Speedster7t AC7t1500 Board Designers Guide (UG101)

Speedster FPGAs

Preliminary Data

Achronix

Data Acceleration

Copyrights, Trademarks and Disclaimers

Preliminary Data

This document contains preliminary information and is subject to change without notice. Information provided herein is based on internal engineering specifications and/or initial characterization data.

Achronix Semiconductor Corporation

2903 Bunker Hill Lane
Santa Clara, CA 95054
USA

Website: www.achronix.com
E-mail : info@achronix.com

Table of Contents

Chapter - 1: Introduction

The Speedster®7t AC7t1500 FPGA includes several advanced interfaces that require careful design in order to operate at their peak performance. This guide is intended as a general overview of PCB design principles that help the designer get the most out of the AC7t1500 FPGA.

This guide is broken down by system components. These include the Ethernet, the PCIe5, the GDDR6 memory and the DDR4 memory interfaces. In addition this document reviews supporting components such as GPIO, reference clocks and power regulation and distribution.

- Ethernet (28G, 56G, 112 Gbps)
- PCIe5 (32 Gbps)
- GDDR6 (16 Gbps)
- DDR4 (3200 MTps)

- GPIO, RefClks
- Power and PDN

The diagram below illustrates the separate interfaces (green rectangles in the figure below) discussed in this guide.
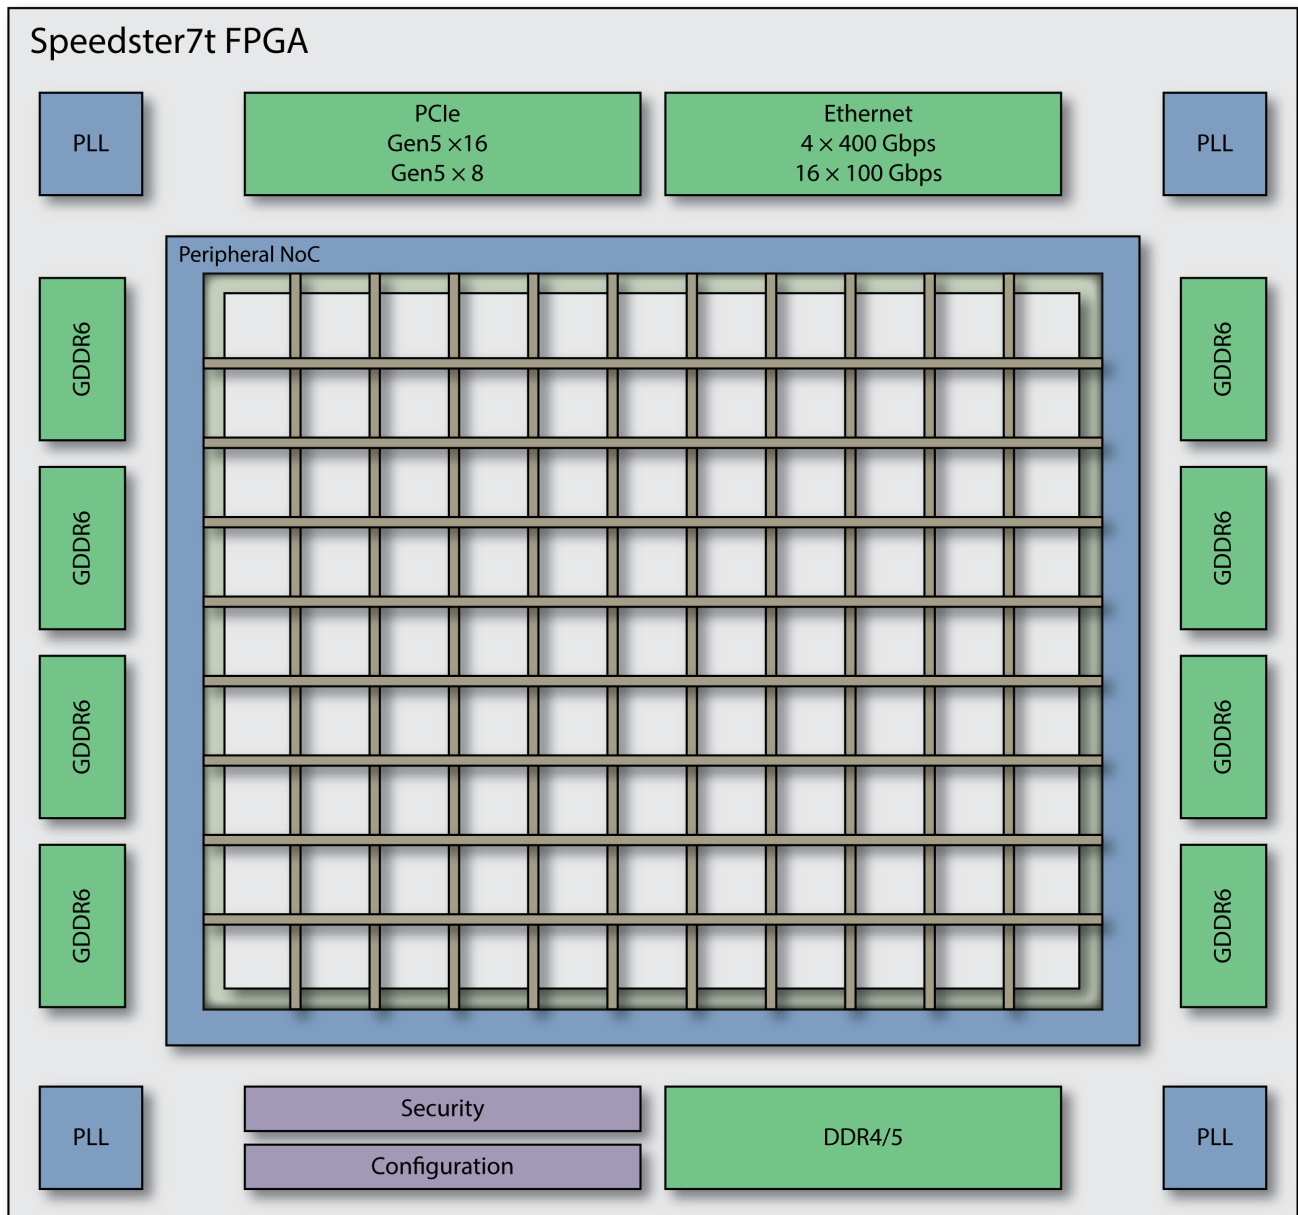
Figure 1: *Speedster7t AC7t1500 FPGA Block Diagram*

Chapter - 2: PCB General Considerations

Much work has been done in the industry to assure successful design implementations of high-speed circuits, while maintaining reasonable material costs. This chapter does not presume to catalog all the best knowledge, but rather to touch on subjects which the industry has identified as best practices for high-speed PCB design.

Board Construction - the Stack-up

The stack-up is the fabrication plan for the PCB. Stack-up design is jointly agreed upon between the designer of the PCB and the fabricator. In designing and evaluating the different available stack-ups for a given layer count, it is crucial to work with a qualified PCB fabricator who has worked with low-loss materials and can provide options for those materials. A good fabricator provides all the technologies needed (via-In-pad, back-drill, fine lines, tight impedance control, etc.) and accurately evaluates the dimensions required to achieve the target impedance. At high speeds, layer-to-layer impedance control plays a vital role in overall signal quality due to reflections from discontinuities in the channel. Using multiple layers increases the board's ability to distribute energy, reduces crosstalk, controls electromagnetic interference and generally supports high-speed signals.

The figure below shows a portion of an example stack-up. Stack-ups vary for each design and application, but are generally symmetrical.

Table 1: *Example Stack-up Section for a High-Speed PCB*

Layer #	Layer Name	Cu Weight	Thickness	Reference Layer	Material	Dielectric Constant	Dissipation Factor	Trace Width/Spacing (mils)			
		(oz)	(mils)			(DK)	(DF)	85Ω DIFF	90Ω DIFF	100Ω DIFF	50Ω SE
	DIELECTRIC_1		1		Solder mask	3.5	0.019				
L1	TOP	0.25 oz	2	L2	Copper, plated						
	DIELECTRIC_2		3.3		*<material name>*	3.4	0.0016				
L2	GND_1	0.25 oz	0.35		Copper						
	DIELECTRIC_3		4		*<material name>*	2.97	0.0014				
L3	SIGNAL1	0.25 oz	0.35	L2, L4	Copper			4.7/4.0	4.1/3.5	3.6/4.0	4.0
	DIELECTRIC_4		4.1		*<material name>*	2.97	0.0014				
L4	GND_2	0.25 oz	0.35		Copper						
	DIELECTRIC_5		3.7		*<material name>*	3.06	0.0017				
L5	SIGNAL2	0.25 oz	0.35	L4, L6	Copper			4.2/4.2	4.0/5.0	3.0/4.95	5.0
	DIELECTRIC_6		4.8		*<material name>*	2.97	0.0014				
L6	GND_3	2.0 oz	2.8		Copper						

> **Table Note**
>
> - Each copper layer, as well as the dielectric between those layers, is described in full, along with the thicknesses and trace dimensions required for each layer.

Stack-up Planning

The selection of a stack-up should take into account the following:

- The types of signals used and their required loss and impedance characteristics
- The number of layers required to route the signals with minimal layer changes
- Sorting and sequence of layers
- Spacing between layers.
- The power delivery needs of the AC7t1500 FPGA
- Mechanical requirements, especially thickness, but also board warpage, mounting and thermal requirements

Manufacturing Tolerances

An impedance manufacturing tolerance of ±10% should be standard for stripline traces. Some vendors can meet a tolerance of ±5%, which is beneficial for SerDes. Tighter tolerances affect manufacturing yield and thus cost.

Crosstalk

Crosstalk and characteristic impedance are interrelated. Crosstalk minimization is accomplished by ensuring the impedance is defined mainly by the distance to the reference plane. As a result, the distance between adjacent traces must be at least 2× the distance of the trace above and below the ground plane. It is recommended that the designer determine the ideal trace separation through modeling,

Microstrip

Signals routed on the surface of the PCB (microstrip) are impacted by the presence of air, which has a different dielectric constant. Different modes of transmission (even, odd and uncoupled) thus have different delays, which impact timing and, ultimately, transfer rate. Stripline (inner layer) construction provides the most satisfactory delay characteristics, and allows closer routing of the traces.

PCB BGA Pad Design

For PCB BGA pad design, consider via-in-pad (VIP) technology to eliminate the inductance of the pin-to-trace connection, to simplify the via structures and better support them with accompanying ground vias. The reflective impact of via stubs increases with frequency and board thickness. An analysis of the via resonance is important to the performance of the overall channel. Via stub mitigation techniques can include:

- Using laser vias to eliminate the stub.
- Using plated through-hole (PTH) vias to send the signal to a lower layer in the board.
- Placing components on opposite sides of the board.
- Back-drilling PTH vias to minimize the stub. Be careful to obtain maximum stub length information from the fab vendor.

Power Distribution

The power distribution network (PDN) for the AC7t1500 FPGA benefits greatly from PWR-GND power plane pairs integrated into the stack-up. These power plane pairs provide local distributed charge wells to support fast transients in the power rails and can also help to reduce the mounting inductance of the decoupling capacitors on each rail. The AC7t1500 FPGA is capable of drawing a large amount of power on multiple rails, depending on the final application. While PDN design is covered in the chapter, Speedster7t AC7t1500 PCB Power and PDN Design (see page 65), it bears repeating that the stack-up must take into account the projected current requirements and provide enough copper to deliver that current with minimal voltage drop. These requirements have an impact on:

- Power efficiency
- Voltage regulator stability
- Copper heating effects

Board Thickness

It is also important to consider the contribution of each layer to the overall thickness of the board, which may be a hard mechanical constraint. There are standard thicknesses of different types of printed circuit board material, each with specific properties which may help when beginning with stack-up review.

Dielectric Materials

For routing very high-speed signals (25 Gbps+), dielectric materials with very low loss may help to control the overall system loss. Variants of Tachyon and Megtron-7 are some of the popular choices for low-loss materials for routing high-speed SerDes signals. Get in touch with your PCB manufacturer to find out the best solution based on your needs.

Fiber Weave

Fiber weave refers to the glass fabric used in the construction of the PCB. Because these glass fibers have a much lower dielectric constant than the surrounding resin, they lower the effective dielectric constant. This situation creates local variations in impedance as an electromagnetic field in a trace passes past it, impacting high-speed signals in two ways. If the trace is routed perpendicular to the weave, it crosses many bundles of fibers at a regular spacing, creating a periodic disturbance which can filter out a narrow band of frequencies directly related to the weave separation. If multiple traces are routed along the fiber bundle, one trace might be closer to the fiber bundle and thus experience a lower dielectric constant. The signals in this trace go faster and timing can be affected. If one side of a differential pair is beside the fiber, and the other side is directly in between two fibers. The two signals of the differential pair experience different propagation delays and common-mode noise is introduced.

There are four possible solutions to the fiber weave problem:

- Rotate either the design or the PCB panel 15%
- Route signals for short distances at random angles.
- For differential pairs, match the weave pitch of the fabric to the pitch of the differential pairs to ensure both sides of the pair experience similar impedance profiles.
- A common solution is to use "spread" glass, which has the fiber bundles of the weave mechanically spread out to reduce the weave effect. Spreading is often the most effective way to reduce fiber weave effect.

Stubs

Stubs are unterminated structures in an electrical path. One of the common ways to create a stub in a PCB is with a through-hole via, where a connection is made from a trace inside the PCB to an outer layer. Since the via goes through the PCB, one side ends on the other side of the PCB without connecting to anything.

In the figure below, the signal splits into two components, A and B. A travels to the unterminated end of the via and reflects, where it either reinforces B or degrades it. Whether it reinforces it and to what degree is determined by the length of the stub. If the stub is equal to one quarter of the wavelength of the signal, component A is 180° out of phase with B, and they cancel each other. Since digital signals are composed of a range of frequencies, some of those frequencies are filtered out by the stub, and information is lost.

Via stubs can be mitigated in one of two ways:

- **Microvias** (also known as laser vias), are vias drilled with a laser to a specific depth (usually only one or two layers). Since these vias only connect one layer to the next successive one or two layers, they do not have stubs caused by back drilling limitations.

- **Back-drilling** is a process where a larger drill is used to drill out the via stub. While this method is effective, it has two drawbacks, First, it is not possible with current technology to fully remove the stub. Secondly the hole created is larger than the barrel of the via, so larger anti-pads are required. Anti-pads can reduce routing capability on the board and also reduce the amount of copper on the power planes, resulting in higher inductance and resistance in the power distribution network.
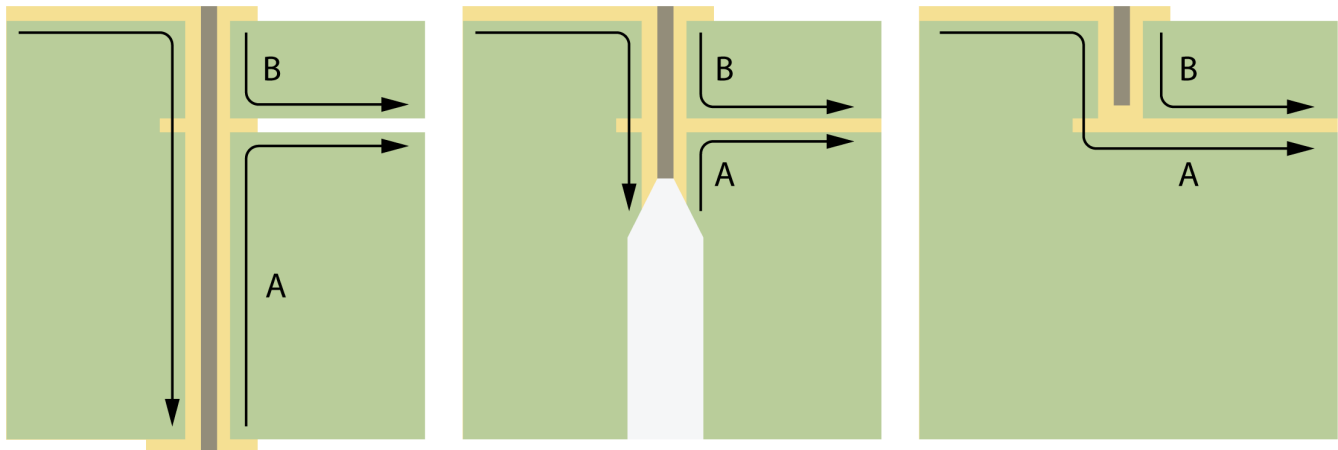
70536804-01.2020.12.14

Figure 2: *Via Stub Resonance and Mitigation by Back-drilling and Laser Vias*

Component Placement

The pin placement of the AC7t1500 BGA package has a major influence on the placement of components. The following figure, illustrating the view of the package from above, shows the locations of the pins for the SerDes quads N0 to N7 as well as the GDDR6 interfaces, the DDR4 interfaces, the FCU, the GPIO and the PLLs. It is clear that the best way to route the signals for all interfaces is directly out of the package.
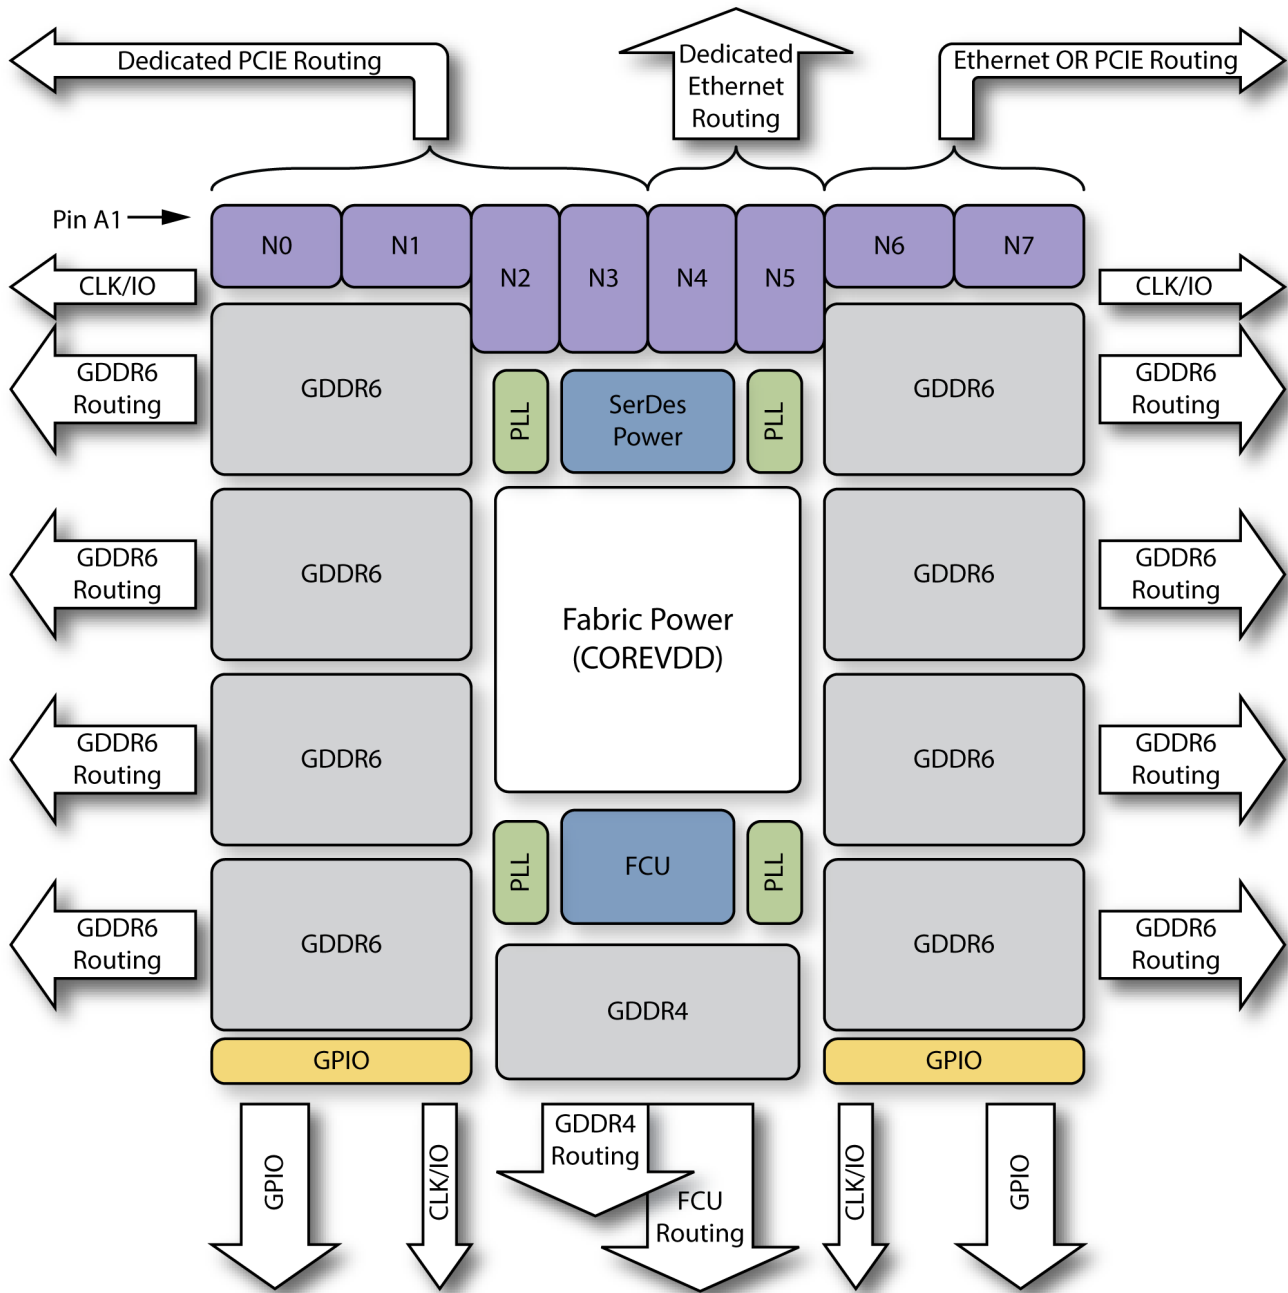
Figure 3: *Speedster7t AC7t1500 FPGA Pin Map and Routing*

The following figure shows one possible configuration of components. On the north side, two SerDes quads (N4 and N5) are connected to a QSFP-DD 1N1 connector, though it could be four QSFP28 connectors. Four quads are connected to a ×16 PCI Express slot and the last four quads (N6 and N7) route to a secondary PCIE slot, possibly for interfacing with an SSD. As well as four GDDR6 ports to the west and four to the east, the device is connected to one or two DDR4 DIMMs on the south side. This arrangement takes advantage of the natural signal flow from the FPGA.

