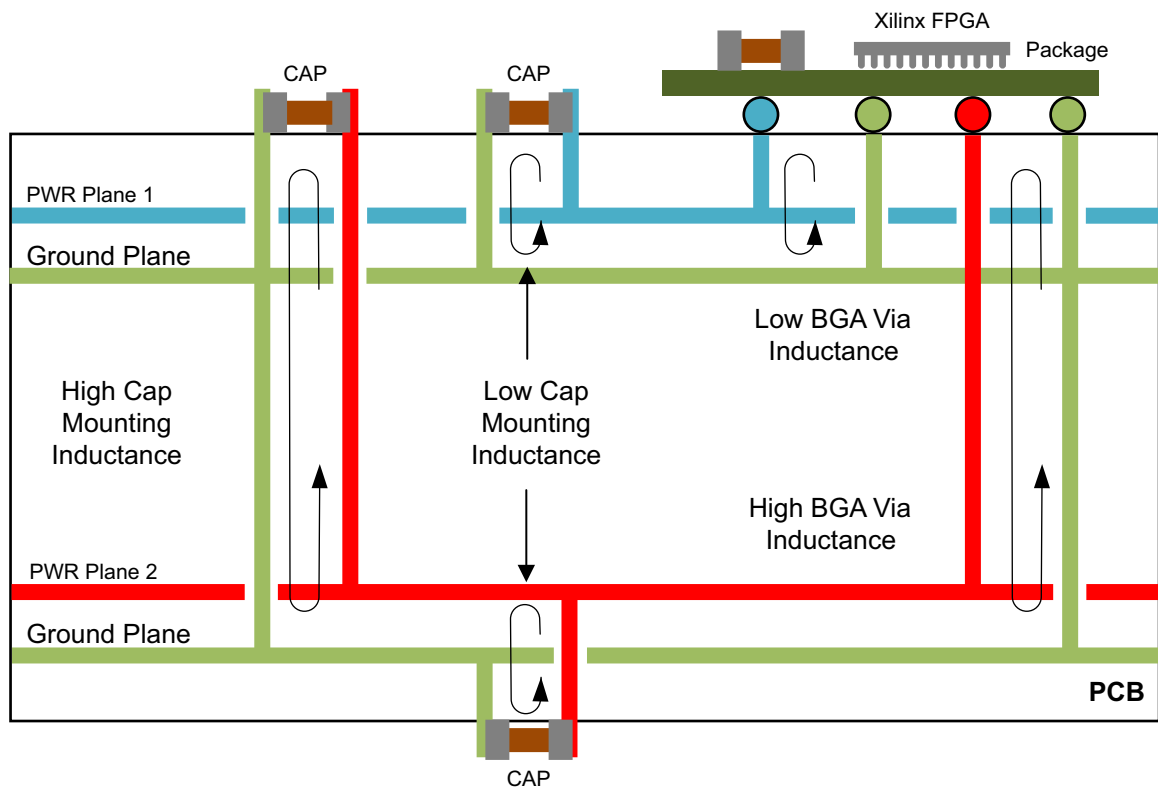
Table 10: Capacitance and Spreading Inductance Values for Different FR4 Dielectric Thickness Power Ground Sandwiches

Dielectric Thickness		Inductance	Capacitance	
(μm)	(mil)	(pH/square)	(pF/in ²)	(pF/cm ²)
102	4	130	225	35
51	2	65	450	70
25	1	32	900	140

The inter-planar capacitance aids in high-frequency decoupling and is inversely proportional to the dielectric thickness between the PWR/GND planes.

The third component in the overall loop inductance is the BGA via inductance. This is the effective inductance from all the vias in parallel underneath the FPGA footprint and is driven by the number of the BGA balls allocated for a given PWR/GND rail because the plated through hole (PTH) breakout on the PCB maps the FPGA pin map. It is also dependent on the overall PCB thickness and the layer assignment for a given PWR rail, because those factors determine the via height.

Figure 30 shows various scenarios for capacitor mounting and BGA via inductance, depending on the choices you make during PDN design.