External clock connections should lead into the corners to connect to one of the four PLL blocks. The FCU connections (JTAG or CPU) must be brought in under the DDR4 interface.
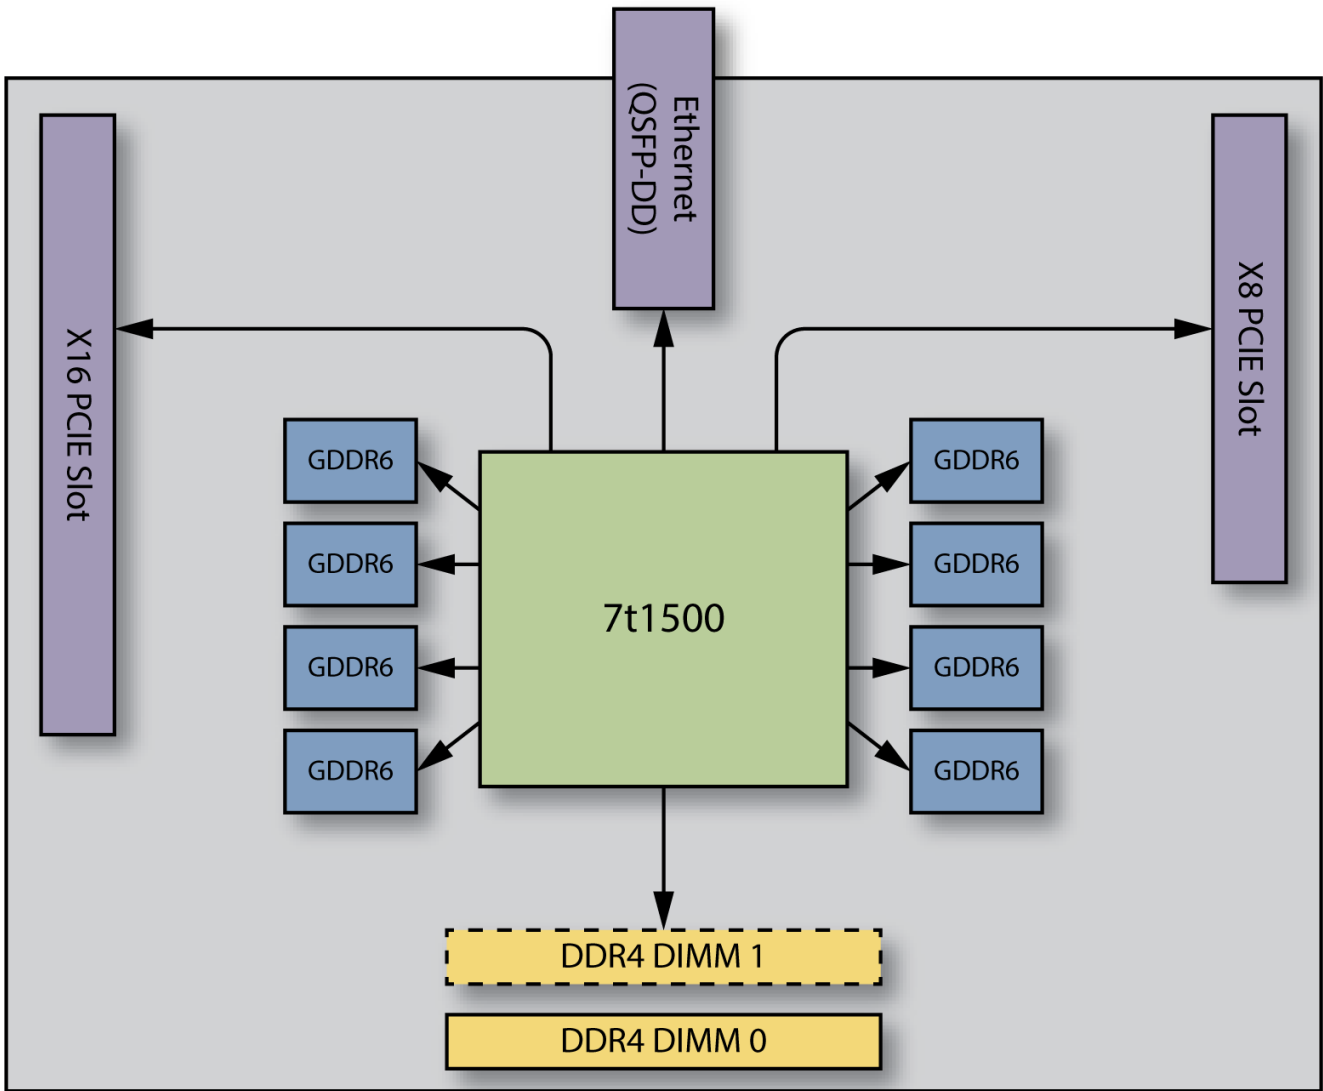
70536804-03.2020.12.14

Figure 4: *Component Placement with One QSFP-DD Ethernet Connector, Two PCIE Connectors, and Two DDR4 DIMMs*

The following figure converts the ×16 PCIE connection to a card-edge connector for a PCIE add-in card, and converts the secondary PCIE interface to a second Ethernet interface using a QSFP-DD. The GDDR6 does not change, but the DDR4 could also be a single SO-DIMM for a more compact layout. Component placements substantially different from these are possible but might require a more expensive PCB with more routing layers. Refer to the interface-related sections that follow for specific guidance on each interface.
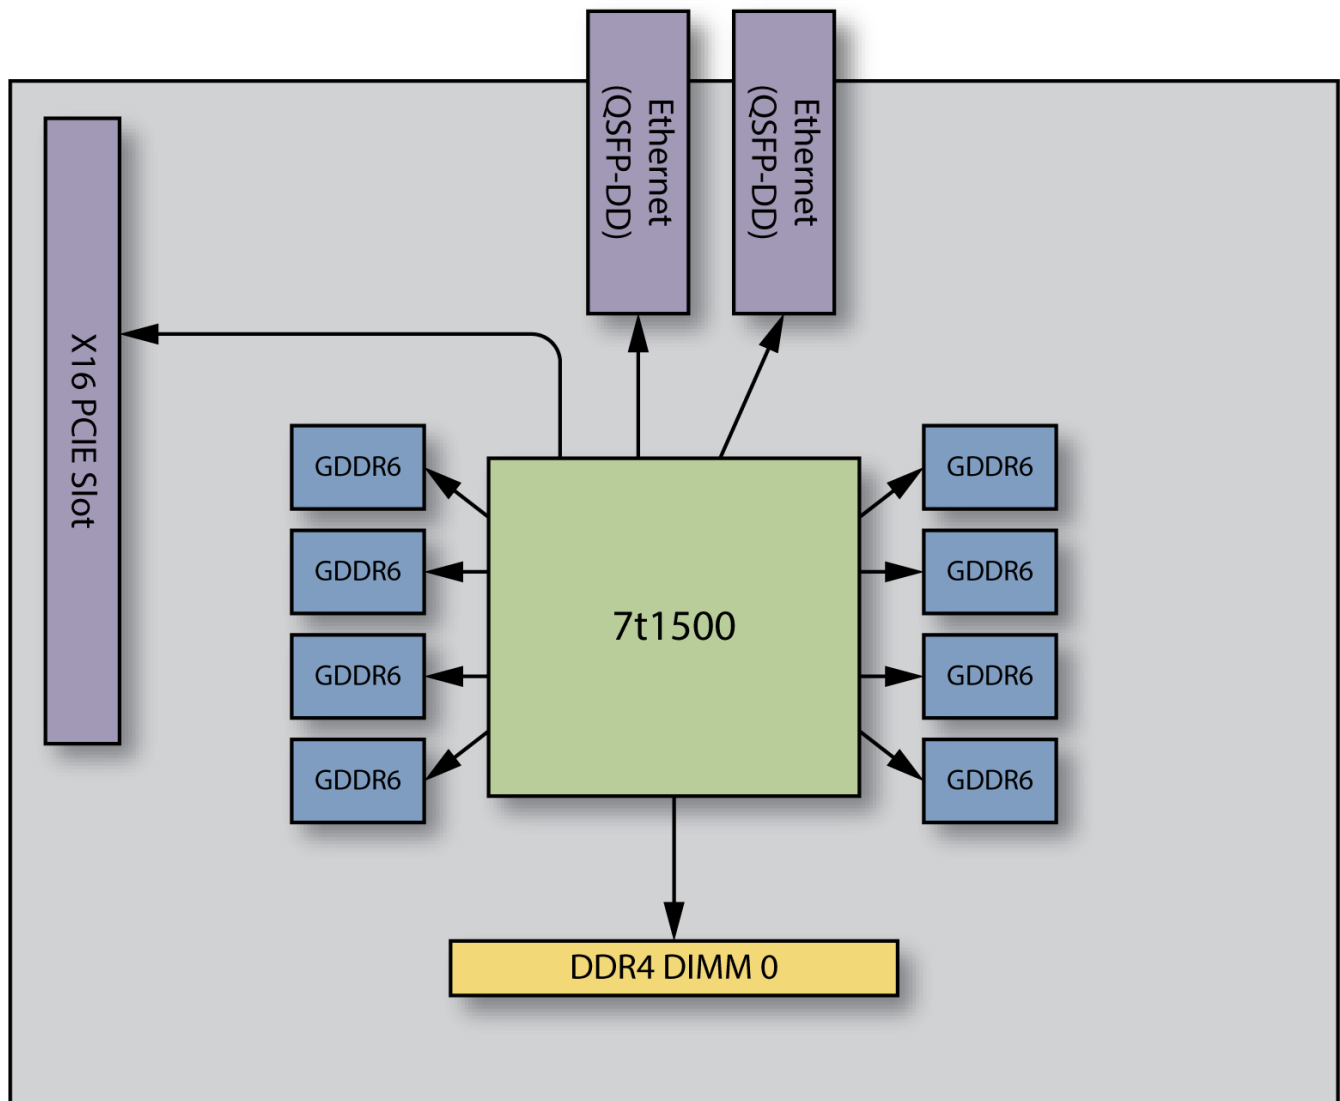
70536804-04.2020.12.14

Figure 5: *Component Placement with Two QSFP-DD Ethernet Connectors, One PCIE Connectors, and One DDR4 SO-DIMM*

Routing Guidelines

The following general guidelines should be followed when routing any PCB:

- Provide sufficient return vias in proper proximity to power vias to reduce the power delivery network loop inductance. Optimize signal via transitions to nominal impedance (50/100 Ω). This operation might require the use of a 3D electromagnetic field solver.

- Ensure traces have a solid reference ground plane over the entire length without interruption in ensuring no impedance discontinuity.

- Return-path discontinuities must be minimized to avoid reflections. Ensure a return current path as close to the signal such that the current loop remains as small as possible.

- Avoid 90^o turns; instead use 45^{o-} angles or curves in traces.
- If a signal transitions from one layer to another; then ensure a good ground reference during the transition, and that the different layers ground are connected. Power and ground return paths must be kept clear of splits or voids that can interrupt returns for differential pairs, and measured lines need to be within their tolerances.
- Differential signals need to be length-matched within the pair for the complete length in the channel. Signal traces should be designed to minimize skew between P and N traces of a differential pair. Limiting the skew to be less than 0.5% of a bit time is recommended. It is also critical to maintain symmetry between the true and complement trace of the differential pair to minimize mode conversion and skew.

Refer to the specific interface specification for length-matching requirements and additional routing guidelines.

Layout Completion Checklist

The following checks are a minimum condition to pass before extracting the interface in a 3D electromagnetic solver for channel signal integrity simulation.

Automated Checks

- Checking for 100% connectivity.
- Checking for dangling lines/antenna vias/floating vias.
- Via-pin alignment and via overlapping checks.
- Shape island checks.
- Trace-trace separation rules.
- Trace-pad rules.
- Length matching. Run net single-pin and no-pin report.
- Shape no net (in this case, all the A1 or fiducial shapes show up in this report which is okay but make sure there are no floating shapes).
- Waived DRC rules.

Visual Checks

- Are all signals on their correct layers with good ground shielding from source point to destination?

- Ground plane on above and below layer of high-speed signal routing.
- Are there sufficient ground stitching vias surrounding the signal traces?
- Signals should not cross splits or breaks or voids in adjacent planes.
- Review the fab notes to make sure they are same as database. Review spacing, physical and impedance rules.

Chapter - 3: High-Speed SerDes Interface

The Speedster7t AC7t1500 SerDes Interface is compliant with many serial interfaces, with data rates up to 56 Gbps NRZ and 112 Gbps PAM4. While Achronix provides many equalization features in the silicon, optimal signal integrity is critical for the transmitted signal to be interpreted reliably by the receiver. Strategic component placement and careful engineering of the routed channel is necessary to minimize electrical parasitic effects and meet the electrical specifications of the interface. This chapter provides guidance for designing PCBs to achieve those goals.

SerDes Channel Topologies

Chip-to-Chip Topology

The chip-to-chip topology supports communication over a printed circuit board for two chips, with one being the AC7t1500 FPGA and the other being a CPU/FPGA or other endpoint. The interconnect between the chip and module is a relatively medium-reach channel over a PCB consisting of traces and vias.

70540326-01.2020.11.22
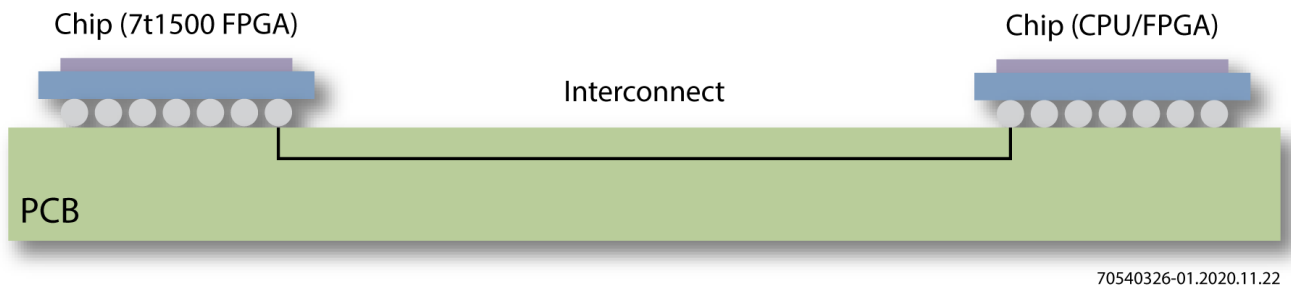
Figure 6: *Chip-to-Chip Topology*

Chip-to-Module Topology

The chip-to-module topology supports pluggable optical or copper cable modules for high-bandwidth switch applications. The interconnect between the chip (AC7t1500 FPGA) and module is a relatively short-reach channel over a PCB consisting of traces, vias and a connector.

70540326-02.2020.11.23
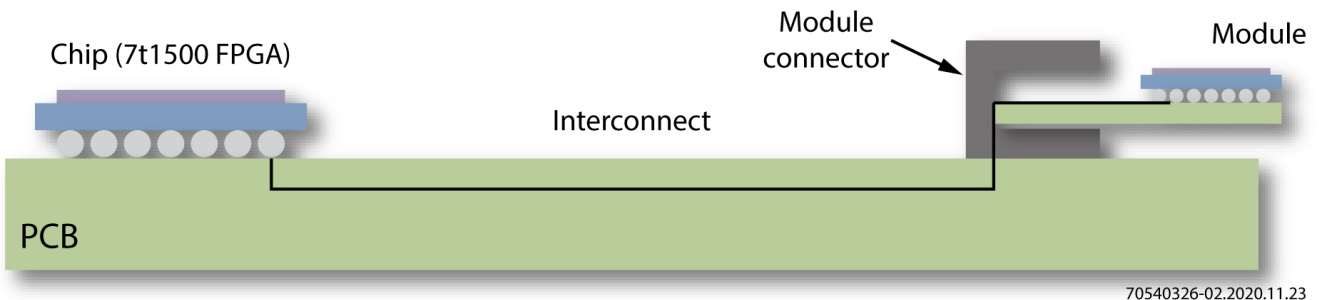
Figure 7: *Chip-to-Module Topology*

Chip-to-Chip over Connector or Cable Topology

This topology supports communication over a printed circuit board and backplane connector or cable for two chips with one being the AC7t1500 FPGA and the other being a CPU/FPGA or other endpoint. The channel between the chip and module is a relatively long-reach channel consisting of interconnect with traces, vias and connectors, or could be a cabled solution or any other advanced PCB technology.
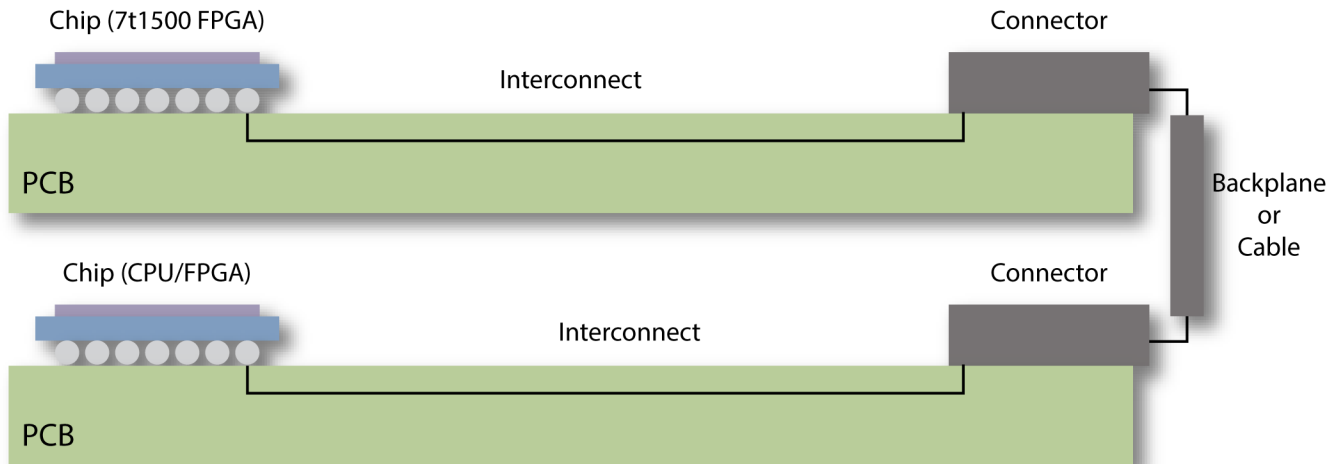
70540326-03.2020.11.23

Figure 8: *Chip-to-chip Over Connector or Cable Topology*

Signal Integrity Specification

Understanding the different high-speed SerDes standards (whether PCIe or Ethernet) is key in building a layout that meets the specification. The PCB designer needs to evaluate for channel loss budgets when deciding on the channel's reach. It is important to evaluate each component of the channel to ensure it is specification compliant. Refer to the electrical specifications of the SerDes standards in meeting the limit lines.

Below are few of the high-speed SerDes standards:

- **PCI Express Gen1, Gen2, Gen3, Gen4 and Gen5** – Refer to the PCI Express standard at https://pcisig. com/.

- **OIF Chip-to-module/chip** – For example, CEI-112G-USR/SR/MR/LR-PAM4 (112 Gbps PAM4 in each channel)

- **Ethernet chip-to-chip/module** – For example, 400 Gbps Ethernet – CDAUI-8 (56 Gbps NRZ on each of 8 channels for a total of 400 Gbps)

- **Optical Chip-Module** – For example, Interlakken (3.125 – 12.5 Gbps, 25 Gbps) consists of 25 Gbps in each of four channels.

Refer to *Speedster7t SerDes User Guide* (UG099) for detailed list of SerDes standards supported by the AC7t1500 FPGA.

Sign-off Simulation

To ensure compliance of the channel with the target specification, the PCB channel must be extracted in a 3D electromagnetic (EM) field solver and simulated using the following sign-off topology:
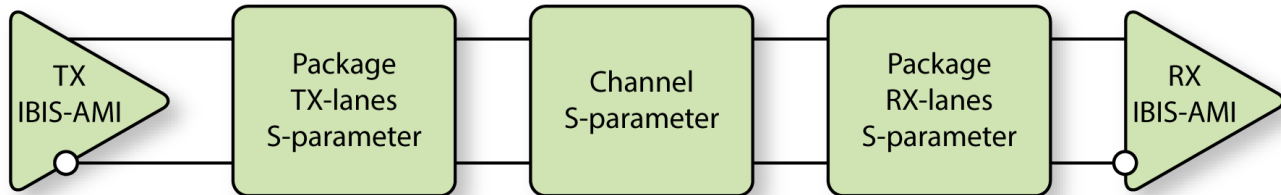
70540326-04.2020.11.23

Figure 9: *Sign-off Simulation Topology*

This topology represents of the entire channel from transmitter to receiver, and includes the following:

- **Transmitter TX IBIS-AMI** behavioral models for the endpoint – In the case of the AC7t1500 FPGA, these models are available from Achronix upon request.

- **Package TX-lanes s-parameter** models for the endpoint – In the case of the AC7t1500 FPGA package, from the silicon bumps to the package pins and including the on-die termination, these models are available from Achronix upon request.

- **Channel S-parameter** model – This model is the extraction of the the target system, and is the concatenation of all models to the system interconnect, including connectors and modules. This model must be provided by the PCB designer.

- **Package RX-lanes s-parameter** models for the endpoint – In the case of the AC7t1500 FPGA package, from the silicon bumps to the package pins and including the on-die termination, these models are available from Achronix upon request.

- **Receiver RX IBIS-AMI** behavioral models for the endpoint – In the case of the AC7t1500 FPGA, these models are available from Achronix upon request.

For final high-speed SerDes PCB sign-off, it is necessary to verify that the channel performance meets the given interface specification for both the transmitter and receiver direction **.**

> ⚠ **Caution!**
>
> **For applications intended to operate at 25 Gbps or above**, the modeling of the package and PCB cannot be regarded as separate entities. It is important to capture the package-to-PCB interface through 3D EM field solver, including the BGA ball and PCB pad for both the signal and ground vias. Contact Achronix Technical Support for assistance when modeling the package-to-PCB interface.

Layout Optimization Guidelines

PCB design for a complex device such as the AC7t1500 FPGA begins with three areas: the stack-up design, component placement, and routing, followed by an iterative process of simulation and adjustment to meet the design goals.

> **Note**
>
> ⓘ Constructing a PCB is an engineering process that considers the system, signaling, power, mechanical and thermal needs of all the components on the PCB. While the needs of the AC7t1500 FPGA are addressed here, it is up to the PCB designer to ensure that all components (including memories, connectors, clocks, VRMs, and CPUs) are addressed in similar fashion.

Component Placement

For general routing guidelines, see the section, "Component Placement (see page 12)".

The figure below presents an overview of the different high-speed SerDes components (PCIe and Ethernet) on a reference board. This presents one possibility of placement of high-speed SerDes connectors (PCIe and Ethernet) on the north side of the FPGA. A key consideration is how far away the connectors can be from the FPGA. While this distance is driven by the loss budget of the chosen interface, it must be validated by a full signal integrity analysis, including extraction and simulation using a 3D electromagnetic simulator.
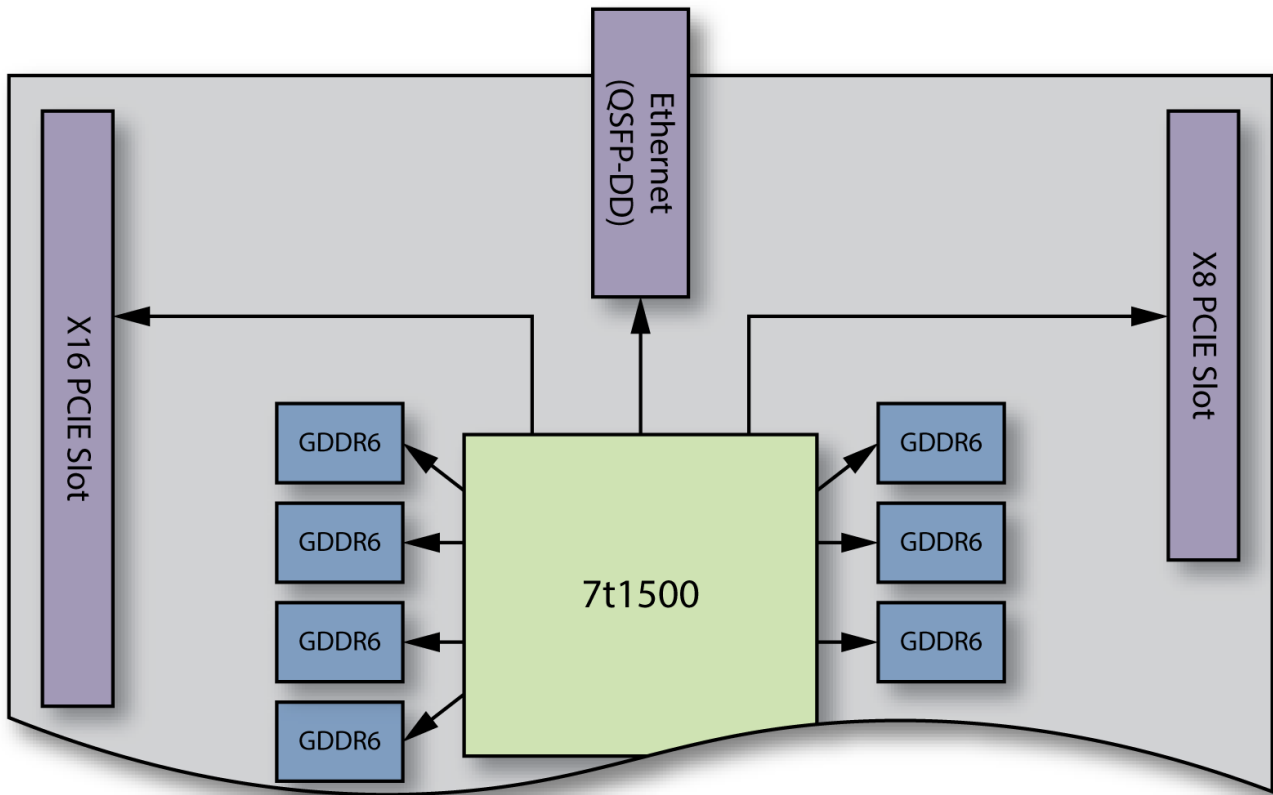
70540326-05.2020.11.23

Figure 10: *Example Placement of High-Speed SerDes Connectors (PCIe and Ethernet)*

The designer of a high-speed SerDes board must carefully engineer the structures where the interconnect transitions from one layer to another, such as within the BGA footprint and at the PCIe or QSFP connector. The goal of such transitional structures is to minimize the impedance discontinuity and thus reduce the loss due to reflections.

The SerDes interface on the AC7t1500 FPGA package is designed with a differential trace impedance of 85Ω, it is important to consider this value when optimizing the different footprints of components and structures along the channel.

Routing Guidelines

For general routing guidelines, see the section, "Routing Guidelines (see page 15)".

Specific to the SerDes interfaces:

- Separate the transmitter and receiver signals to separate routing layers. This separation makes signal breakout and routing easier and reduces near-end crosstalk (NEXT).

  - Transmitter signals are the outermost pins, so they belong on the upper layer.

  - Receiver signals break out of the BGA via array on the lower layer. Separate single-ended and differential signal routing layers in making routing easy. Evaluate for crosstalk of signals routed in parallel for a long distance.

- The AC7t1500 FPGA supports simpler routing of Ethernet connections through a pin-mapping feature that allows the reassignment of the channels in a quad to support specific connector types. Refer to the chapter, "Speedster7t AC7t1500 PCB Pin Mapping on the AC7t1500 SerDes Interface (see page 25)" for more information.

- A common observed practice is **via-In-pad** (VIP) construction by using a laser via in freeing up routing space between signal pins. Signal vias accompanying with ground vias helps to control via inductance. Via constructed using the VIP process are referred to as **VIP vias**.