X19348-082317

Figure 30: Different capacitor Mounting Scenarios

PDN Layer Priority Guidelines for Kintex and Virtex UltraScale and UltraScale+ FPGAs

- Assign the highest priority from a layer assignment standpoint to the sensitive transceiver rails by placing them close to the FPGA and having the decoupling capacitors on the top layer—assuming the FPGA is also on the top layer to minimize the loop inductance. It is recommended not to combine the transceiver power rails with the other non-transceiver digital rails to avoid noise coupling, because any noise coupled onto these rails has a direct impact on the transceiver jitter performance. The critical MGT rails like MGTAVCC and MGTAVTT have a significant number of on-package decoupling (OPD) capacitors, thereby limiting the number of capacitors needed on the PCB.
- You can assign second priority to the I/O bank power rail (V_{CCO}) for the DDR4 interface along with the associated power rails like the V_{TT} rail for the fly-by termination. The V_{CCINT_IO} rail should also be treated on par with the V_{CCO} rail if this rail is powered separately from the V_{CCINT} rail (assuming the device being used is a -2LE, -1LI where the V_{CCINT} has the flexibility to operate at V_{nom} (0.85V) or V_{low} (0.72V)).
- The V_{CCINT} power rail, which is typically the highest current-consuming rail, can be accorded a lower priority compared to the transceiver rails (MGTAVCC, MGTAVTT) and the V_{CCO} rail for the DDR4 interface. There are a few reasons for this:
 - Typically the V_{CCINT} rail on the UltraScale and UltraScale+ devices have a large number of BGA balls allocated for distributing the current. As an example, the number of BGA ball pairs for the V_{CCINT} rail on the VU9P_B2104 package is 78, which is fairly high. All the PTH via pairs in parallel under the FPGA result in very low BGA via inductance, irrespective of the plane placement on the PCB. It is good design practice, especially for high-current designs, to have the decoupling capacitors right underneath the FPGA footprint for this rail to take advantage of the low BGA via inductance, thereby avoiding the plane spreading inductance. It is recommended not to share the PWR/GND vias in the FPGA breakout region. The bulk capacitors can be placed beyond the FPGA footprint on the top/bottom layer depending on the plane location. In most cases, the PCB back-brace behind the heat sink that is recommended from a thermal perspective should not impede your ability to place decaps under the FPGA. However, this criteria should also be taken into account during the PDN solution for the V_{CCINT} rail.
 - The V_{CCINT} rail on UltraScale and UltraScale+ devices has a significant number of on-package capacitors, along with optimum amounts of die capacitance on the FPGA, thereby requiring you to focus on the mid-/low-frequency decoupling.
- Place other power rails like the V_{CCAUX} and V_{CCAUX_IO} along with any other remaining power rails wherever convenient and decouple accordingly by following the decoupling guidelines in the *UltraScale Architecture PCB Design User Guide* (UG583) [Ref 7]. Xilinx PCB capacitor requirements for V_{CCINT} and V_{CCO} and the other rails is in the *UltraScale Architecture PCB Design User Guide* (UG583). The GTH and GTY requirements can be found in the *UltraScale Architecture GTH Transceivers User Guide* (UG576) [Ref 10] and *UltraScale Architecture GTY Transceivers User Guide* (UG578) [Ref 11], respectively.

Figure 31 shows the layer priority recommendations for the Kintex and Virtex UltraScale and UltraScale+ families of devices.