- A **ground fence** (GF) is a row of ground vias between the high-speed traces. It is intended to reduce crosstalk by containing the electromagnetic fields, much like a Faraday cage. When utilizing a ground fence, it is important to set the distance between the vias such that it does not create a resonance. Keep GF vias at a pitch equal to ½ the wavelength of the signal or smaller.

- **Ground return Vias** are ground vias close to the signal via. The purpose of ground return vias is to provide continuous return current path when signals switch layers to minimize impedance discontinuity. Use ground return vias wherever a signal transitions from one layer to another.

- **Via stub resonance** occurs when a signal transitions between two layers, but the via construction is longer than the distance between the layers. This phenomenon can be controlled through the use of laser vias or through **back-drilling**, the practice of removing the excess via stub by drilling out the barrel of the via.

Component Footprint Optimization Guidelines

FPGA BGA Pad-to-Trace Breakout Footprint

It is important to minimize the impedance discontinuity at the BGA pad-to-trace breakout to minimize loss due to reflection. In the screen captures below, one particular layout optimization is shown that led to reduced impedance discontinuities for BGA pad on layer 01 to layer 03 trace breakout on a reference board:

- On layer 01, a via-In-pad transfers the signal from the BGA footprint directly from layer 01 to layer 03. The VIP allows for much simpler routing by soldering directly over the via. Further, it is important to note the antipads which surround the differential pair signal pads and ground pads.

- On layer 02, ground vias surrounding the differential signal pair provide a controlled current return path accompany the signal as it transits to the layer 02 to layer 03. This use of ground vias minimizes the inductance of the via and provides a predictable return path for the signal. Again, there are antipads surrounding the differential pair signal pads. These antipads should be sized to balance the inductance of the structure in order to match the impedance of the traces. Layer 02 acts as a ground reference plane for the signal routed on layer 03.

- On layer 03, the differential signal breaks out. When routing differential pairs, it is important to avoid routing the pairs close to each other for a long distance as it leads to crosstalk. Further, ground stitching vias are used for routing of the differential pairs to minimize the crosstalk between the pairs.

- On layer 04, a ground plane serves as a reference layer for the signal routed on layer 03.
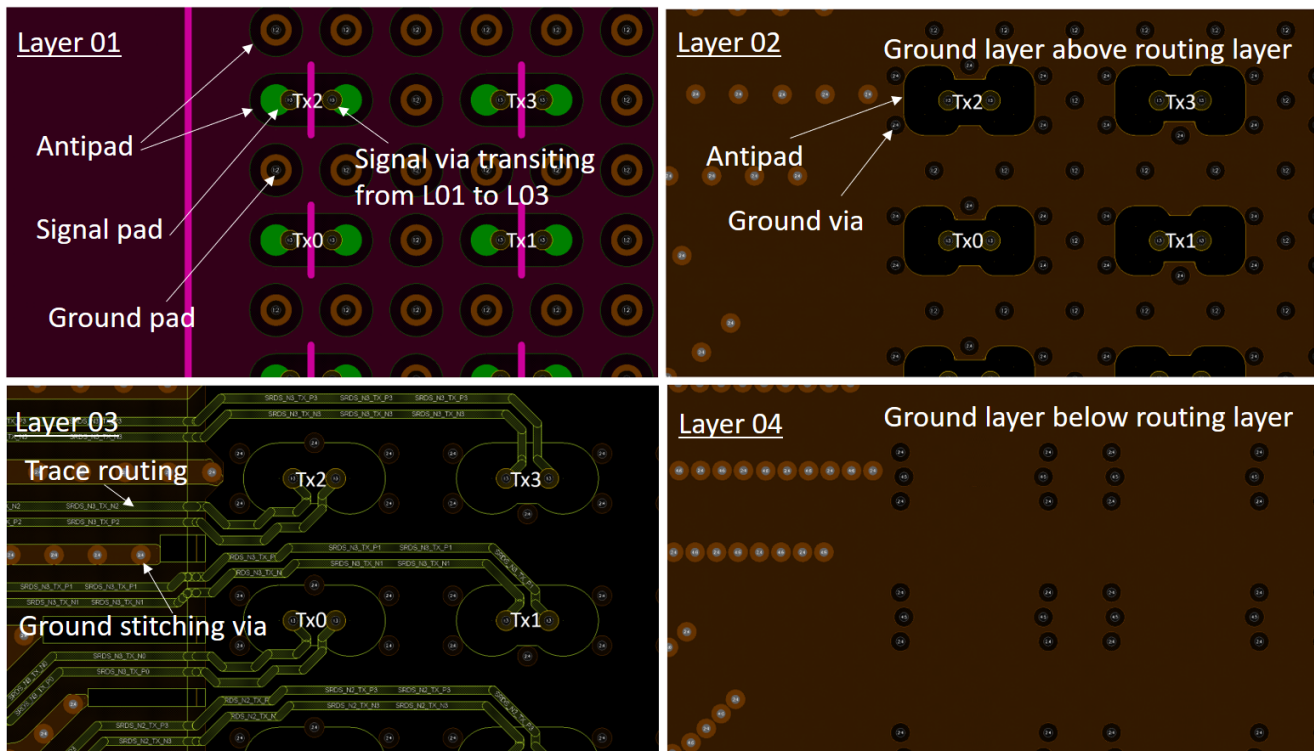
Figure 11: *FPGA BGA Pad Breakout to Layer 3*

DC Blocking Capacitor Footprint

DC blocking capacitors present multiple opportunities for impedance discontinuities where the trace meets the capacitor. Parasitic capacitance between the capacitor pad and the ground plane below is controlled with a cutout of the first and second ground reference plane underneath the capacitor pads. Simulations using a full-wave electromagnetic solver are used to evaluate the number of layers to void under the capacitors. As an example. one particular layout optimization is shown below. This implementation, from an Achronix test board, reduced the impedance discontinuity of the DC blocking capacitor after a transition from the lower layer 03.

- On layer 01, there are capacitor pads with traces routing to the connector. A via carries the signals from the inner layer 03 to layer 01.

- On layer 02, ground vias surround the differential signal pair, providing a ground current return path and accompanying the signal as it transits to the layer 02 to layer 03. The use of ground vias provide a low-inductance return path. Note the antipads which surround the differential pair signal pads. These antipads must be sized to provide the correct capacitance to balance the inductance of the via. Layer 02 acts as a ground reference plane for the signal routed on layer 03.

- On layer 03, the differential signal breaks out to a differential trace of the channel.

- On layer 04, a ground plane provides another reference for the signal routed on layer 03.
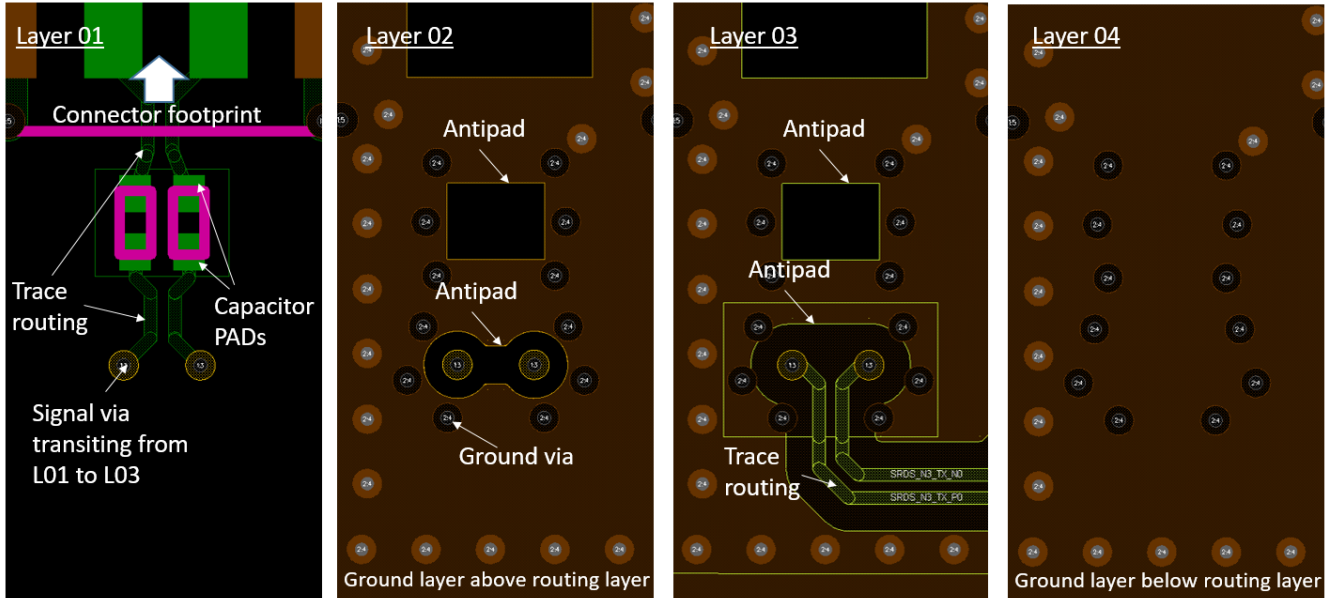
Figure 12: *Example Structure of DC Blocking Capacitor at a PCIe Connector*

Connector Footprint

Refer to the connector manufacturer's specific design recommendations for the best connector performance. If the manufacturer provides no recommendation, perform simulations to determine the best layout footprint optimization. A common practice is to introduce a cutout of the first and second ground reference plane underneath the connector signal pads to compensate for the large parasitic pad capacitance. Simulations using a full-wave solver are to be used in evaluating how many layers underneath the connector footprint are to be cutout. In the snapshots shared below, one particular layout optimization is shown that led to reduced impedance discontinuities for a connector footprint.

- On layer 01, connector pads have traces going to the signal via, which is a laser via to eliminate stubs. The signal transits from layer 01 to layer 03. Ground vias are placed close to the ground pads of this connector to minimize the inductance of the ground pin connection.

- On layer 02, it can be seen that ground vias surround the differential signal pair in providing a ground current return path and accompany the signal as it transits from layer 02 to layer 03. The use of ground vias help in providing a low-inductance return path. Further, antipads surround the differential pair signal pads. These antipads must be sized to control the capacitance of the via. Layer 02 is a ground reference plane for the signal routed on layer 03.

- On layer 03, the differential signal breaks out to the differential stripline of the channel.

- On layer 04, a second ground plane serves as the lower reference for the signal routed on layer 03.
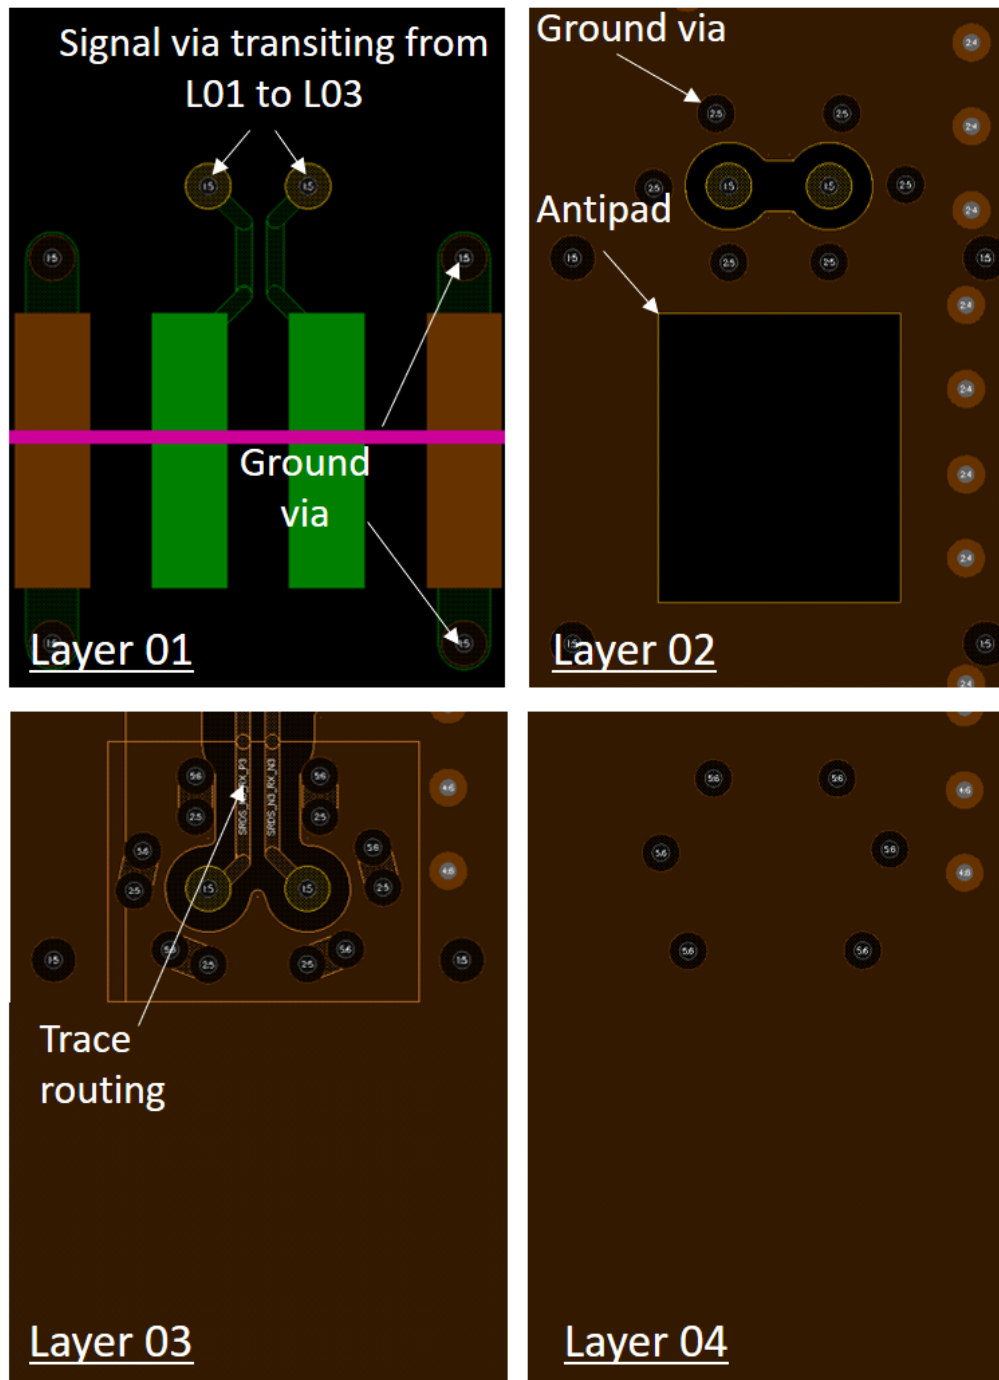
Figure 13: *Example Structure for Transition Via to Connector Pins from Layer 03*

Chapter - 4: SerDes Pin Mapping for Ethernet Connectivity

The PCS interface within the AC7t1500 allows limited reassignment (or re-mapping) of the signals to the package pins. The goal of the re-mapping is to enable routing of signals to QSFP, QSFP28, and QSFP-DD 1N connectors using only two layers, one for transmit and one for receive. Routing opposite-direction signals on different layers limits near-end crosstalk (NEXT) in the trace portion of the interconnect.

> **Note**
>
> Quads are configured in pairs. Mapping applies to both quads in the pair.

Table 2: *Lane Remapping*

	0	1	2	3	4	5	6	7
Linear	0	1	2	3	4	5	6	7
QSFP-28	0	2	1	3	4	6	5	7
QSFP-DD	0	2	4	6	5	7	1	3

There is a timing penalty with using the non-linear maps, which slightly increases latency:

- QSFP-28 has a single flop on the data path
- QSFP-DD has two flops on the data path

It is recommended to only use the mapping that is appropriate to the application. Selection of the SerDes pin mapping is performed through ACE, refer to the section " SerDes Lane Mapping " in the *Speedster7t Ethernet User Guide* (UG097) for details on lane mapping.

SerDes Connectivity – Linear Mapping

Linear mapping is generally most useful for PCIE interfaces on Quads 0, 1, 2 & 3 and 6 & 7. If a flyover cable connection to a connector is used, this mapping is also the best choice, as any re-numbering of the traces can be performed in the flyover cable.
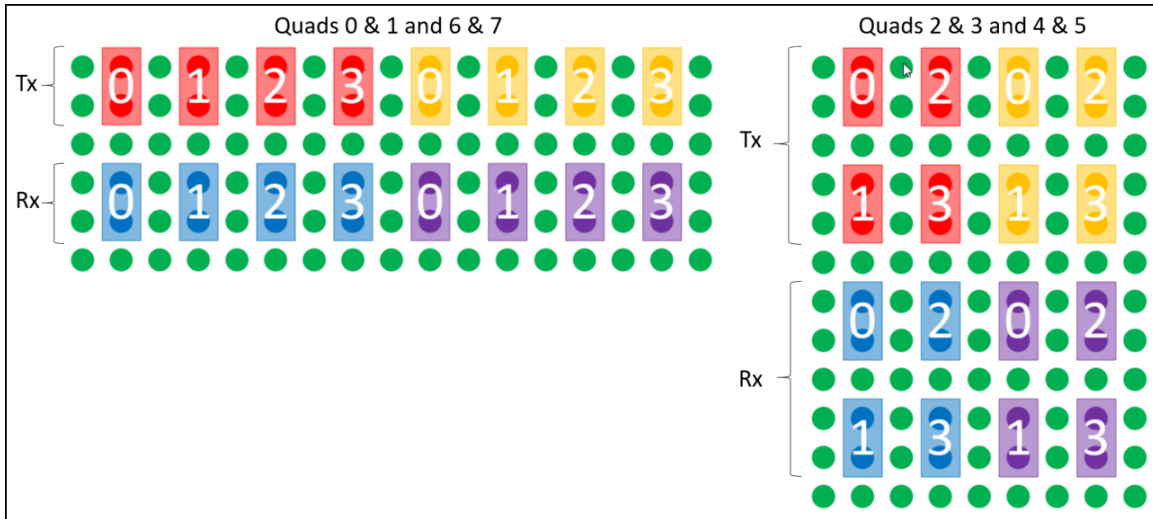
Figure 14: *Linear Pin Mapping*

SerDes Connectivity – QSFP-28 Mapping

QSFP-28 mapping swaps lanes 2 and 3 in each of the two quads to enable two-layer routing to a QSFP or QSFP-28 style connector,
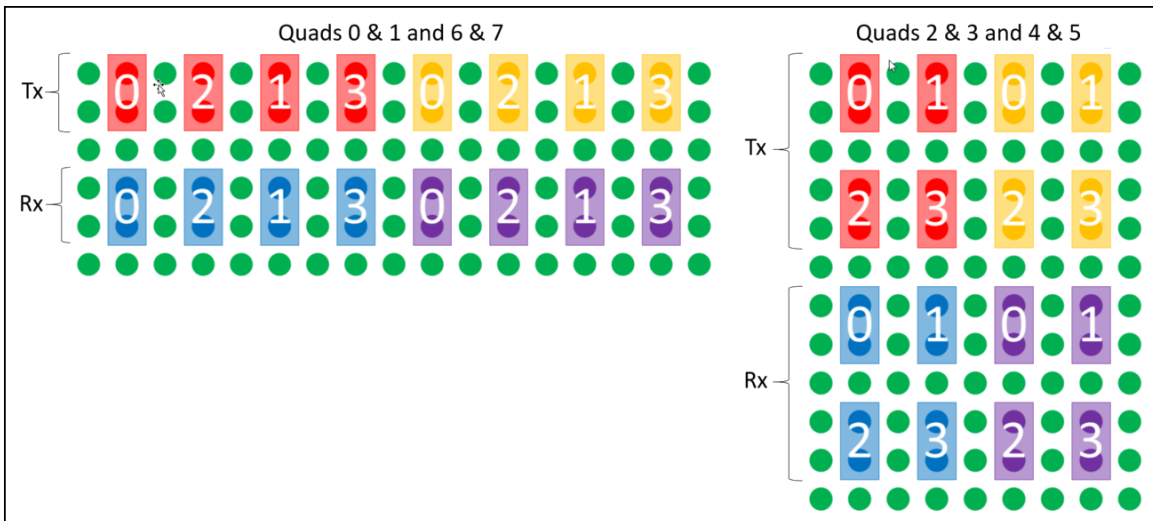
Figure 15: *QSFP-28 Mapping*

SerDes Connectivity – QSFP-DD Mapping

QSFP-DD mapping treats two quads as one octal group. It swaps the lanes 2 and 3, then splits the lower order quad and inserts the higher order quad inside it. This strategy makes it easier to route the eight channels required for QSFP-DD on two layers.
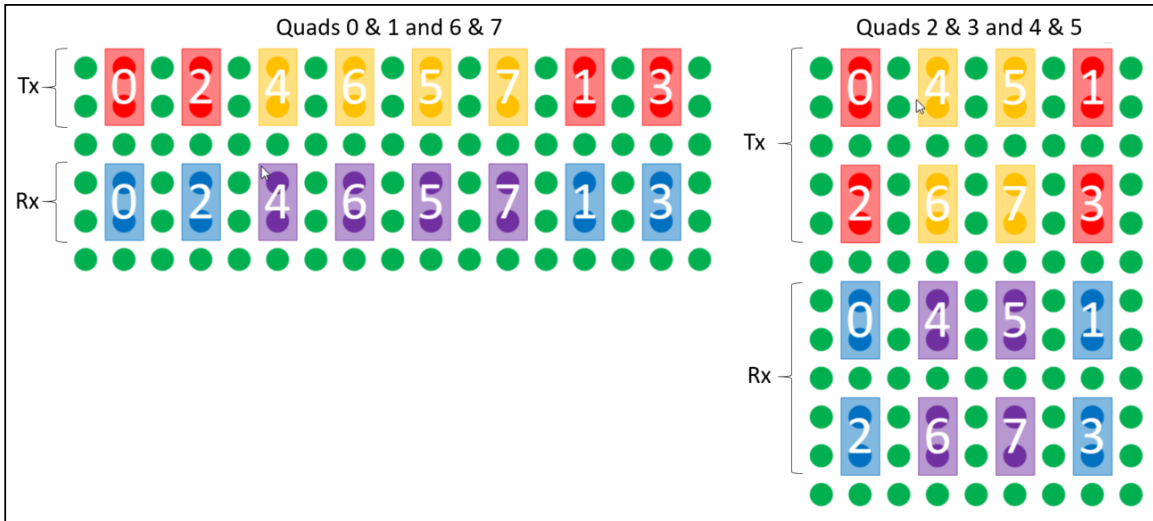
Figure 16: *QSFP-DD Mapping*