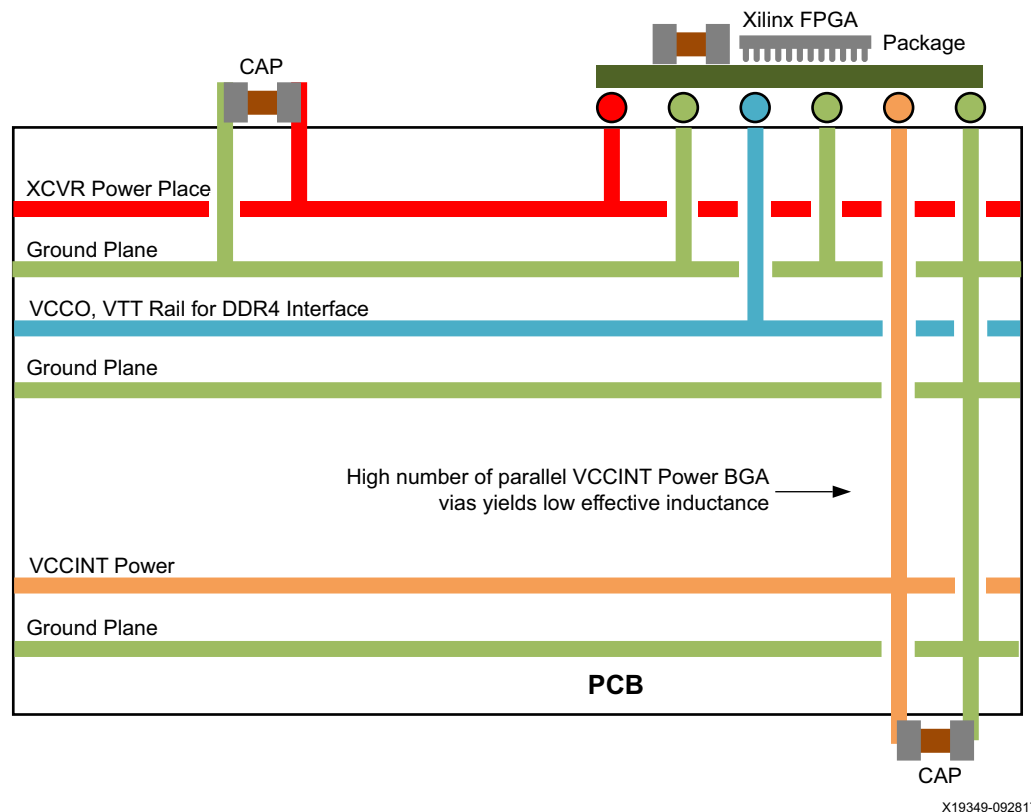


Figure 31: Layer Priority Recommendations for Kintex and Virtex UltraScale and UltraScale+ FPGAs


Thermal and Mechanical Guidelines

Considerations for the thermal management and overall mechanical design must also be a part of the overall considerations and sometimes trade-offs for the card. FPGA junction temperature must remain below the maximum specified by the data sheets, generally 100°C and thus often requires careful thought of the heat sink, heat sink mounting and attachment, and the associated airflow to ensure adequate thermal dissipation at the maximum power draw within the tolerances of the system.

Each thermal design is unique based on system requirements, expected power draw, device selection, and specific dimensions and layout of the card, so the purpose of this application note is not to convey a specific thermal design, but instead discuss best practices used for thermal and mechanical design.

The first step to understanding the thermal characteristics of a card is to estimate the maximum power draw of the devices. Xilinx provides a Xilinx Power Estimator (XPE) as a means for customers to perform early power estimations of a device based on the specified resource and performance metrics. The XPE tool and its documentation can be downloaded from www.xilinx.com/power.

In most cases, forced convection across a heat sink is used (the airflow is either provided by the server or a local fan). Based on expected airflow direction and velocity, an initial heat sink solution should be derived, maximizing air contact to the heat sink. Careful consideration of the mounting of the heat sink to the card should also be performed early because mounting holes, screws, and brackets often need keepout areas to accommodate them. It is suggested to have at least four mounting points using dynamic attachments such as spring-loaded push pins connected to a back brace mounting bracket as shown in Figure 32. This allows for an even distribution of force across the package while minimizing possible board flex and warping due to the heat sink attachment. It is also suggested to evenly and thoroughly apply a good quality, low thermal resistance thermal interface material (TIM) and employ between 20 and 40 psi of force to the package to ensure good thermal contact between the device and heat sink.



X19351-122018

Figure 32: Typical Heat Sink Mounting Considerations

After the initial thermal design is established, it is highly encouraged to model and use computational fluid dynamics (CFD) simulation of the system as a means to understand and estimate the maximum junction temperature of the FPGA based on the calculated air flow, device power consumption, device placement, heat sink characteristics, local heating due to other devices, as well as many other factors. Xilinx provides DELPHI multi-resistance thermal models at www.xilinx.com/power pre-compiled for ANSYS Icepak and a Mentor FloTHERM. If another thermal design software package is used, the documentation within the models explain how to create a multi-resistance model in other software packages. Thermal simulation is often an iterative process of changing placements, and heat sink characteristics among other factors, to ensure the highest thermal margin within the constraints of the system.

Reference Design

Download the reference design files for this application note from these Xilinx websites:

- [UltraScale Signal and Power Integrity Lounge](#)
- [UltraScale+ Signal and Power Integrity Lounge](#)

Note: A non-disclosure agreement is needed to access the lounges.

[Table 11](#) shows the reference design matrix.

Table 11: Reference Design Matrix

Parameter	Description
General	
Developer name	Ravindra Gali
Target devices	Virtex and Kintex UltraScale and UltraScale+ FPGAs, Zynq UltraScale+ MPSoCs
Source code provided	Yes
Source code format	IBIS models, simulation workspace
Design uses code and IP from existing Xilinx application note and reference designs or third party	No
Simulation	
Functional simulation performed	No
Timing simulation performed	Yes
Test bench used for functional and timing simulations	N/A
Test bench format	IBIS models, S parameter files
Simulator software/version used	Any EDA simulator
SPICE/IBIS simulations	Yes
Implementation	
Synthesis software tools/versions used	N/A
Implementation software tools/versions used	N/A
Static timing analysis performed	N/A
Hardware Verification	
Hardware verified	N/A
Hardware platform used for verification	N/A

References

1. *PCI Express Card Electromechanical Specification, Revision 3.0* (pcisig.com)
2. Bogatin, Eric. "Loss in a channel: Rule of Thumb #9." EDN Network. April 4, 2014. Retrieved from www.edn.com/electronics-blogs/all-aboard-/4429778/Loss-in-a-channel--Rule-of-Thumb--9
3. Liew, Elaine. Alam, Shah. Okubo, Taka-aki. Sudo, Toshio. Hosoi, Toshihiro. Tsuyoshi, Hiroaki. Kuwaki, Fujio. *As originally published in the IPC APEX EXPO Conference Proceedings. "Signal Transmission Loss due to Copper Surface Roughness in High Frequency Region."* Retrieved from www.circuitinsight.com/pdf/signal_transmission_loss_copper_surface_roughness_ipc.pdf
4. McMorrow, Scott et al, DesignCon 2005. "The Impact of PCB Laminate Weave on the Electrical Performance of Differential Signaling at Multi-gigabit Rates."
5. Ritchey, Lee W. "Materials Manufacturing, PCB Fabrication, PCB Stackup Design," DesignCon 2017
6. Marshall, John A. "Measuring Copper Surface Roughness for High Speed Applications." Retrieved from www.circuitinsight.com/pdf/measuring_copper_surface_roughness_ipc.pdf
7. *UltraScale Architecture PCB Design User Guide* (UG583)
8. Wang, Yong et al. DesignCon 2015. "UltraScale FPGA DDR4 2400 Mbps System Level Design Optimization and Validation." Retrieved from www.xilinx.com/publications/events/designcon/2015/we2paper-ultrascale-ddr4.pdf
9. Klokotov, Dmitry. Shi, Jin. Wang, Yong. DesignCon 2014. Xilinx. "Distributed Modeling and Characterization of On-Chip/System Level PDN and Jitter Impact." Retrieved from www.xilinx.com/publications/events/designcon/2014/1_WE5Paper_DistributedModelingan dCharacterizationofOnChipSystem.pdf
10. *UltraScale Architecture GTH Transceivers User Guide* (UG576)
11. *UltraScale Architecture GTY Transceivers User Guide* (UG578)
12. Mirshafiei, Sean. Isola. Enos, Dan. Northrop Grumman, Springfield, MO. "Signal Integrity Analysis Techniques Use to Characterize PCB Substrates." Retrieved from www.isola-group.com/wp-content/uploads/Improved-High-Speed-Low-Loss-Materials-for-Lead-Free-Assembly-Compatibility.pdf

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
10/05/2017	1.0	Initial Xilinx release.
12/21/2018	1.1	Updated second paragraph in Blind and Buried Vias . Updated Figure 27 title.

Please Read: Important Legal Notices

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <https://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <https://www.xilinx.com/legal.htm#tos>.

AUTOMOTIVE APPLICATIONS DISCLAIMER

AUTOMOTIVE PRODUCTS (IDENTIFIED AS "XA" IN THE PART NUMBER) ARE NOT WARRANTED FOR USE IN THE DEPLOYMENT OF AIRBAGS OR FOR USE IN APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE ("SAFETY APPLICATION") UNLESS THERE IS A SAFETY CONCEPT OR REDUNDANCY FEATURE CONSISTENT WITH THE ISO 26262 AUTOMOTIVE SAFETY STANDARD ("SAFETY DESIGN"). CUSTOMER SHALL, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE PRODUCTS, THOROUGHLY TEST SUCH SYSTEMS FOR SAFETY PURPOSES. USE OF PRODUCTS IN A SAFETY APPLICATION WITHOUT A SAFETY DESIGN IS FULLY AT THE RISK OF CUSTOMER, SUBJECT ONLY TO APPLICABLE LAWS AND REGULATIONS GOVERNING LIMITATIONS ON PRODUCT LIABILITY.

© Copyright 2017–2018 Xilinx, Inc. Xilinx, the Xilinx logo, Artix, ISE, Kintex, Spartan, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. PCI, PCIe, and PCI Express are trademarks of PCI-SIG and used under license. All other trademarks are the property of their respective owners.