Mapping Examples – QSFP-28

The QSFP-28 mapping mode works better on Quads 2, 3, 4 and 5, where the pins for each quad are arranged in a 2 × 2 formation. As can be seen in the two figures below, the 2 × 2 quads can be routed in two layers, while the 1 × 4 pin quads either require four layers or extra vias to keep them in two layers:
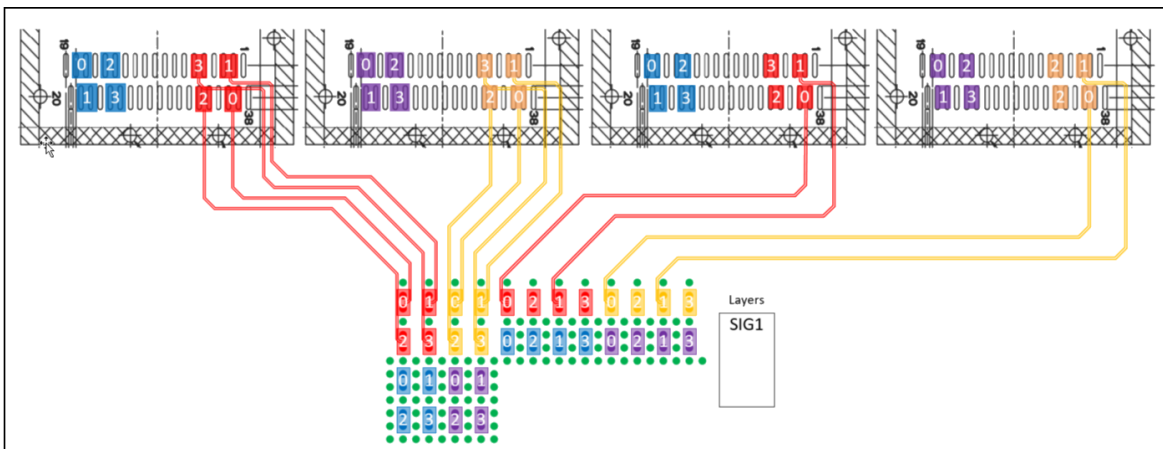
Figure 17: *Routing to QSFP-28 Connectors – Layer 1*
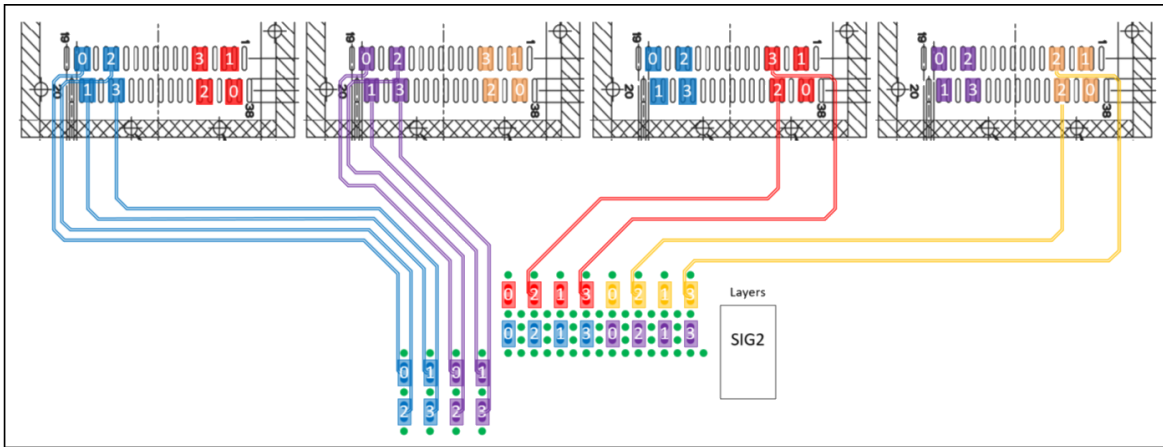
Figure 18: *Routing to QSFP-28 Connectors – Layer 2*
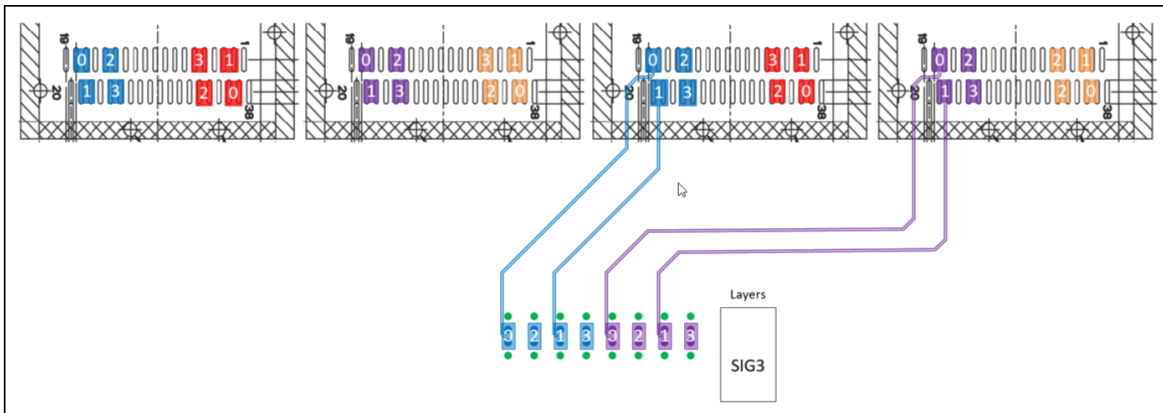
Figure 19: *Routing to QSFP-28 Connectors – Layer 3*
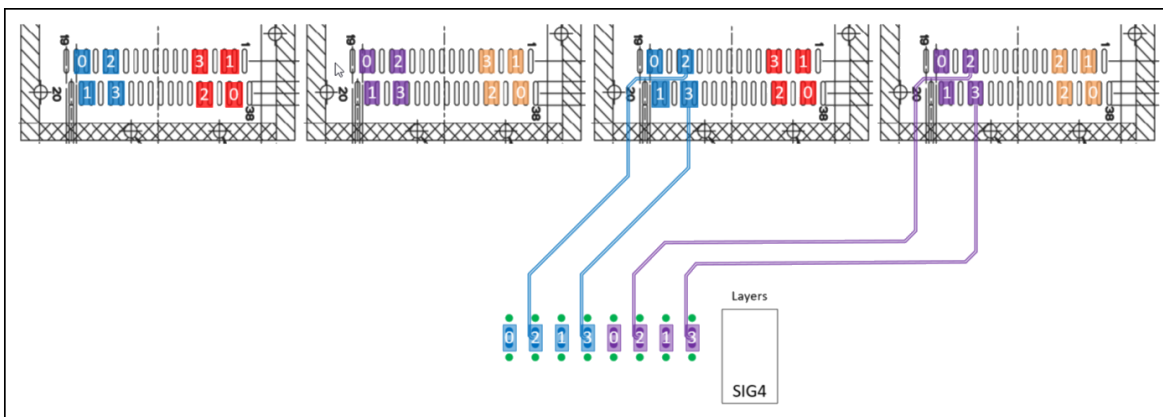
Figure 20: *Routing to QSFP-28 Connectors – Layer 4*

Mapping Example – QSFP-DD

The QSFP-DD mapping works very well for QSFP-DD 1N1 connectors (single DD module). However, for a QSFP-DD 2N1 connector (one allowing two modules to be inserted vertically), four layers will still be required.
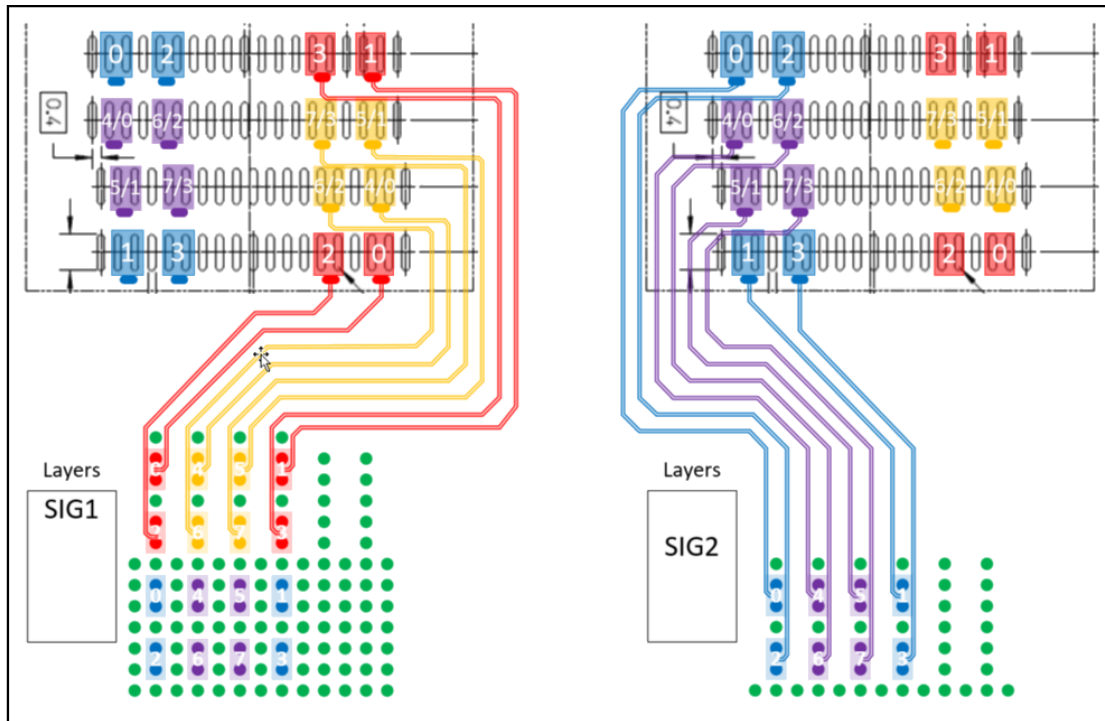
Figure 21: *Routing to QSFP-DD Connectors – Layer 1 & 2*

Chapter - 5: Speedster7t AC7t1500 PCB GDDR6 Interface

The GDDR6 interface supports a maximum data rate of 16 Gbps and is targeted at systems that require low-latency and high-bandwidth memory solutions. The high frequency of operation requires the package and PCB design to be optimized for minimal losses and minimal crosstalk. GDDR6 is a high-speed SDRAM communication protocol designed to support applications requiring high bandwidth such as high-performance computing and machine learning operations.

> **Note**
>
> It is up to the PCB designer to extract the design in an appropriate electromagnetic modeling tool and simulate to verify operation against the GDDR6 specification.

GDDR6 Channel Topologies

GDDR6 SDRAM based memory systems are typically divided into channels. GDDR6 is designed around a 16-bit wide channel. A channel can be comprised of a single device operated in ×16 configuration, or two devices each operated in ×8 configuration. Both these configurations are supported by an AC7t1500 device:

- A ×16 configuration (two independent ×16 bit data channels)
- A ×8 configuration (two devices, each with ×8 channels, in a back-to-back clamshell configuration)

By-16 Configuration

Data connection is point to point for the ×16 mode. Both channels acts as separate devices and communicate independently to the memory controller. The ×16 mode's point-to-point topology supports communication over a PCB for two chips with one being the AC7t1500 FPGA and the other being a DRAM on the other endpoint.

70540420-01.2020.11.23
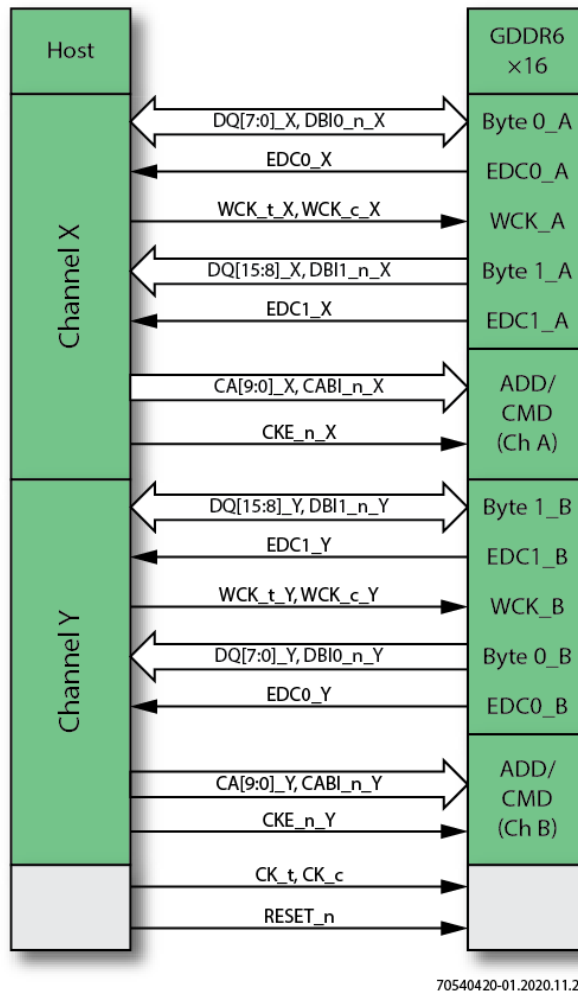
Figure 22: *System View of a ×16 Configuration*

70540420-02.2020.11.23
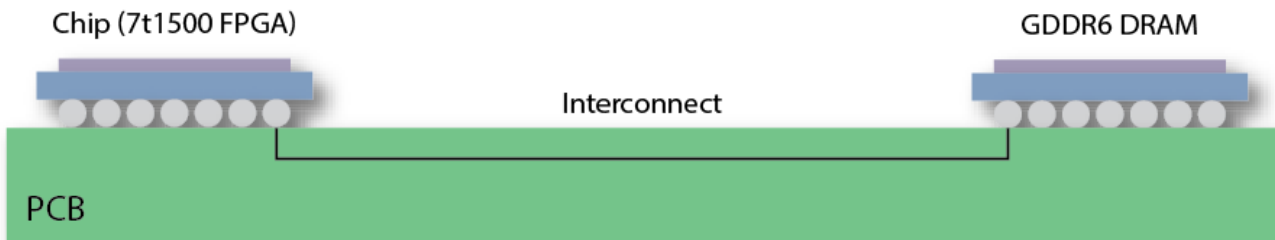
Figure 23: *Data Connection in ×16 Mode*

By-8 Configuration

For ×8 mode the devices are typically assembled on opposite sides of the PCB (one device on the top layer and the other on the bottom layer) in what is referred to as a clamshell layout.
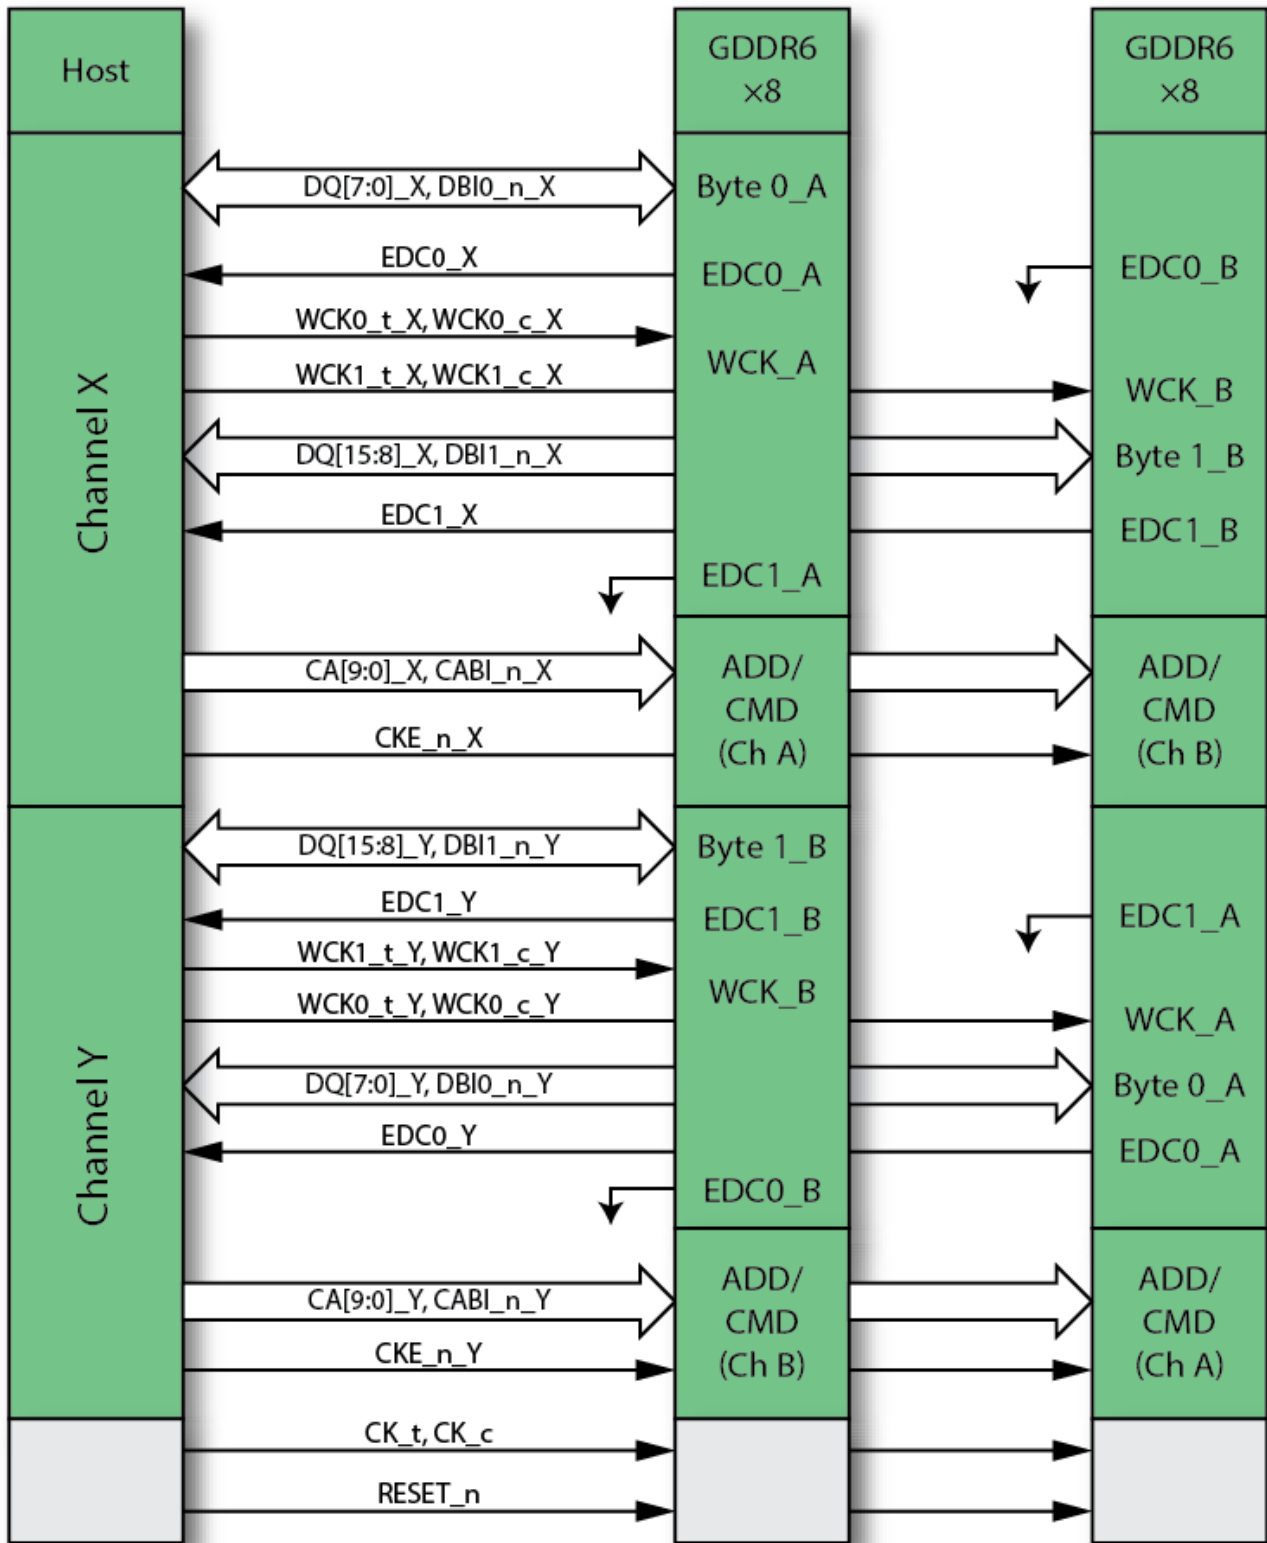
Figure 24: *System View of a ×8 Clamshell (Dual-Memory) Configuration*

The figure below clarifies the use of ×8 mode and how the bytes are enabled/disabled to give the controller the view of the same bytes that a controller sees with a single ×16 device. For a 16-bit channel using two devices in a clamshell design, byte 0 comes from channel A from the top device and byte 1 comes from channel B from the bottom device and look equivalent to the ×16 mode at the controller.
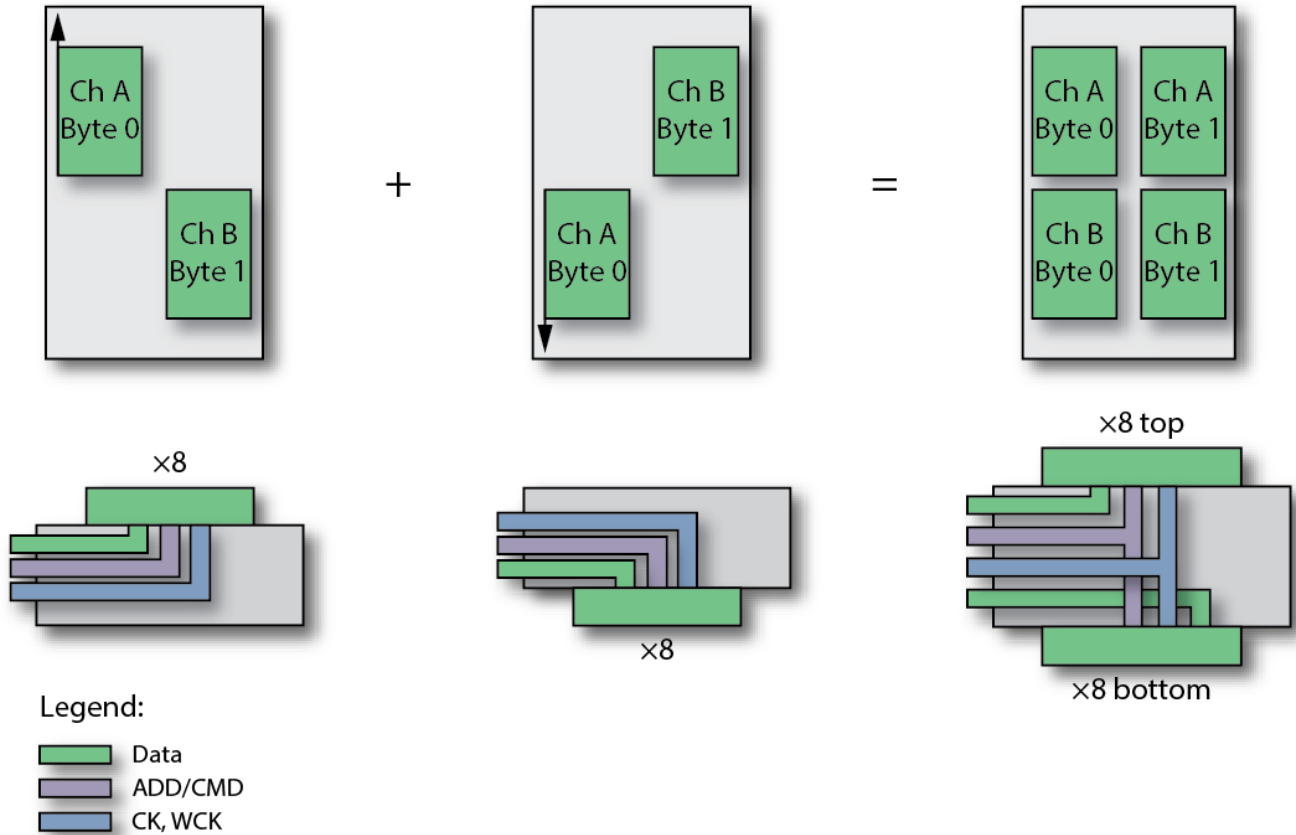
70540420-04.2020.11.25

Figure 25: *Layer Assignments for Clamshell Mode*

GDDR6 supports the ×8 configuration due to the dual-channel architecture of GDDR6. In ×8 mode for data connections, only one of the two data bytes per channel is enabled (byte 0 of channel 0 and byte 1 of channel 1), while the other two data bytes are disabled during data transfer.
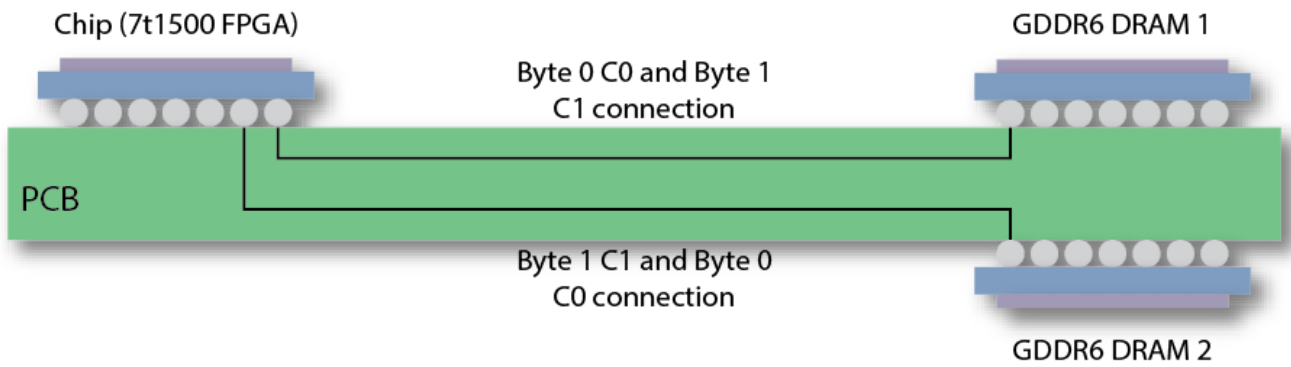
Figure 26: *Data Connection in ×8 Mode*

in ×8 configuration follow a "T" topology as shown in the figure below. command/address (CA) bytes for both channels are routed together to both the DRAMs present on the bottom and top layers.
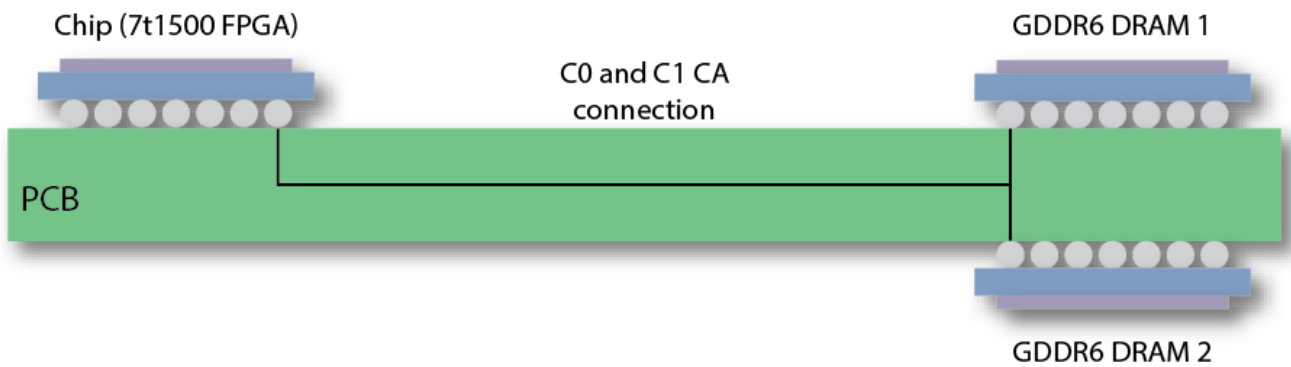
Figure 27: *CA Connection in ×8 Mode*

GDDR6 Mapping from ACE

There is a difference between how the eight GDDR6 interfaces are referenced in ACE and how they are referenced in the physical domain (schematic/layout). Internally, the interfaces are assigned sequential numbers. At the PCB level, the interfaces are assigned labels according to their "map view" — the view of the device as seen from above the PCB — on the east and west sides (pin A1 is considered the northwest corner). In that case, the interfaces are numbered east and west: E0 to E3 and W0 to W3.

Table 3: *GDDR6 Mapping to External References*

Internal Label	External Label	Internal Label	External Label
0	W3	7	E3
1	W2	6	E2
2	W1	5	E1
3	W0	4	E0

As an example, the **W0** data bus, e.g. GDDR6_**W0**_C0_DQ[7..0], is controlled by the registers referenced by GDDR6 port **3**.

Signal Integrity Specification

Signal Integrity PCB Design Specification Guidelines

The following PCB-level specifications are recommended for the GDDR6 signals:

Table 4: *GDDR6 Channel Specifications*

Parameter	Specification
Insertion loss	DQ > −2dB @ 8 GHz
	WCK pair > −2 dB @ 8 GHz
	CA > −2dB @ 2 GHz
	CLK_P/N > −2dB @ 2 GHz
Return loss	DQ < −20 dB @ 8 GHz
	WCK pair < −20 dB @ 8 GHz
	CA < −20 dB @ 8 GHz
	CLK_P/N < −20 dB @ 8 GHz
DQ-DQ power sum crosstalk (cumulative sum of crosstalk from all other DQ signals on a victim DQ pin)	< -23 dB @ 8 GHz (mandatory) < -25 dB @ 8 GHz (stretch goal)

Parameter	Specification
DQ-CA/CA-DQ power sum crosstalk (cumulative sum of crosstalk from all DQ signals on a victim CA pin/cumulative sum of crosstalk from all CA signals on a victim DQ pin)	< -23 dB @ 8 GHz (mandatory)
CA-CA power sum crosstalk (cumulative sum of crosstalk from all other CA signals on a victim CA pin)	< -24 dB @ 2 GHz (mandatory) < -25 dB @ 2 GHz (desirable) < -21 dB @ 4 GHz (stretch goal)
EDC-DQ individual crosstalk (individual crosstalk between any EDC pin and any DQ pin)	< -33 dB @ 4 GHz and < -30 dB at 8 GHz
EDC-CA individual crosstalk (individual crosstalk between any EDC pin and any CA pin)	< -33 dB @ 4 GHz and < -30 dB at 8 GHz
DQ/CA-WCK power sum crosstalk (cumulative crosstalk on any WCK pair because of DQ and CA signals)	Single ended WCKP/N to DQ/CA: < -25dB @ 8 GHz (stretch goal < -27 dB @ 8 GHz). Diff WCKP/N to DQ/CA: < -30 dB @ 8 GHz
DQ/CA-CLK power sum crosstalk (cumulative crosstalk on any CLK pair because of DQ and CA signals)	Single ended CLKP/N to DQ/CA: < -25 dB @ 8 GHz (stretch goal < -27 dB @ 8 GHz). Diff. CLKP/N to DQ/CA: < -30 dB @ 8 GHz
Recommended Impedance	Single-Ended: 50Ω DQ[7:0], EDC, DBI, CA
	Differential: 100Ω WCK_P/WCK_N, CLK_P/CLK_N

Table Note

- Impedance recommendations are based on package impedances of 50Ω and 100Ω for single-ended and differential signals, respectively. Designers are advised to run SI simulations to determine their optimal impedance settings.

Delay Matching

The GDDR6 controller is able to adjust the timing of individual lines to enable precise timing alignment. While some length adjustment is required, it is generally to within some multiple of the DQ signal unit interval (UI). For instance, when the trace lengths of the WCK and SD_CLK lines have been matched within the target specified, the DLL inside the GDDR6 controller can further tune the signals to meet the setup and hold requirements of the interface.

The table below specifies the skew adjustment requirements for the GDDR6 signals at the DRAM component, i. e., these skew requirements have to be met for the AC7t1500 package and the PCB interconnect length.

Delay tuning requires the addition of serpentine traces, which consume routing space and can affect impedance and increase crosstalk. It is recommended that the designer balance the the amount of delay tuning against the requirements of the interface. In other words, do not match to exactly the requirements stated below, but exceed them to a degree allowed by the routing resources available in order to preserve operating margin.

Achronix can provide delay numbers for the GDDR6 traces in the package if requested. These numbers can be embedded in the PCB design to allow silicon-to-memory device pin delay tuning.

Table 5: *GDDR6 Length Matching Targets*

Signal to Target	Match Requirement	Suggested Target1
DQ [[7:0]] to WCK0 DQ [[15:8]] to WCK1	±1 UI of DQ = ±62.5 ps	±0.8 UI of DQ = ±50 ps
CA[[9:0]] to SD_CLK	±0.5 UI of DQ = ±31.25 ps	±0.5 UI of DQ = ±31.25 ps
WCK[[1:0]] to SD_CLK	±2 UI of DQ = ±125 ps	±1UI of DQ = ±62.5 ps

Table Note

- Achronix has been successful in achieving these matching targets within its own designs. They serve as a useful but flexible design target for designers.

Intra-pair Skew

To avoid timing errors and common-mode conversion (a source of EMI), it is recommended that both components of each differential signal be as close to equal in length as reasonably possible.

- **WCK_C and WCK_T** – To be length matched within 50 μm (2 mils).
- **CLK_P and CLK_N** – To be length matched within 50 μm (2 mils).

GDDR6 Signal Integrity Sign-off Simulations

When simulating the read- and write-cycles of the GDDR6 transaction, it is necessary to collect eye diagrams both before and after the DFE. While the eye after the DFE is the measurement of concern, the eye before the DFE can be useful in debugging the channel.

Signal Integrity Sign-off

Signal integrity designers have to decide the worst-case simulations to qualify their system. At the minimum, Achronix recommends to run the following sign-off simulations:

- Write-cycle channel simulations for data/WCK signals
- Read-cycle channel simulation to confirm BER compliance
- Command/address transient/crosstalk simulation for the topology chosen (×8/×16)

Following sections describe all three simulations in detail.

GDDR6 DQ Write Channel Simulations

As GDDR6 has BER requirements, it is mandatory to run channel simulations and ensure that the eye height /width values recommended by JEDEC/DRAM vendor at the DRAM component are met.

GDDR6 DQ Write Channel Simulation Setup

It is recommended to identify all the aggressors to DQ signals and simulate them together. Typically, DQ signals have coupling with other DQ signals and CA signals. The complete cluster of signals that have crosstalk impact on each other have to be simulated so that the impact of crosstalk is correctly captured. Running a complete instance of GDDR6 in a single run is the recommended sign-off simulation as this captures all the crosstalk impact. However, if the run times are huge, depending on the design, a subset of the instance can be run such as two DQ bytes and the CA byte if it is believed that the crosstalk is well captured with this approach.

For both ×8 and clamshell mode, the data byte connection is always point to point and hence, the simulation setup is similar.
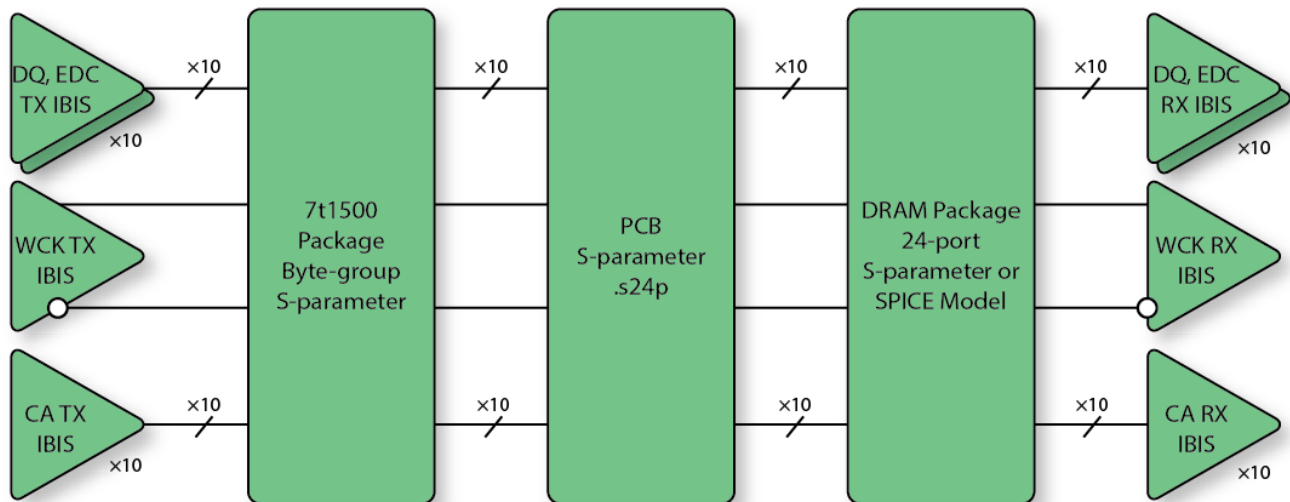
Figure 28: *GDDR6 DQ Write Simulation Setup*

This topology represents of the entire channel from transmitter to receiver, and includes the following:

- **DQ and WCK IBIS Model** – These IBIS models mimic the I/O behavior of the AC7t1500's GDDR6 DQ/DBI /EDC and WCK I/O. These models are available from Achronix upon request.

- **AC7t1500 Package S-parameter** – This parameter is the AC7t1500 package S-parameter model. In the figure above, a 24-port S-parameter model is shown which represents the DQ0-7/DQ8-15, corresponding DBI and EDC signals, along with a corresponding WCK pair for the AC7t1500's package. This package S-parameter models the signal interconnects from the silicon bumps to the package BGA pins, and is available from Achronix upon request.

- **PCB S-parameter** – The PCB model should capture the interconnect between the FPGA package and the DRAM package. The modelling of the PCB must be performed using a 3D field solver for better accuracy.

- **DRAM Package** – This model is provided by the GDDR6 DRAM device manufacturer.

- **DQ and WCK RX IBIS** – These IBIS models for the DRAM receiver I/O behavioral are provided by the GDDR6 DRAM manufacturer.

GDDR6 DQ Write Channel Simulation Sign-off Eye Requirement

> **Note**
>
> Refer to JEDEC specifications and your DRAM vendor for the eye mask requirement at the DRAM component.

Eye mask at the DRAM has to be adjusted to account for the transmitter jitter.

The jitter contribution of the AC7t1500 device (excluding AC7t1500 parasitics):

Tx Total Jitter (TJ) limit for write cycle = 12.5 ps @ BER1E-10

There are a few important simulation considerations for designers:

- GDDR6 receive I/O typically have DFE support. The DFE functionality might need to be modeled separately (IBIS models do not support DFE modelling) if the DRAM vendor requires the eye probe to be used on silicon after the receiver I/O.

- Channel simulations do not take power distribution noise into account which might lead to optimistic simulation results. One way to handle this optimism is to run system-level signal power transient simulations for sufficient cycles with worst-case power noise injected so that a mask adjustment factor may be determined. This mask adjustment factor should increase the eye mask to compensate for the power noise. Meeting this mask gives more confidence to the designer in the signal integrity of the system.

- Channel simulations might not always capture the worst-case crosstalk, and hence, it might be necessary to either run separate transient simulations to capture the crosstalk between signals accurately and perform some mask adjustment, or, with a user-defined input pattern, excite worst-case crosstalk between signals.

GDDR6 DQ Read Channel Simulations

Similar to GDDR6 DQ write simulations, read simulations have a BER requirement and hence, channel simulation is mandatory.

GDDR6 DQ Read Channel Simulation Setup

DQ signals have coupling with other DQ signals and CA signals. So, it is recommended to identify all the aggressors to DQ signals and simulate them together. Typically, the complete cluster of signals that have crosstalk impact on each other have to be simulated so that the impact of crosstalk is correctly captured. Running a complete instance of GDDR6 in a single run is the recommended sign-off simulation, as this captures all the crosstalk impact. However, if the run times are huge, depending on the design, a subset of the instance can be run such as two DQ bytes and the CA byte if it is believed that the crosstalk is well captured with this approach. For both ×8 and clamshell mode, the data byte connection is always point to point and hence, the simulation setup is similar. The topology shown represents of the entire channel from transmitter to receiver and uses the same models as for the write simulation above.

The modelling of the entire signal interconnect group must be performed using a 3D field solver as the frequency of operation is high.
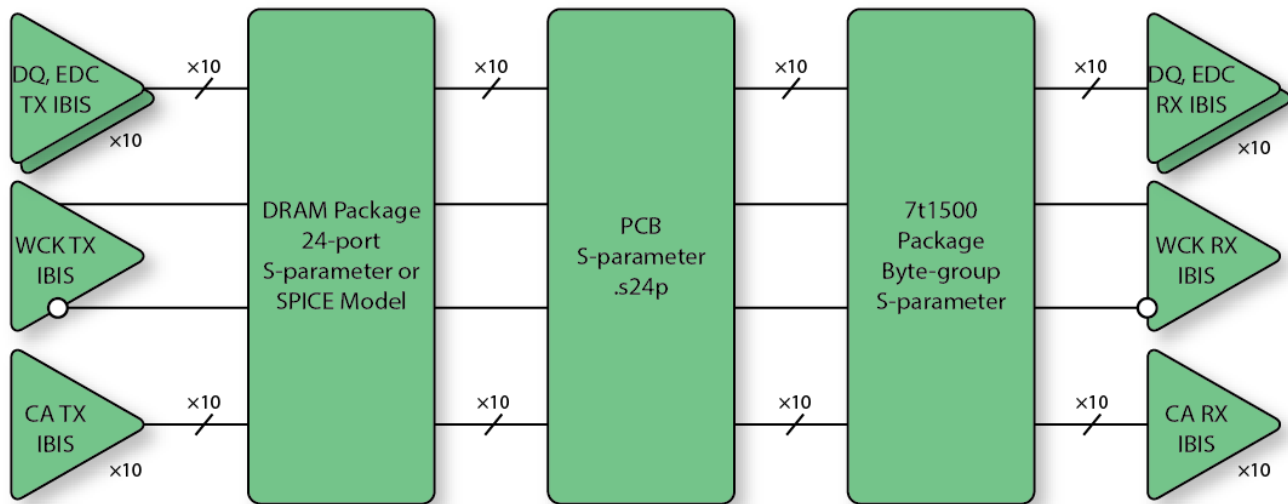
70540420-08.2021.10.22

Figure 29: *GDDR6 Read Simulation Setup*

GDDR6 DQ Read Channel Simulation sign off eye requirement:

For the read cycle, the following mask has to be met:

- Probe point is at the input of the AC7t1500 receiver
- Eye height requirement = 270 mV (AC) and 180 mV (DC)
- Eye width requirement = 25 ps (0.4 UI of DQ @ 16 gbps)

Contact Achronix Support for more details on creation of the eye mask for the read/write cases in specific scenarios.
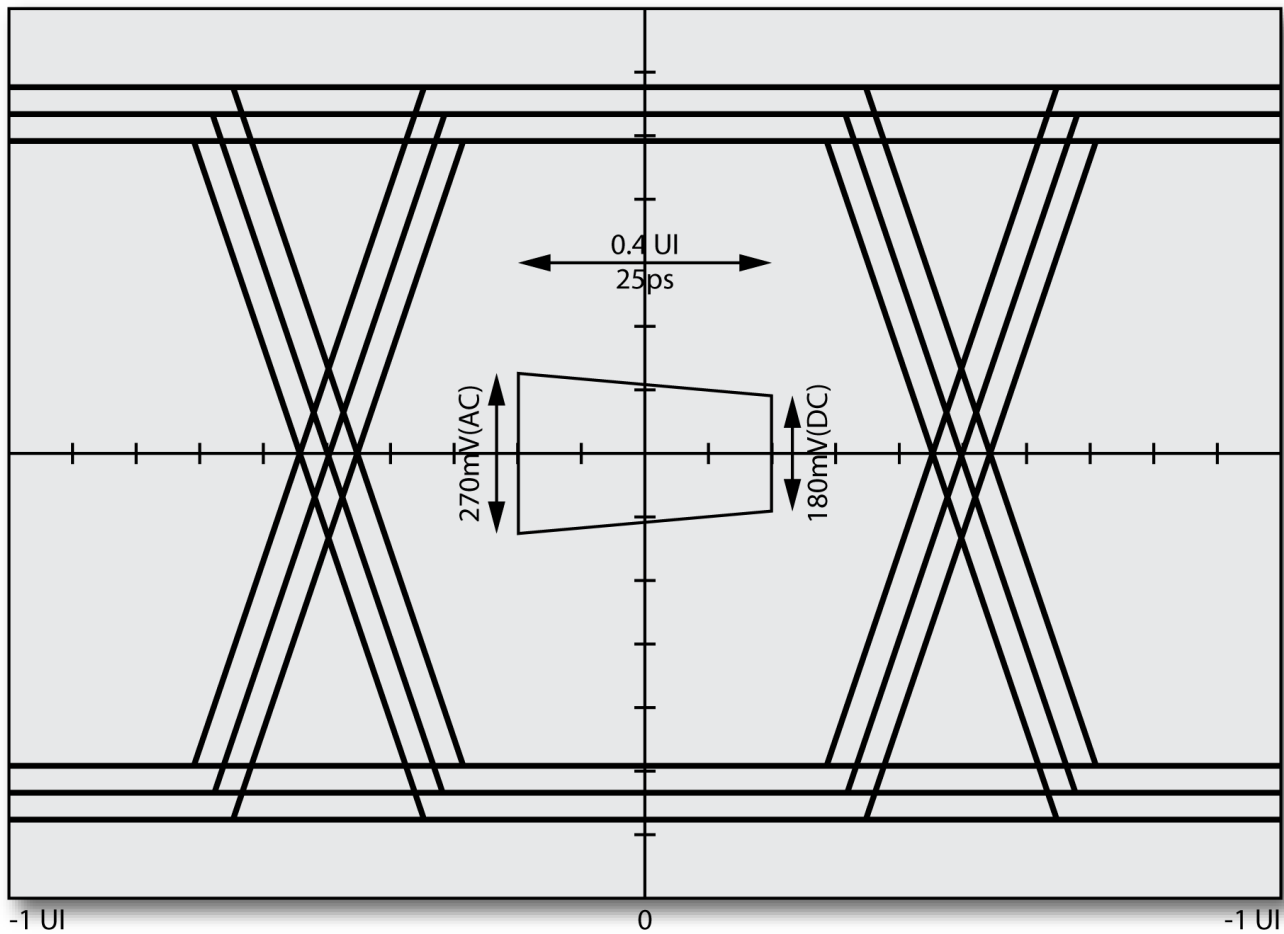
Figure 30: *Read-Eye Mask*

GDDR6 Command-Address Transient Simulations

Transient simulations are recommended to check signal integrity of the command address signals on the DRAM component.

GDDR6 Command-Address Transient Simulation Setup

It is recommended to simulate all command address signals of one channel (CA0-CA9, CKE_N, CABI_N) along with CLK_P/N for the GDDR6 instance in a single simulation testbench (see the figure below). While for the ×16 mode, command address signal connections are point to point, for clamshell mode, command address signals follow a "T" topology and connect to two receivers on separate DRAMs.
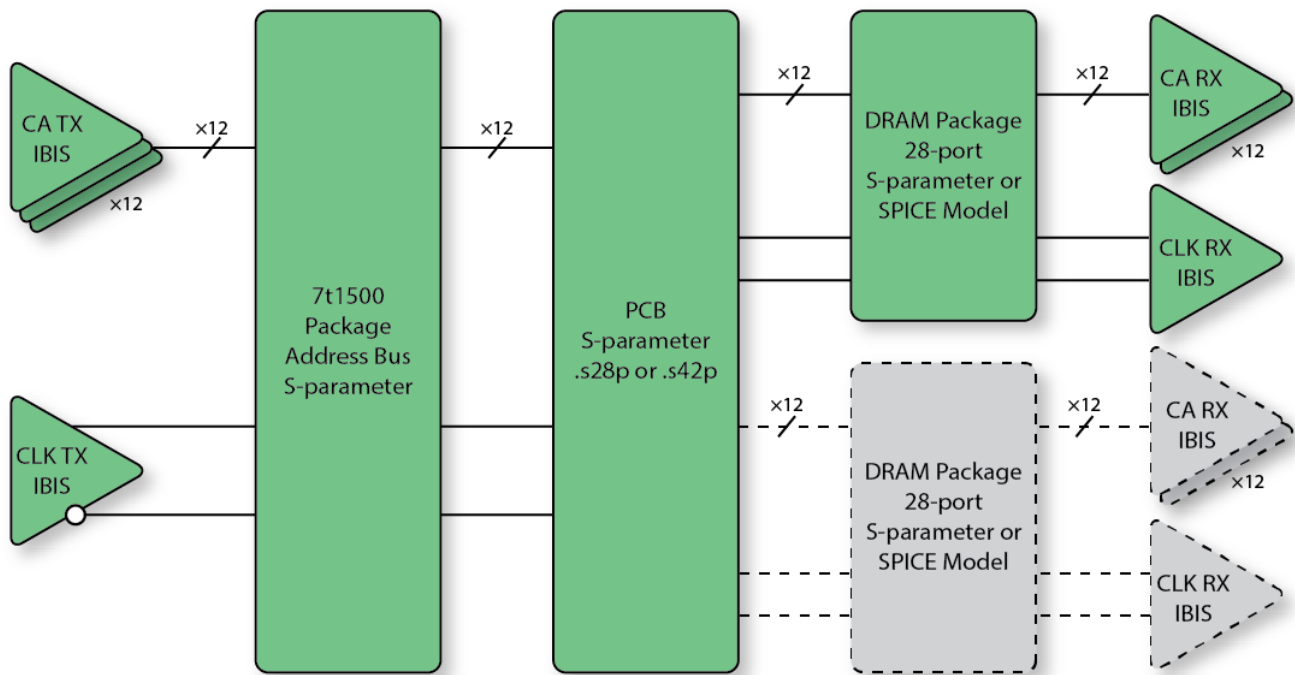
70540420-10.2020.11.25

Figure 31: *GDDR6 Command-Address Bus Simulation Setup*

This topology represents the entire channel from transmitter to receiver and includes the following:

- **CA and CLK IBIS Model** – These IBIS models mimic the I/O behavior of AC7t1500's GDDR6 CA and CLK pair signals. These models are available from Achronix upon request.

- **AC7t1500 Package S-parameter** – This parameter is the AC7t1500 package S-parameter model. In the figure above, a 28-port s-parameter model is shown which represents the s-parameter model for CA0-CA9, CKE_N, CABI_N and CLK_P/N for the AC7t1500's package. This package S-parameter models the signal interconnects from the silicon bumps to the package BGA pins, and are available from Achronix upon request.

- **PCB S-parameter** – The PCB model should capture the interconnect between the FPGA package and the DRAM package. For a single-chip (×16) configuration, each CA TX I/O is connected to a single receiver whereas for ×8 mode (clamshell mode), each CA transmit I/O is connected to two DRAMs. In the topology, a 24-port S-parameter model represents a ×16 mode and a 42-port S-parameter model represents a ×8 mode model. The modelling of the entire signal interconnect group must be performed using a 3D field solver for better accuracy.

- **DRAM Package** – This model is provided by the GDDR6 DRAM device manufacturer.

- **CA and SDCLK RX IBIS** – The DRAM receiver I/O behavioral models should be provided by the GDDR6 DRAM manufacturer.

GDDR6 Command Address Transient Simulation Sign-off Eye Requirement

For command address transient simulation, the following mask has to be met:

- Probe point is at the input of the AC7t1500 receiver

- Eye height requirement = 360 mV (AC) and 270 mV (DC)

- Eye width requirement = 125 ps (0.5 UI of CA @ 4 Gbps)

> **Note**
>
> (i) For 16 Gbps GDDR6 operation, GDDR6 CA signals run at 4 Gbps

Contact Achronix Support for more details on the creation of the eye mask for the read/write cases in specific scenarios.
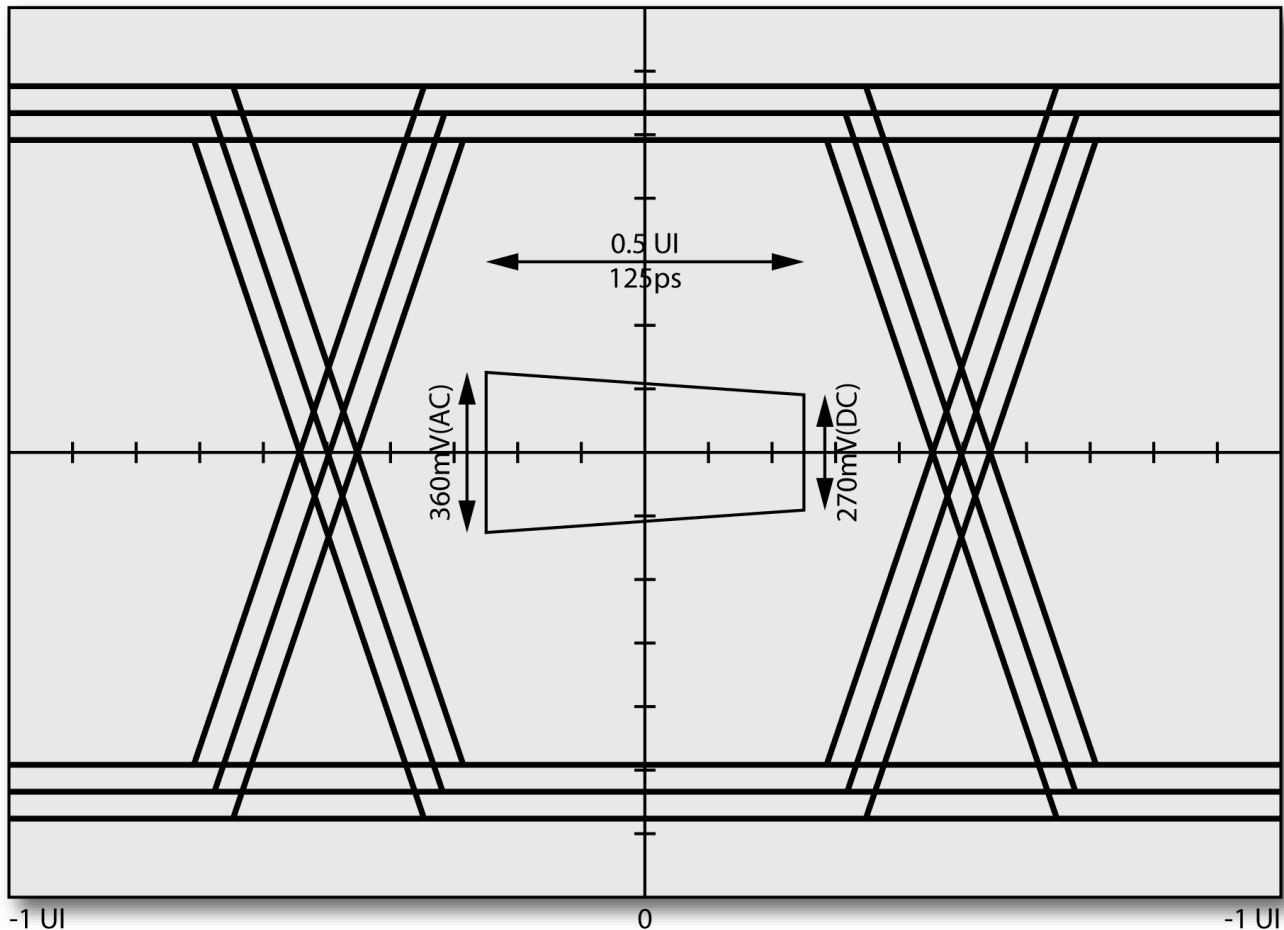
70540420-11.2020.11.30

Figure 32: *CA Eye Mask*

Layout Optimization Guidelines

Stack-up Guidelines

Achronix recommends routing DRAM channels using stripline (inner) traces. Since each routing layer has propagation delay and impedance variations, signals within a given functional group should route using the same layer and geometry. See the section, "Board Construction - the Stack-up (see page 8)", for general direction on constructing the PCB stack-up

Design the stack-up to keep via stubs to a minimum (less than 1/4 of a wavelength at the Nyquist frequency) in order to minimize return loss and impedance discontinuities. Route signals strategically to minimize stubs and, when they are unavoidable, backdrill vias in the FPGA area.

Placement Considerations - Crosstalk Optimization

The GDDR6 interface supports two channels, each with 16 bits for a total data width of 32 bits. Due to a large number of data bits, routing of these signals requires special attention or else high crosstalk may be observed. There are three sources of crosstalk: trace to trace, via coupling between signals at the FPGA end and via coupling at the DRAM end.

- **Trace to Trace Crosstalk Optimization** – It is recommended to space out the traces as much as possible to reduce crosstalk between them. It might be a good idea to spread the signals in multiple layers to reduce trace to trace crosstalk. A crosstalk threshold of at least -30 dB for both far-end and near-end crosstalk is recommended,

- **Via Crosstalk Optimization** – If the crosstalk is high, it is a good idea to assess the via-to-via coupling and use ground microvias between the signals to reduce via crosstalk. As shown in the figure below, a ground microvia has been added between diagonally placed data signals to provide ground shielding and reduce crosstalk. It might not be possible to add ground shielding vias between every two signals, and therefore, the designer must assess the crosstalk for all signals and add ground shielding vias for cases where crosstalk between signals is high and meeting the crosstalk budget is otherwise difficult.

> **Note**
>
> The figure below shows one ground via for four signal vias which is not a guideline from Achronix. The number of ground vias required should be determined based on crosstalk assessment of the PCB layout.
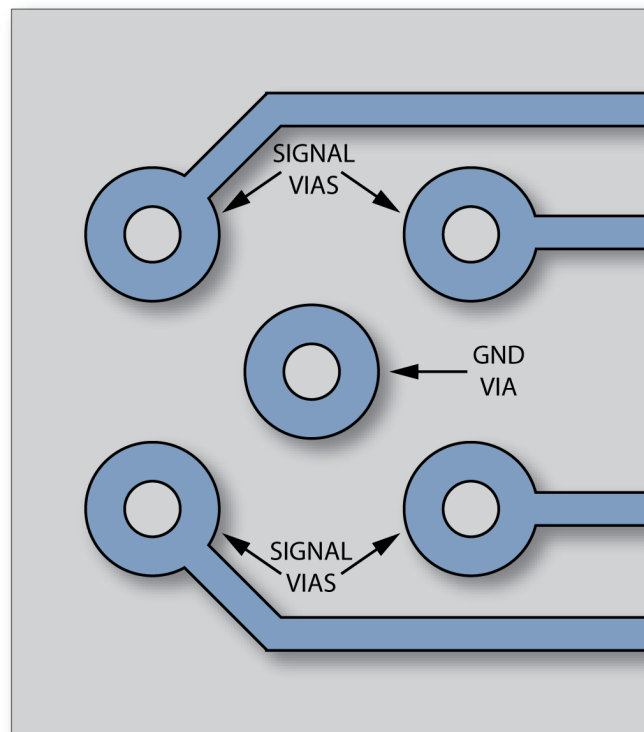
70540420-12.2020.11.30

Figure 33: *Via Crosstalk Optimization*

Routing Guidelines

The following guidelines are specific to the routing of the GDDR6 high-speed interface:

- It is recommended that single-ended data traces be routed to a characteristic impedance of 50Ω. As the package trace impedance is also 50Ω, having 50Ω on the PCB ensures impedance does not change between package and PCB. However, designers may run signal integrity simulations to find out the best system performance.

- Crosstalk mitigation is important. If via crosstalk is difficult to contain, improve crosstalk performance by trace separation and/or routing on different layers. Routing signals on different layers is recommended only if the two layers are sandwiched between the same dielectric medium. If the dielectric mediums are different, a careful assessment of propagation delay is needed to ensure skew limits are not violated.

- Provide sufficient return vias in proper proximity to power vias to reduce the power delivery network loop inductance. Optimize signal via transitions to nominal impedance (50Ω/100Ω). Ensuring the requirement might necessitate use of a 3D electromagnetic field solver.

- Differential signals must be length matched within the pair for the complete length in the channel. Any skew generated on the differential pair should be addressed at the earliest possible place in the layout. Signal traces should be designed to minimize skew between P and N traces of a differential pair. It is also critical to maintain symmetry between the true and complement trace of the differential pair to minimize mode conversion and skew.

> ⚠ **Caution!**
>
> DRAM routing strategies often include swapping bits within a byte lane. *Do not swap bits for GDDR6 as it is not designed to support this strategy.*

For general routing guidelines, see the section, "Routing Guidelines (see page 15)".

Chapter - 6: Speedster7t AC7t1500 PCB DDR4 Interface

DDR4 SDRAM has various advantages over its predecessors, including higher module density and lower voltage requirements, as well as higher data rate transfer speeds. The DDR4 standard allows for DIMMs of up to 64 GB in capacity and transfer rates as high as 3200 MT/sec. This high frequency of operation requires the package and PCB design to be optimized for minimal losses and minimal crosstalk.

This chapter provides PCB design guidance on definition, placement, and routing of the DDR4 interface and focuses on electrical design parameters to optimize the PCB design requirements to meet the electrical specs and reduce the electrical parasitic effects for the interface. While designers have the option of designing their memory onto the main board (the same PCB as the FPGA), this discussion addresses memory modules, specifically unbuffered DIMMs (UDIMM), registered DIMMs (RDIMM), load-reduced DIMMs (LRDIMM) and small-outline DIMMs (SO-DIMMS) that can provide signal integrity advantages, but the designer is directed to JEDEC Standard 21C for more detail on those DIMM technologies.

> **Note**
>
> ⓘ Designers working with soldered-down memory solutions might want to consider replicating the DIMM layout on their PCB as close as reasonable, remembering that DIMMs are very thin and can have much smaller vias. The resulting impact on signal integrity should not be ignored.

DDR4 Channel Topologies

The different signal groups, data and strobe vs. Add/Cmd/Ctrl and clock, have different loading configurations leading to different routing requirements. A few topologies for DDR4 configuration, for DQ/DM and Address/CMD signals are discussed below:
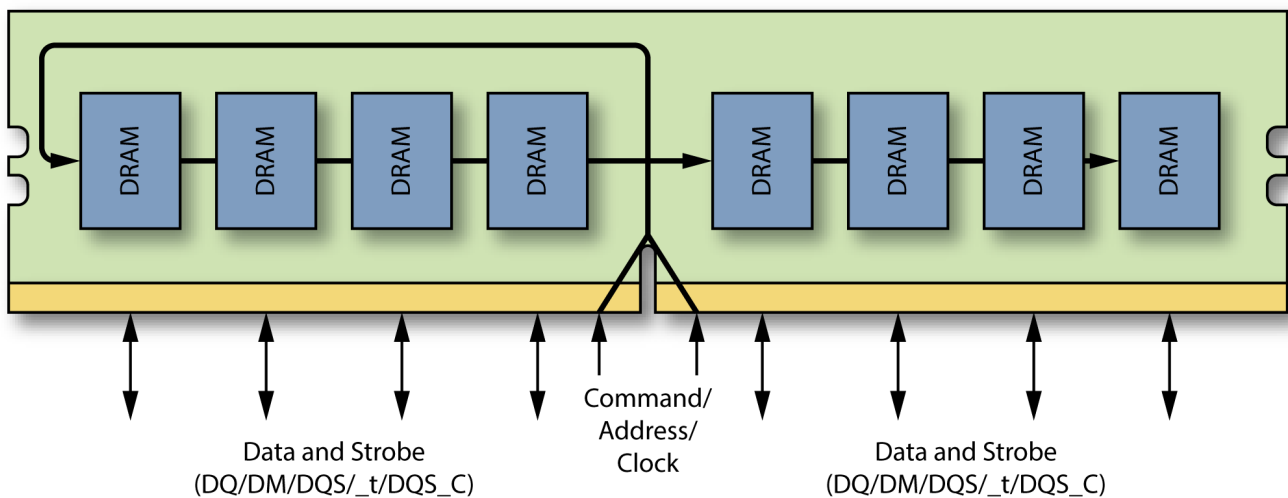
70540288-01.2020.11.23

Figure 34: *Fly-By Routing Used in Unbuffered DIMMs (Single Rank Shown)*

DQ/DM/DQS Topology

For single-rank applications, the topology is point-to-point and can be implemented with little special consideration — this layout can readily be placed on a four layer board.

In order to maintain the source-synchronous relationship between DQ, DM and DQS, it is important that each signal be routed in the same layer and contain the same number of vias if applicable. Two-rank applications require more scrutiny and are implemented using T or fly-by topology. The figure below shows point-to-point topology to be used for Rank1/Rank2 memory for data signals.
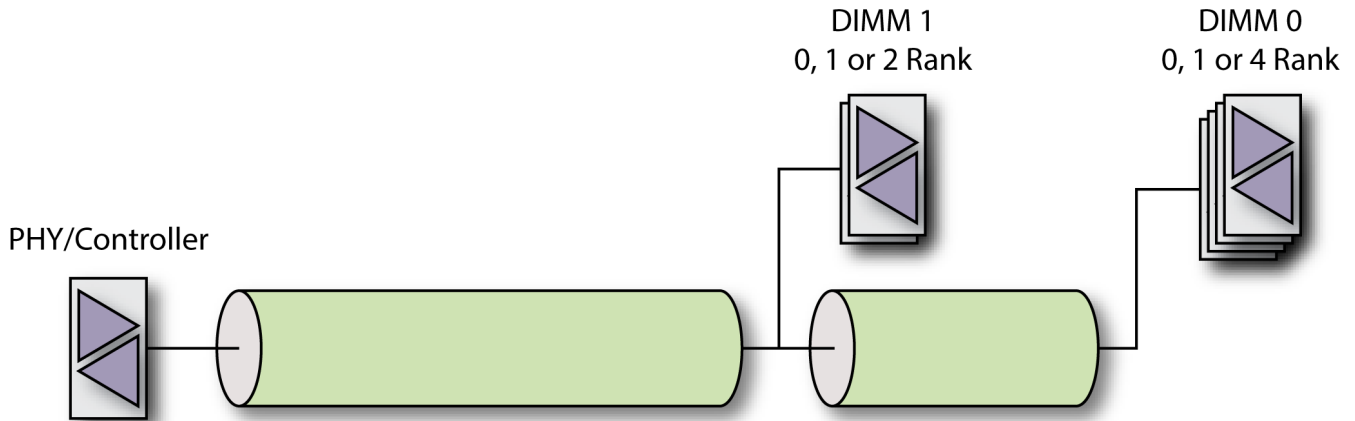
Figure 35: *DQ/DQS/DM Point to Point*

Address and Command Topology

The DDR4 specification includes functionality that allows Add/Cmd/Ctrl signals and CK/CK# signals to be routed as long daisy chains using fly-by termination. The figure below illustrates how these daisy chain (fly-by) and "T" routings are accomplished on an two-rank memory. Care must be taken with the fly-by technique to adequately match the impedance of this line to the termination at the far end of the net to minimize reflections at the devices along the route.
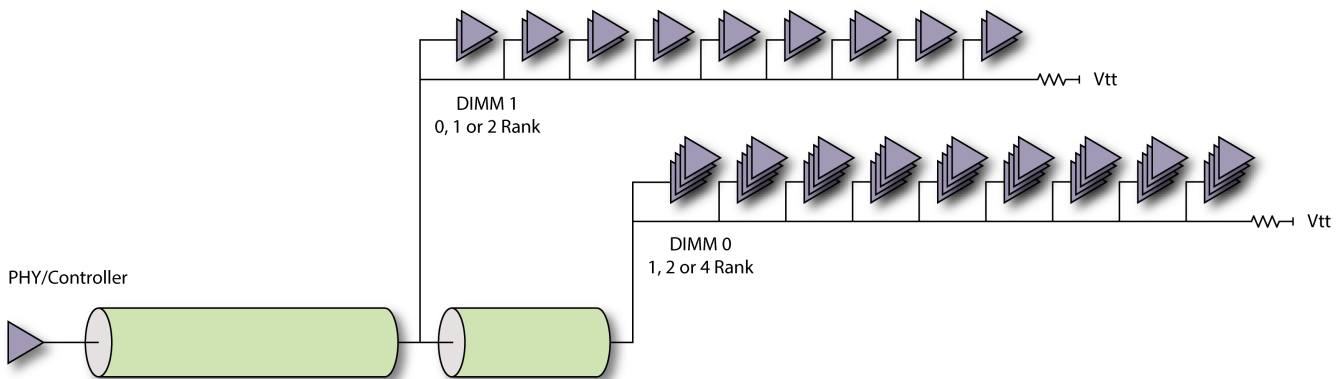
Figure 36: *Unbuffered DIMMs Implemented in Two-Slot Daisy-chain Configuration*

Differential Clocks CK/CK# Topology

For the clock network, a daisy-chain topology is preferred for both single- and dual-DIMM configurations. The figure below shows the clock network for a dual-DIMM configuration. Differential termination resistors referenced to V_{CC} are placed at the far ends of both DIMMs.
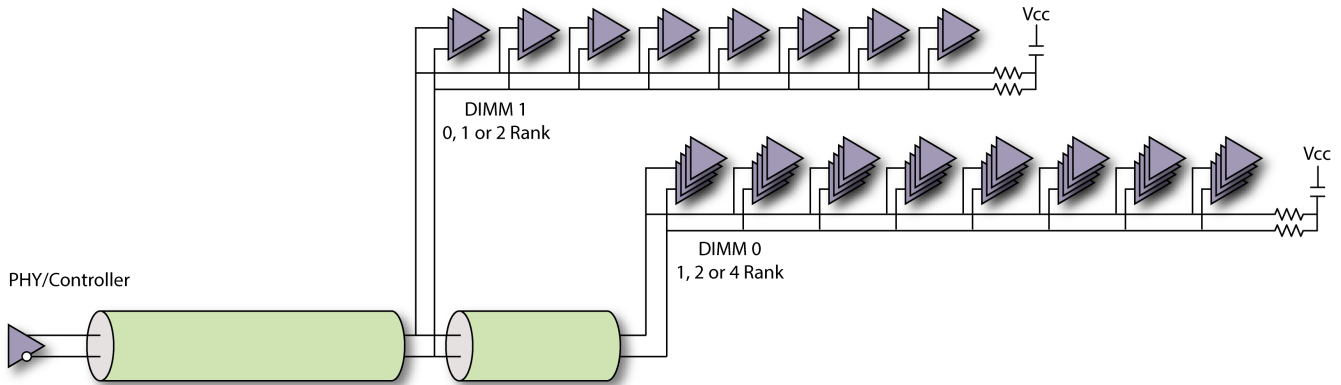
Figure 37: *Daisy-Chain (Fly-By) Topology for Differential Clock in a UDIMM*

Signal Integrity Specification

Signal Integrity PCB Design Specification Guidelines

Table 6: *Signal Integrity Specifications for 3200 MTps Single-Rank DDR4*

Memory Configuration	Signal Group	PCB Maximum Insertion Loss	PCB Maximum Return Loss	Near-End Crosstalk (NEXT)	Far-End Crosstalk (FEXT)
Single-rank 3200MTps	Data	> −1 dB @ 3.2GHz	< −20 dB @1.6 Ghz	−35 dB @ 1.6 GHz	
	DQS				
	CA	> −1 dB @ 1.6GHz	< −17 dB @ 1.6 GHz		
	CLK	> −1 dB @ 3.2GHz	< −20 dB @ 1.6 GHz		

Delay Matching

Delay matching on the DDR4 bus is governed by the JEDEC DDR4 specification, which specifies timing at the device and controller pins. Achronix provides package delays so that the delay of the entire path can be matched, which results in a more complete solution to the timing problem. For 3200 MTps single-rank operation (from the JEDEC Spec), the key specifications are:

- The data UI is $1/3200 \times 10^6$ = 312.5 ps. The command/address (CA) UI is twice that, or 625 ps.

- t_{DQ2DQ} is the receive mask DQ-to-DQ offset, or the overall matching of the DQ signals, and is specified at 0.125 UI, or 39.1 ps. Some of this offset is consumed by the module and connector, so it is suggested that DQ be matched within a byte at less than 20 ps on the PCB.

- t_{DQS2DQ} is the receive mask DQS-to-DQ offset,and is specified at ±0.220 UI, or ±68.8 ps. This value can be seen as the worst-case placement (in time) of the DQS relative to the data. Within a byte, it is suggested that DQS be matched to the data within ±10 ps on the PCB.

- t_{DQSS} is the relationship of the clock to the DQS. The DQS is required to be within ±0.270 UI (using the clock UI), or 168.8 ps. This relationship is complicated in the case of UDIMMs, which employ a fly-by topology for the CA bus and clocks. The Achronix controller provides a DLL which adjusts the DQS-to-CLK delay from -UI to 5.97 UI. In order to obtain maximum flexibility and accommodate the greatest number of DIMMs, it is recommended to keep each DQS within ±85 ps of the clock on the PCB.

- t_{IS} and t_{IH} are the CA setup and hold times. There is, at the time of this writing, no specification for these values at 3200 MTps, but they are extrapolated from lower specs to be 108 ps each. Dividing the data UI by two and subtracting the setup and hold times results in an overall matching requirement of <204.5 ps, or ±102 ps. The controller groups the CA and clock signals into bit groups of four signals each, which can be delayed independently. It is suggested the delay matching within a bit group be kept to less than 3 ps, but of the significant routing uncertainty of the fly-by CA routing, and overall PCB delay matching for the CA traces of ±25 ps is recommended.

Table 7: *Bit Groups for Address Matching*

Group No.	Signals
1	CKE0, CKE1, CKE2, CKE3
2	BG0, BG1, ACT_N, A09
3	A12, A11, A07, A08
4	A06, A05, A04, A03
5	CLK0_T, CLK0_C, CLK2_T, CLK2_C
6	CLK1_T, CLK1_C, CLK3_T, CLK3_C
7	A02, A01, BA1, PAR
8	A13, BA0, A10, A00
9	CAS_N, WE_N, RAS_N, C0
10	C1, C2, A17, CS_N0

DDR4 Signal Integrity Sign Off Simulations

Signal Integrity Sign-off

To determine compliance of the PCB to the specification, a worst-case PCB model for data signals should be used. Models of fly-by topology with signals connected to the furthest DIMM side are the worst-case board model.

- Read and write cycle channel simulations to confirm BER compliance and eye mask. Run single ranks at 3.2 Gbps, dual ranks at 2.4 Gbps and quad ranks at 1.6 Gbps.
- Command/address transient/crosstalk simulation for the topology chosen.

DDR4 Write Cycle

Simulation Setup

This topology represents the entire channel from transmitter (Achronix FPGA) to receiver (DRAM). The topology shown below has a single DIMM. If a second DIMM is used, the topology must be adjusted accordingly. The topology includes the following:

- **Transmitter (TX) IBISbehavioral models** – These models are specific to the ACt1500, and are available from Achronix upon request.

- **Package S-parameter models** – These models cover the AC7t1500 package, from the silicon bumps to the package pins and consist of one data byte, and are available from Achronix upon request.

- **PCB S-parameter model** – These models cover from PCB pin to DRAM module pin and must be provided by the PCB designer.

- **Connector S-parameter** - This model covers the DIMM/SO-DIMM connector, if present, and must be provided by the connector manufacturer.

- **Module EBD models** – This extracted model covers the topology of the module (if present), and must be provided by the DRAM module vendor.

- **Receiver (RX) IBISbehavioral models** – These models are specific to the DRAM vendor and are available from Achronix upon request.

For final high-speed SerDes PCB sign-off, it is necessary to verify that the channel performance meets the given interface specification in both the transmitter and receiver direction.

For final high-speed SerDes PCB sign-off, it is necessary to verify that the channel performance meets the given interface specification in both the transmitter and receiver direction.

70540288-05.2020.12.13
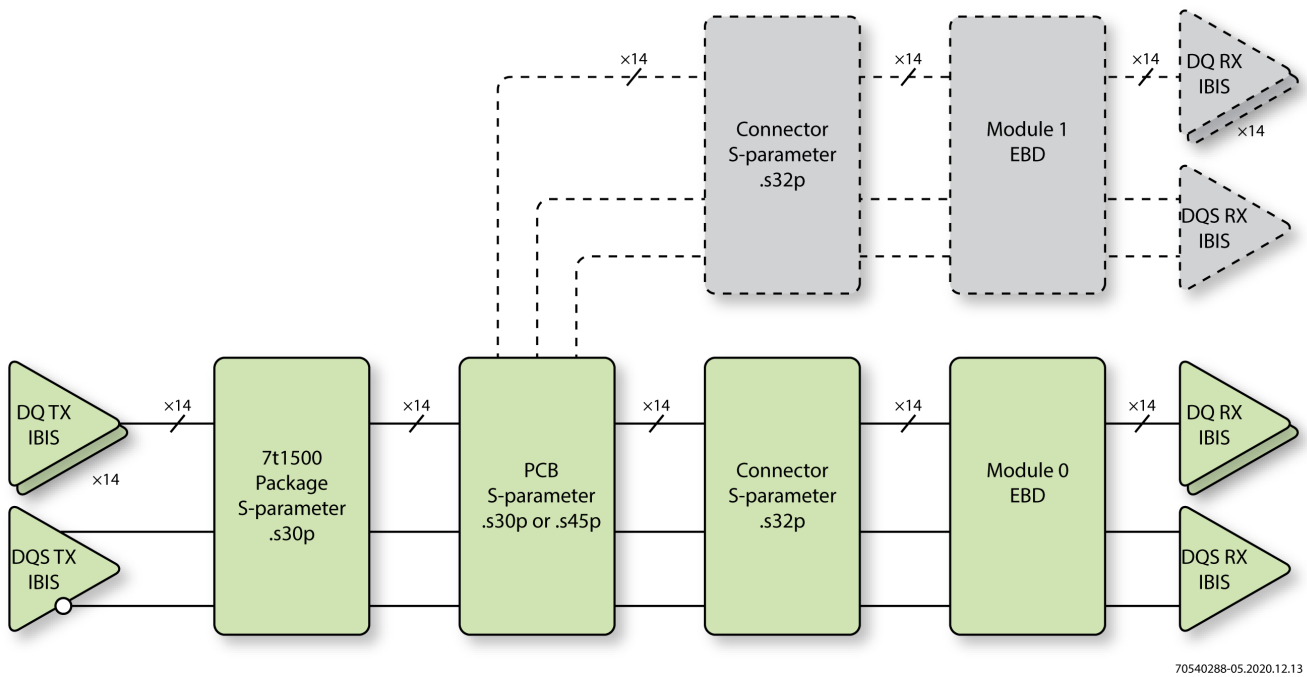
> **Figure Note**
>
> Solid lines represent singe-module configuration. Dashed lines represent an optional second module.

Figure 38: *DDR4 Write Simulation Setup*

Signal Integrity Specification for Data Write Cycle

The eye mask shown below is taken from JEDEC spec (JESD79-4B). Refer to latest JEDEC Spec for the eye mask requirement at the DRAM component.
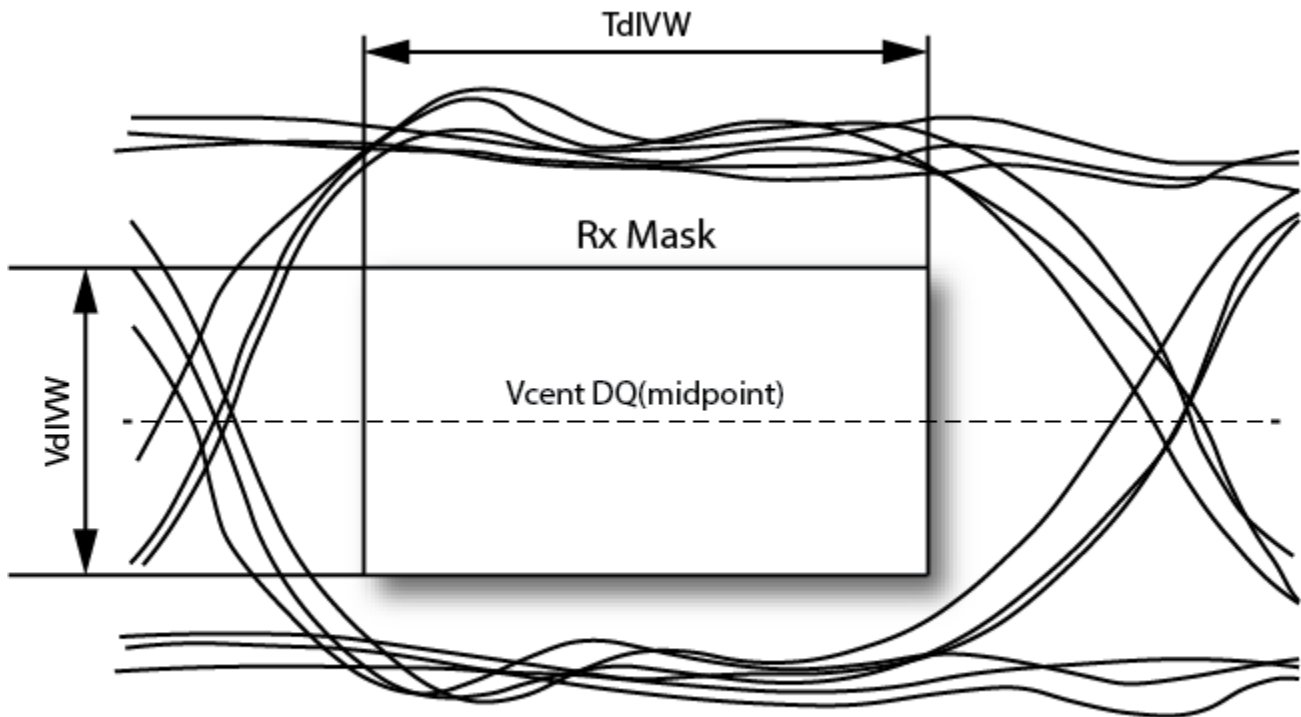
70540288-06.2021.10.10

Figure 39: *Receive Compliance Requirement*

Key DDR4 Eye Diagram Terms:

- V_{CENT} **Value** – Defined as the midpoint between the largest DQ reference voltage and the smallest DQ reference voltage level, computed using voltage training.
- **UI (Unit Interval**) – 312.5 ps for 3.2 Gbps operation of DDR4
- **VdivW (RX Mask Voltage)** – Greater than 110 mV
- **TdivW (RX Timing Window)** – Greater than 0.23 UI = 71.875 ps

Jitter Analysis

Table 8: *Transmit (Write) Jitter Budget Tables*

Component	Setup (ps)	Hold (ps)	Notes
UI width at 3200 Mb/s	156.3	156.3	
Data-dependent jitter			
Output rise/fall mismatch	5.6	5.6	Delay differences between rising and falling edges after training.
V_{DD} -PSIJ (±2.5%)	9.3	9.3	Jitter induced by noise on the V_{DD} rail.
V_{DDQ} -PSIJ (±5%)	6.0	6.0	Jitter induced by noise on the V_{DDQ} rail.
Training error			
Strobe alignment error	13.8	18.9	Alignment of strobe in data.
Aging	0.4		Aging of delay lines.
PLL jitter	2.5	2.5	RefClk feed-through, internal noise, supply modulation.
Total transmit components	**37.5**	**42.2**	

> **Table Note**
>
> - V_{DD} = 0.808V (0.85V Nominal) — Linear Summation

Table 9: *Receive (Read) Cycle Budget Tables*

Component	Setup (ps)	Hold (ps)	Notes
UI width at 3200 Mb/s	156.3	156.3	
Data-dependent jitter			
Input rise/fall mismatch	1.9	1.9	Difference between rising and falling edge delay at input.
V_{DD} -PSIJ (±2.5%)	12.3	12.3	Jitter induced by noise on the V_{DD} rail.
V_{DDQ} -PSIJ (±5%)	1.0	1.0	Jitter induced by noise on the V_{DDQ} rail.
Training Error			
Strobe alignment error	13.6	18.4	Alignment of strobe in data.
Aging	2.9	0.0	Aging of delay lines.
V_{REF} Error	2.0	2.0	V_{CENT} accuracy and supply noise.
Flop SU/Hd requirements	7.0	4.0	
Total receive components	**40.7**	**39.6**	
DRAM transmitter data-valid window	43.8	43.8	Uncertainty based on $(1 - t_{DVWp})$ × UI from JEDEC standard.
Input Period Jitter Derating	16.0	16.0	Derate by the negative peak specification.
Total DRAM transmit components	**59.8**	**59.8**	
Interconnect allowance	55.8	56.9	Interconnect uncertainty from simulation measured at PHY pad V_{REF} crossing.

> **Table Notes**
>
> - V_{DD} = 0.808V (0.85V Nominal) — Linear Summation
> - Budget assumes all error terms correlated at worst case values — a highly pessimistic assumption. Silicon performance is better than the numbers reflected here.

DDR4 Read Cycle

Simulation Setup

This topology represents of the entire channel from transmitter(DRAM) to receiver (Achronix FPGA) and includes the following:

- **Transmitter (TX) IBIS behavioral models** – These model are specific to the DRAM vendor and are available from Achronix upon request.

- **Package One-Byte S-parameter models** – These models cover the AC7t1500 package, from the silicon bumps to the package pins and consist of one data byte, and are available from Achronix upon request.

- **Board S-parameter model** – These models cover the AC7t1500 package, from PCB pin to DRAM present on the board and must be provided by the PCB designer.

- **EBD/DRAM S-parameter models** – This extracted model covers of the topology of the receiver system and must be provided by the DRAM vendor.

- **Receiver (RX) IBIS behavioral models** – These models are specific to the AC7t1500 FPGA and are available from Achronix upon request.
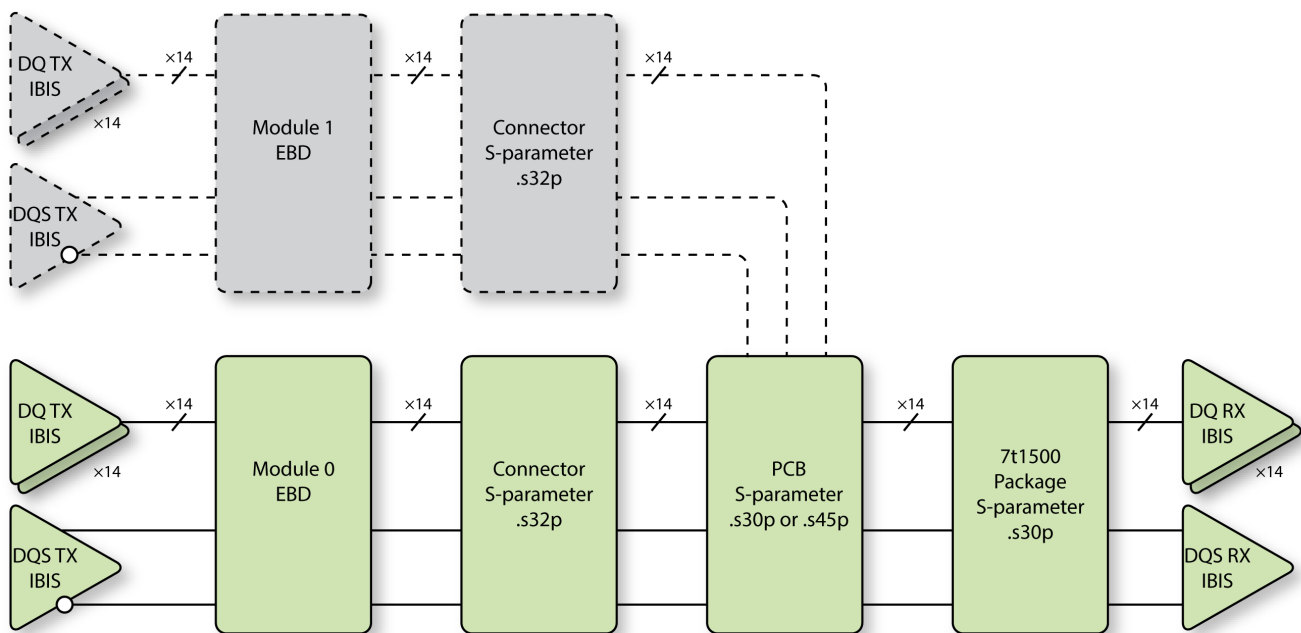
Figure 40: *DDR4 Read Simulation Setup*

Signal Integrity Specification for Data Read Cycle

GDDR6 data signals received by the AC7t1500 FPGA must comply to the following eye mask:

- Minimum eye height requirement (at AC7t1500 device die pin level) @ 1e-16 BER = 110 mV

- Minimum eye width (setup+hold) requirement (at the AC7t1500 device die pin level) @ 1e-16 BER = 199.8 ps

Contact Achronix Support for direction on creation of eye mask for the read/write cases in specific scenarios.
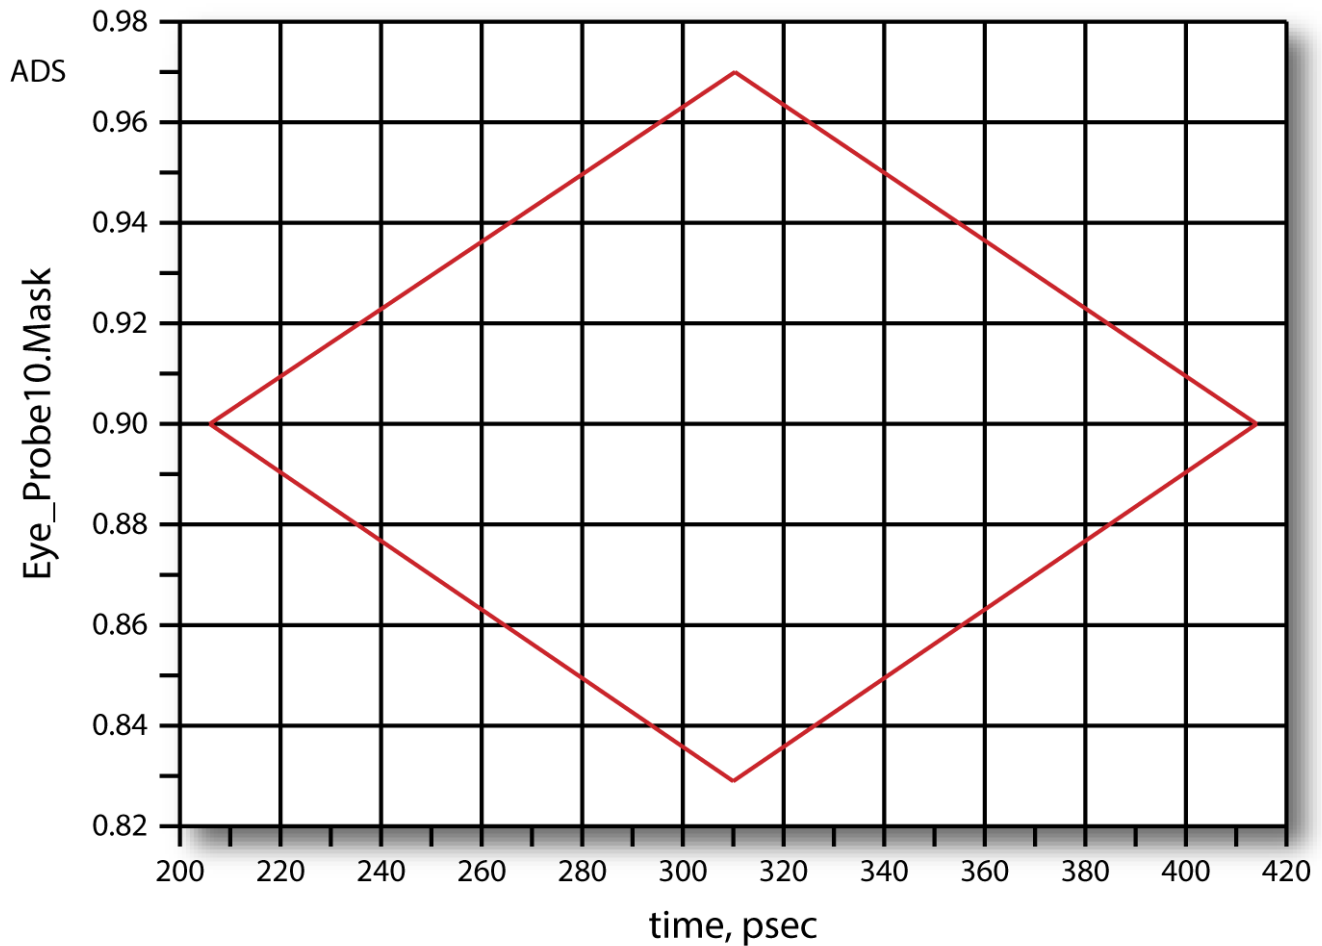
70540288-08.2020.12.13

Figure 41: *Read Cycle Eye Mask*

DDR4 Command Address Write Cycle

Sign-Off Simulation

This topology represents of the entire command/address channel from transmitter (Achronix FPGA) to receiver (DRAM). The topology shown has a single–DIMM configuration in solid lines. An optional second DIMM is presented in dashed lines. The topology includes the following:

- **CA Transmitter (TX) IBISbehavioral models** – These models are specific to the AC7t1500 FPGA and are available from Achronix upon request.

- **Package S-parameter models** – These models cover the AC7t1500 package, from the silicon bumps to the package pins and consist of one data byte, and are available from Achronix upon request.

- **PCB S-parameter model** – These models cover from PCB pin to DRAM module pin and must be provided by the PCB designer.

- **Connector S-parameter** - This model covers the DIMM/SO-DIMM connector, if present, and must be provided by the connector manufacturer.

- **Module EBD models** – This extracted model covers the topology of the module (if present), and must be provided by the DRAM module vendor.

- **Receiver (RX) IBISbehavioral models** – These models are specific to the DRAM vendor and are available from Achronix upon request.

For final high-speed SerDes PCB sign-off, it is necessary to verify that the channel performance meets the given interface specification in both the transmitter and receiver direction.

Final sign-off criteria must also meet the specifications provided by JEDEC at the DRAM component.
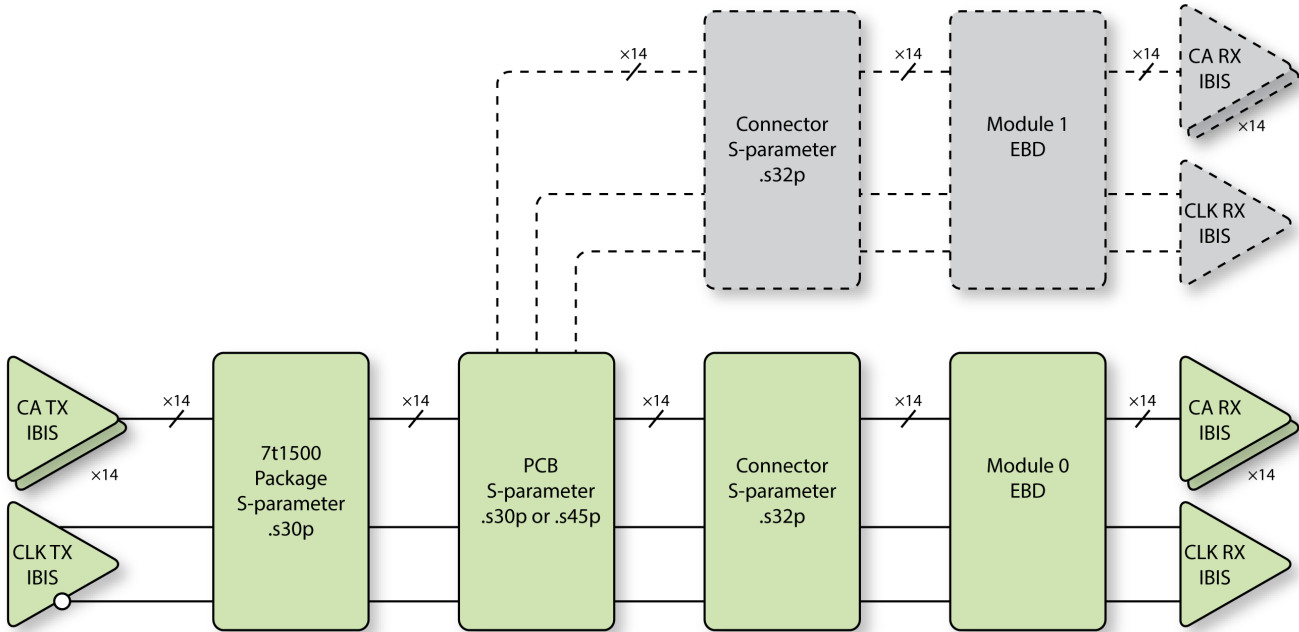
Figure 42: *DDR4 Command Address Write Setup*

Layout Optimization Guidelines

Stack-Up Guidelines

DDR4 has special considerations for stack-up planning:

- Since each routing layer has propagation delay and impedance variations, signals within a given functional group should route on the same layer with the same geometry. Otherwise, precautions must be taken to ensure delay and impedance matching.

- Reference planes:
    - DDR4 data signal layers are recommended to be sandwiched between ground planes.
    - The command/address bus is referenced to V_{DDQ}. It is recommended that the CA bus be sandwiched between a ground plane and a V_{DDQ} plane.

- Single-ended DDR4 signals should have an impedance of 50Ω, and differential signals to be designed for 100Ω impedance.

- Multiple factors such as dielectric and conductor material, channel length and waveguide geometry affect the trace impedance. It is better to use a low-loss PCB material to reduce the insertion loss.

See the section, "Board Construction - the Stack-up (see page 8)", for general direction on constructing the PCB stack-up

Routing Guidelines

The following guidelines are specific to DDR4 routing:

- It is recommended that the single-ended signals on the PCB be routed to a characteristic impedance of 50Ω, and the differential signals on the PCB be routed to a characteristic impedance of 100Ω (these impedance values match those of the AC7t1500 package traces). However, based on other constraints and design requirements, the PCB designer can choose different impedance values if the signal integrity performance is acceptable.

- Optimize selection of drive strength and on-die termination to achieve the best signal integrity performance.

For general routing guidelines, see the section, "Routing Guidelines (see page 15)".

Chapter - 7: GPIO, SPIO, CLKIO, and Miscellaneous Signals

GPIO, SPIO, CLKIO Interfaces

This chapter discusses board design considerations for the general-purpose I/O (GPIO), special-purpose I/O (SPIO) and CLKIO (REFIO and MSIO) . It includes best practices for these interfaces.

GPIO Interface

GPIO pins enable communication between the FPGA fabric and external components. Speedster7t FPGAs provide a variety of GPIO features and supported I/O standards. These features are detailed in the *Speedster7t GPIO User Guide* (currently in progress). Some features which may impact a board designer are included below:

Table 10: *Supported Signaling Schemes*

Category	Signaling Scheme	Comments
LVCMOS	LVCMOS18 LVCMOS15 LVCMOS12 LVCMOS11	Voltage levels (V_{OH}, V_{IH}, V_{OL}, V_{IL}) vary according to V_{DD} (see JEDEC specs). For example, LVCMOS12 V_{OH} may not reach the LVCMOS18 V_{IH} threshold. Take caution when crossing LVCMOS V_{DD} domains.
SSTL	SSTL18 (Class I and II) SSTL15, SSTL135, SSTL12 (Class I) DIFF_SSTL18 (Class I and II) DIFF_SSTL15 (Class I) DIFF_SSTL_135 (Class I) DIFF_SSTL12 (Class I)	Support for this standard requires two adjacent MSIO macros configured to form a pseudo-differential transmit or receive pair.
HSTL	HSTL_18 HSTL (Class I and II) DIFF_HSTL_18 DIFF_HSTL	
HSUL	HSUL_12 DIFF_HSUL_12	

Table 11: *GPIO Transmit Configurability*

Transmit Lane Parameter	Design Specification
LVCMOS Drive Strength	Pin programmable drive strength (2/4/6/8/10/12/14/16mA) at 1.8V I/O supply.
ODT for Class II Operation	Transmit supports split Thévenin ODT in order to enable Class II (parallel termination at both the transmit and receive ends) operation.
Driver Impedance	Transmit supports driver impedance calibration in order to provide on-chip equivalent of series termination.
High-Z, Pull-up, Pull-down	The driver supports tri-state (high-impedance) output mode. While in High-Z mode, pull-up or pull-down impedances can also be enabled.
Pseudo-Differential Mode	Two MSIO macros are required to form a pseudo-differential pair: • One configured as the differential master (idat_i_a active) • One configured as the differential slave (idat_i_a disabled)
Slew-rate control	In addition to output drive strength/impedance configurability, the MSIO has two bits of pre-driver slew-rate control.

Table 12: *GPIO Receive Configurability*

Receive Lane Parameter	Design Specification
ODT for Class I/II Operation	Receive supports split Thévenin ODT in order to enable Class I and Class II operation.
Schmitt Trigger	Receive supports pin configurable Schmitt trigger capability.
Pseudo-Differential Mode	Two MSIO macros are required to form a pseudo-differential pair: • One configured as the differential master (odat_c_a active, output level translator driven by outputs from both master and slave) • One configured as the differential slave (odat_c_a disabled)

GPIO ZCAL

ZCAL configuration allows control of an output driver's impedance, greatly assisting driving controlled impedance traces/transmission lines. Single-ended impedances of 40Ω/50Ω/60Ω can be programmed. If a signal is designated as controlled impedance, this feature should certainly be put to use.

For differential transmission lines, the impedance can be programmed to 100Ω. Refer to the *Speedster7t GPIO User Guide* (currently in progress) for further details.

GPIO VREF

The V_{REF} feature provides a flexible way to generate a reference voltage for logic level thresholds. V_{REF} can be generated in several ways, including from a Thévenin voltage divider for $V_{DDH}/2$, and through program control, from $V_{DDH} \times 0.3$ up to $V_{DDH} \times 0.7$. Refer to the *Speedster7t GPIO User Guide* (currently in progress) for further details.

SPIO Interface

The DDR4 controller can be placed in bypass mode, which allows DDR4 I/O to be repurposed as SPIO. Refer to the *Speedster7t GPIO User Guide* (currently in progress) for further details.

CLKIO Interfaces

REFIO interface

The reference clock differential I/O (REFIO) interfaces provide inputs capable of differential LVCMOS_18, VCMOS_15, LVDS_18, LVDS_15, and LVPECL, up to a frequency of 600 MHz with a buffer jitter of less than 0.3 ps RMS from 12 kHz to 75 MHz. As outputs, the REFIO interfaces can clock up to 1 GHz.

> **Note**
>
> When using LVPECL logic, the signals must be AC coupled on the board, i.e., DC blocking caps must be used between the clock source and the FPGA.

MSIO interface

The multi-standard I/O (MSIO) interface provide inputs that run up to a frequency of 500 MHz. Each MSIO pad can be independently configured to input one clock or one reset, or output one clock. The MSIO pads may also be configured as a pseudo-differential pair. The interface supports LVCMOS_18, LVCMOS_15, HSTL_18, SSTL_18, and SSTL_15 I/O standards.

Refer to the *Speedster7t Clock and Reset Architecture User Guide* (UG083) for further details.

Miscellaneous

FTDI Interface for FPGA Programming

Achronix ACE software can program the Speedster7t AC7t1500 FPGA via a USB cable. To facilitate this, it might be useful to include a simple USB interface. Consult Achronix Support for a proven USB interface design with control, data and clock signals tied to specific pins of an FTDI interface chip for FPGA and flash memory programming. Please refer to the *Speedster7t Configuration User Guide* (UG094) and the *Bitstream Programming and Debug Interface User Guide* (UG004) for details on FTDI and JTAG programming.

Table 13: *FTDI Connections to the AC7t1500 FPGA*

AC7t1500 Signal	Pin	FTDI Signal	Pin
FCU_CPU_DQ1	AF36	FTDI_AD1	17
FCU_CPU_DQ9	AT35	FTDI_AD2	18
FCU_CPU_DQ17	BA32	FTDI_AD3	19
FCU_CPU_DQ25	AW31	FTDI_AD4	20

JTAG Interface features

The AC7t1500 FPGA supports industry-standard JTAG, with the following signals:

Table 14: *AC7t1500 JTAG Pin*

Signal	Pin	Direction	Pull-up/Pull-down
JTAG_TCK	AR27	Input	
JTAG_TDI	AW25	Input	
JTAG_TDO	AU27	OUT	
JTAG_TMS	AP26	Input	
JTAG_TRSTN	AV26	Input	Weak (100 kΩ)

Board Layout Concerns for GPIO, SPIO and CLKIO Interfaces

Layout and routing of low-speed signals (i.e. <300 MHz, or >1 ns edge rate) is much simpler than the multi-gigabit signals of the faster Speedster7t interfaces. Nonetheless, careful attention to best practices, even at "low speeds", assure a working design.

- Unless a signal is controlling something simple, slow, and/or short, such as a blinking LED, it is best to specify an impedance for the signal. Specifying a controlled impedance assures a predictable waveform at the receiver every time, and helps to control EMI.

- Avoid routing signals with stubs. For slow serial interfaces such as I^2C, this is not so important for data (SDA), but is surprisingly important for clock (SCL). A long stub (for which the propagation delay is longer than the transition time of the signal) can cause a non-monotonic clock edge, meaning the receiver may see multiple transitions across the threshold voltage, and the logic can "double-clock". With today's smaller and smaller silicon geometries, receivers are more and more sensitive to this issue. It is recommended to simply guarantee that the clock signal cannot be seen at the receiver as anything but a single rising or falling edge. Again, it cannot be overstated that careful design of clock routes, even in "slow" interfaces, is essential.

- Control the current drive of the output. The current drive determines the signal's edge rate, which can impact reflections, crosstalk, and monotonic edges. A data line may tolerate a fast edge, but it may cause EMI and crosstalk with other data channels and with RF circuits. The Speedster7t I/O can be programmed to control output current to a fine degree. For even tighter control, consider terminating with a series resistor close to the clock output pin to match the output to the trace impedance by effectively controlling the driver's current.

- Always route differential pairs with a designed geometry in order to achieve a controlled impedance. For the GPIO, SPIO and CLKIO differential pairs, the output impedance is 100Ω. This requirement can easily be met in designs already using 50Ω single-ended traces. Standard differential pair routing rules apply.

- When using multiple LVCMOS logic levels (i.e. 1.8V, 1.5V, 1.2V, 1.1V), pay attention to logic input and output levels, i.e., V_{IL}/V_{IH} and V_{OL}/V_{OH}. An LVCMOS11 output, at 1.1V V_{DD}, may not cross the V_{IH} logic threshold of an LVCMOS18 input, at 1.8V. In such instances it is necessary to incorporate the use of logic-level translators between the driver and receiver to ensure error-free transmission.

Chapter – 8: Power and PDN Design

Power Distribution Network

The power distribution network (PDN) is the circuitry engineered to provide all the power requirements of the digital load, including the instantaneous current demand of thousands of transistors switching simultaneously. It encompasses everything from power conversion, copper planes to distribute the current, capacitors that serve as local charge storage, the package, including the pins, vias and copper planes, the bumps of the die, the metal layers dedicated to distributing the current, and on-die capacitance. Each voltage required by the load demands a separate PDN.

The figure below illustrates a typical PDN.

70536065-01.2020.12.26
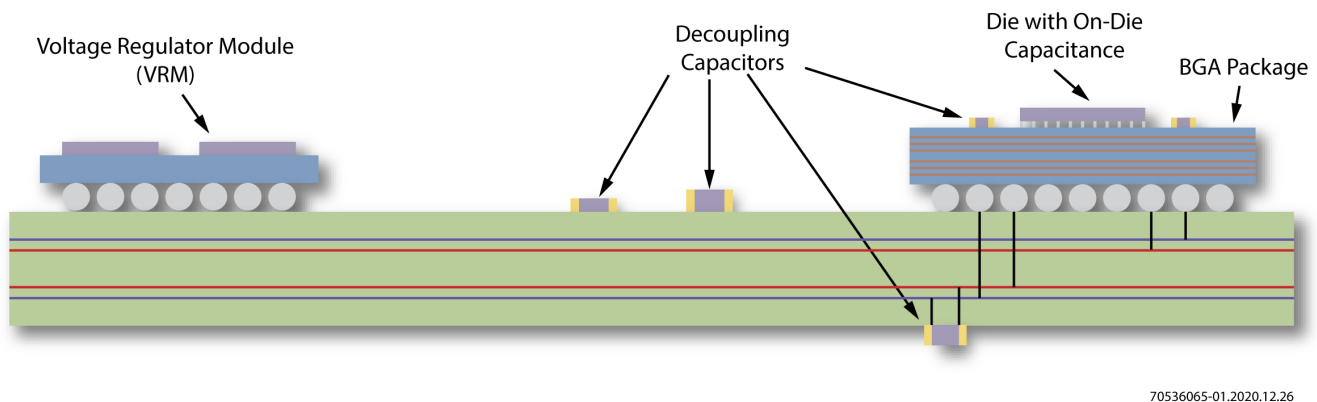
Figure 43: *A Typical Power Distribution Network*

In the figure above:

- **Voltage Regulator Module (VRM)** – The VRM is responsible for generating the required voltage, generally converting it from a higher voltage source. It generates all the current required at that voltage, and must be very low impedance up to several hundred kHz.

- **PCB** – The PCB connects the VRM to the load. In an advanced PDN, the PCB has thick copper planes to carry the substantial current, with additional copper ground planes to carry the return current. Vias conduct current down to the planes from the VRM.

- **Decoupling Capacitors** – Various sizes and values of decoupling capacitors are used to provide wells of charge that respond faster than the VRM. Vias connect these capacitors to the planes.

- **BGA Package** – The package carries the silicon die which has the switching circuitry being supplied (the load). It can be viewed as a small PCB, with copper planes, decoupling capacitors and vias of its own.

- **Silicon Die** – The FPGA die carries the switching circuitry, connected to the PDN by bumps and metal layers. There is also additional on-die capacitance to provide the smallest but fastest-response charge wells.

Robust PDN Design Steps

Frequency Domain Target Impedance

In order to meet ripple specifications below, the PDN should be designed according to the "Frequency Domain Target Impedance" methodology. This methodology is well documented both in conference papers (DesignCon) and in academic text books. The methodology considers the frequency response of the PDN, and holds that the PDN must meet the needs of the load at all frequencies (not just DC). Those needs are defined by the Target Impedance, which is the ratio of the maximum allowed deviation from the nominal voltage to the maximum instantaneous current:

$$Z_{Target} = \Delta V_{MAX} / \Delta I_{MAX}$$

The electrical equivalent model for the PDN can be shown as illustrated below.
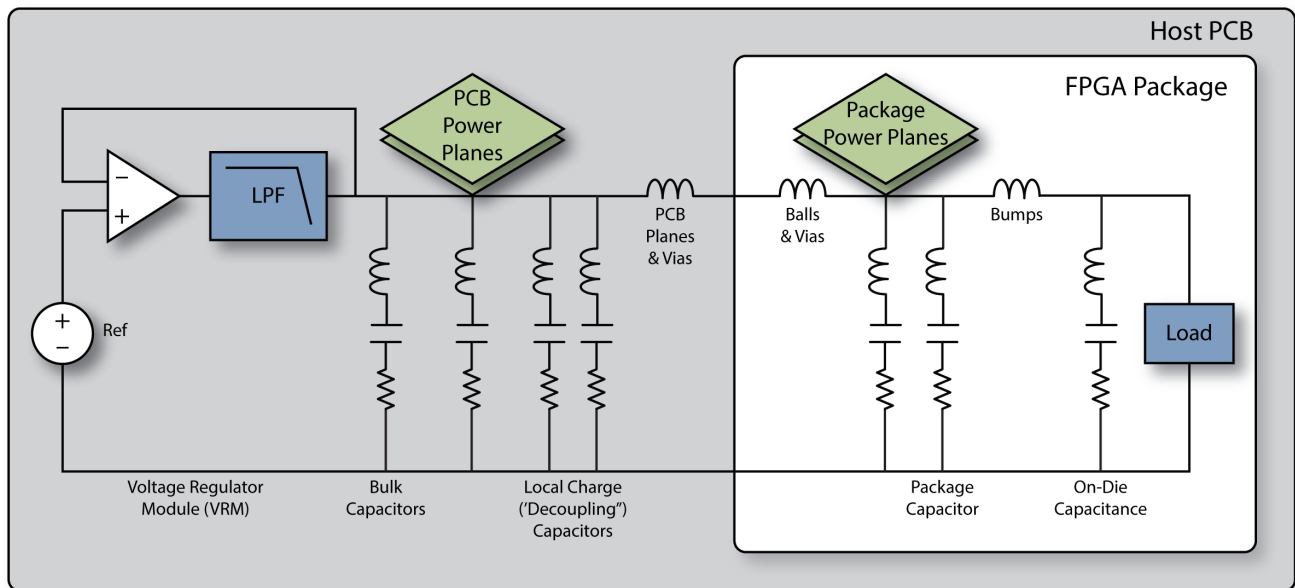
70536065-02.2020.12.26

Figure 44: *Typical Power Distribution Network Equivalent Circuit*

The complete PDN design for PCB can be broadly broken down into the following steps:

1. Determine the power requirements of the load. From this information the designer can determine the required target impedance. This load should broken down into:
    a. Voltage and voltage tolerance
    b. Total current
    c. Worst-case dynamic current profile
2. Design the voltage regulator
3. Design PCB PDN to meet the required target impedance
4. Attach the package and die model to the PCB model to extract system-level model
5. Run transient analysis to meet the ripple noise spec defined at die pin level

Each of these steps is covered in further detail below.

Determining Power Requirements

Using the $CORE_{VDD}$ PDN as an example, this PDN must deliver a maximum of 49.3A of static current and up to 113.7A of dynamic current, for a total DC current of 163A, at the chosen voltage of 0.75V, 0.5V or 0.95V. To solve the impedance formula above, the designer needs for each supply:

- The maximum voltage excursion from the nominal (the ΔV_{MAX})

- The maximum instantaneous current (or ΔI_{MAX}).

> **Note**
>
> ⓘ I_{MAX} is not the dynamic current, which is specified at DC, but rather at the switching rate of the transistors.

In the case of the $COREV_{DD}$ supply, ΔV_{MAX} is given as 36 mV, and ΔI_{MAX} is observed, from the worst-case current profile, to be 18A. This results in a target impedance of 0.036/18 = 0.002Ω, or 2 mΩ.

Achronix provides the impedance profile for each rail. As can be seen below, this profile is modified to relax the requirement at higher frequencies.

Designing/Specifying the Voltage Regulator Module

The design of the voltage regulator model (VRM) is critical to ensure a clean power supply to the chip. The regulator must be designed keeping the maximum total (static plus dynamic) current in consideration. Since large amounts of current are typically required, a switching-mode power supply (SMPS) is typically selected, both for its ability to deliver large currents, but also its efficiency, The SMPS VRM is a classic negative-feedback control system. As a result, attention must be paid to ensure stability through compensation networks.

Commonly, SMPS VRM ICs and even fully developed modules are available from reliable manufacturers who can also provide guidance on placement and routing of the VRM. In the best case, the manufacturer can provide a SPICE model of the regulator to plug into the model.

An important feature of the VRM is the sense line, which must be routed from the load (usually the pins of the BGA) back to the controller. In the case of the $COREV_{DD}$ supply, sense pins are provided on the BGA, which allows die-level control of the voltage. The sense pins allow the controller to control the voltage at the point-of-load, which enables the PCB to have some loss between the regulator and the load, generally expressed as "IR drop". However, the designer is cautioned against allowing too much loss. As it generates heat, this causes the regulator to be less efficient, and can even cause the regulator to be less stable or limit its effective operating frequency range. In the case of the $COREV_{DD}$, a reasonable voltage drop might be 200 mV. This correlates to a DC plane resistance 0.200/163 = 1.226 mΩ, and a power loss (P = IV) of 163A × 0.2V = 32.6W, a significant source of heat. Clearly, the lower the IR drop that can be designed into the PCB the better, but that takes a considerable amount of copper. This trade-off is a design decision that must be weighed carefully.

A designer is well advised to work with the manufacturer of their VRM components to select appropriate parts, and to lay those components out carefully according to the manufacturer's recommendations. It is critical when laying out a VRM that conductive loops carrying high transient current (the output return loop to ground, the output return loop to the supply, and the input current loop) be kept as small and short as possible.

Designing the PCB PDN

The PCB PDN should be designed in a way to ensure that the bulk capacitors, decoupling capacitors and routing from the regulator to the AC7t1500 device are sufficient to meet the impedance profile for the load.

Bulk Capacitance Selection

The bulk capacitance should be enough to ensure that the frequency at which the regulator output impedance crosses the target impedance, the bulk capacitors present provide a lesser impedance than the target impedance.

$$C_{BULK} > L_{VRM}/(Z_{Target})$$
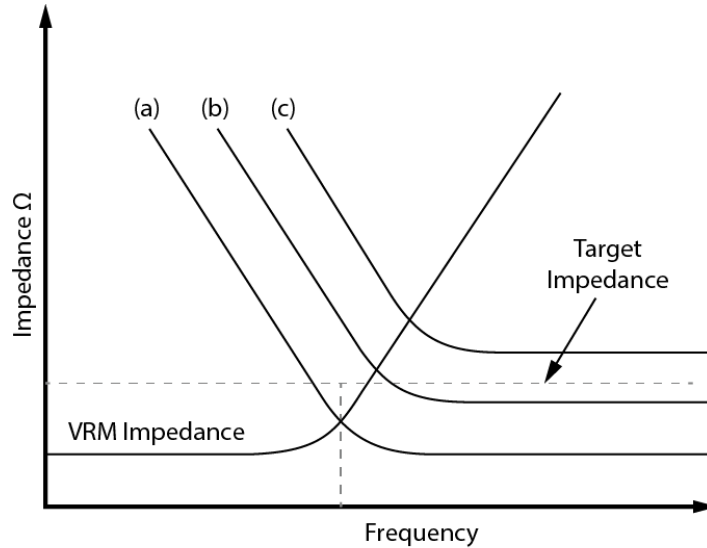
70536065-01.2021.09.26

Figure 45: *Selecting Bulk Capacitance Values*

The figure above shows the equivalent output impedance for the voltage regulator model. The VRM impedance crosses the target impedance at frequency F1.

- Curve **a**, **b** and **c** show the impedance plot for different bulk capacitors.
- Case **c** has insufficient bulk capacitance as the curve intersects the VRM impedance above the target impedance line.
- Case **b** has the minimum bulk capacitance needed to keep the impedance below target impedance after frequency F1.
- Case **c** has more bulk capacitance than the minimum bulk capacitance requirement.

PCB Decoupling Capacitor Selection

PCB decoupling capacitors are to be selected based on the target impedance. Target impedance value can be calculated as explained in Determining Power Requirements (see page 67). Alternatively, Achronix provides impedance targets for each power supply. Refer to the *Speedster7t Power User Guide* (UG087) for the impedance targets for the supplies.

Some important considerations regarding PCB decoupling capacitors:

- The impedance targets are defined for the worst-case operations and might need scaling based on the use case scenario.

- The decoupling capacitors should be placed as close as possible to the FPGA. Refer to PDN Layout Guidelines (see page 71) for more details.

- Ultra-low-ESR capacitors are not always optimal. The equivalent series resistance (ESR) of capacitors should bring the overall impedance close to the target impedance. For example, the number of caps (*n*) of the same value C that are required to keep the impedance value below target impedance is:
  $Z_{Target} \geq ESR/n$
  Where: ESR is the the intrinsic + spreading resistance of the decoupling capacitor, C.

This effect can be seen in the figure (see page 68) above, where each capacitor of larger value also has a lower ESR.

AC Impedance Analysis

The PCB AC impedance must be simulated in an AC simulator to verify that the target impedance is met at the frequency of interest. The optimization of capacitors might need multiple iterations as each capacitor addition can impact the impedance well beyond the frequency of interest.

Some EDA tools provide assisted capacitor selection to help optimize the PDN based on target impedance and other constraints. However, the reliability of these tools to provide the correct optimization may not always be trusted. Hence, it is recommended to analyze the selection of capacitors from these capacitor optimization tools and improvise the selection if needed.

The target impedance may be modified somewhat. For instance, the target impedance of the COREV$_{DD}$ supply is specified by Achronix documentation to be:

- ≤2 mΩ for the range of 50 KHz to 30 Mhz

- ≤7.6 mΩ at 100 MHz

…and can linearly increase between 30 MHz to 100 MHz (see the following graph).
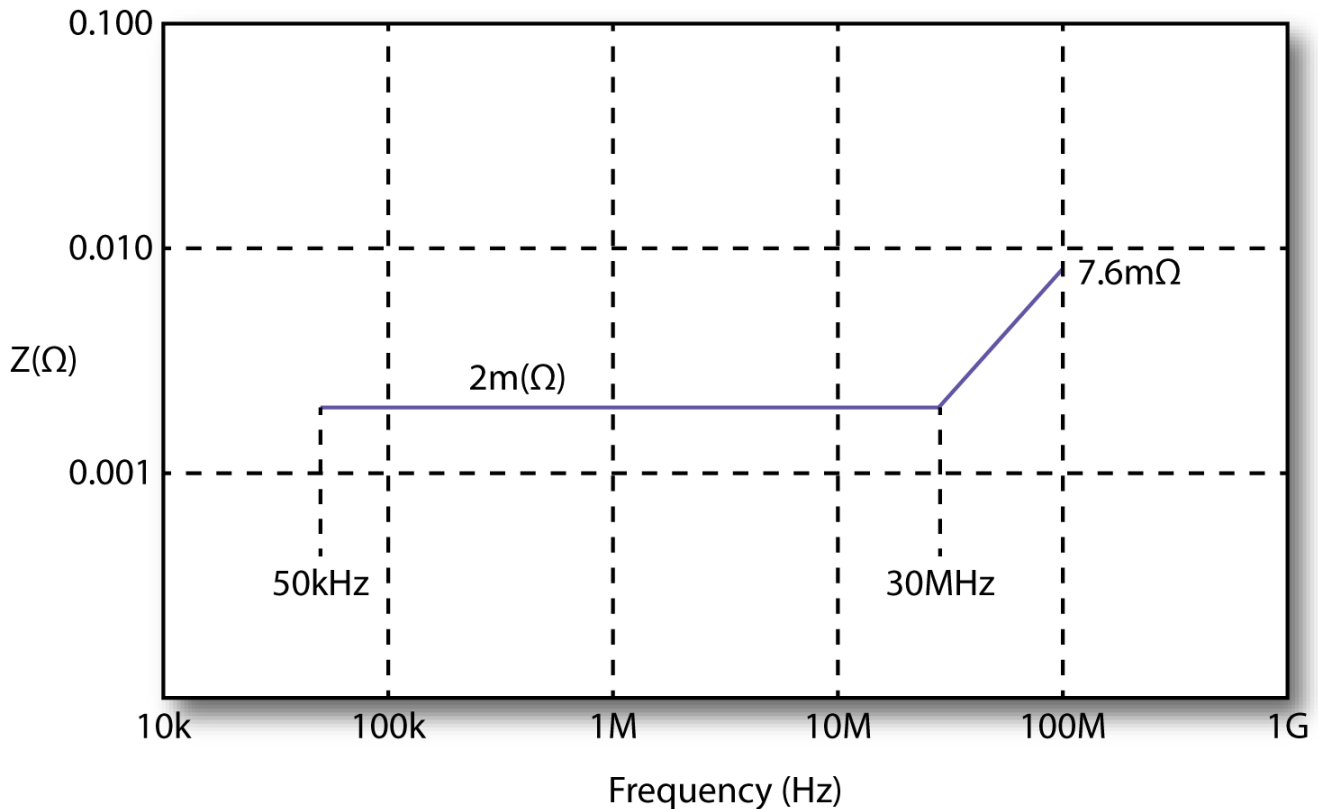
70536065-04.2020.12.27

Figure 46: *Impedance Envelope of COREV$_{DD}$ Supply*

The impedance of the COREV$_{DD}$ PDN at the device bumps must be lower than the line shown in the graph above. For further details on specific power rails, contact Achronix support.
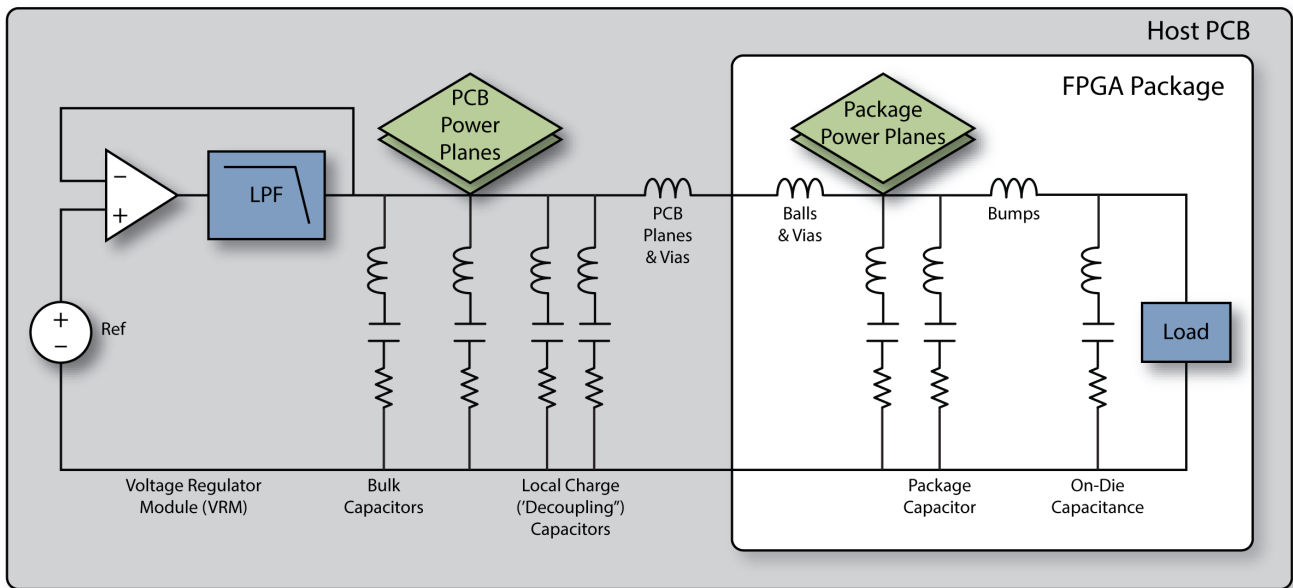
System-Level Modelling for the PDN

System-level modeling requires stitching together the equivalent model of each component to mimic the system-level electrical circuit for the PDN. The following models are required for creating this system-level PDN model:

- **Voltage regulator** – Modelling a voltage regulator for PDN design can be a complex task. In the most simplified form, a voltage regulator can be represented with a lumped series R-L circuit. However, this model does not capture the feedback provided to the regulator from the device. This model may still be sufficient for PDN analysis, but for modelling a voltage regulator, a more detailed analysis might be required. Contact the voltage regulator vendor for help with modeling.

- **PCB model** – PCB layout models can be extracted in any standard EDA 2.5D EM tool. The PCB model must capture the complete path from the VRM to the AC7t1500 device, and it must have models for all components in the PCB.

- **Package and die model** – Contact Achronix Support for package and die models.

Modeling PDN System-Level Transients

To run system-level transient simulations, the system-Level PDN model is required. Also, current profiles for the PDN are required.

70536065-05.2020.12.27

Figure 47: *System-Level Transient Model*


ⓘ **Note**

In the figure above, the load needs to be replaced by the on-die current profile.

PDN Layout Guidelines

The following is a list of layout guidelines for the PDN:

- Do not split the ground plane into separate planes for analog, digital and power pins. A single contiguous ground plane is recommended.

- Stagger the vias placement to avoid creating long gap (plane chokes) in the plane due to via voids.

- Place and route the components within their own power plane.

- Power planes can handle more current than traces and the broad planes helps in lowering the operating temperature of the board.

- Power planes coupled to ground planes immediately above or below reduce the capacitor mounting inductance and thus help to minimize dynamic noise. The thinner the dielectric between power and ground planes, the better.

- Add more power stitching vias for better connectivity and to reduce loop inductance.

- Power vias should have ground vias in close proximity to reduce the current loop as much as possible.

Decoupling capacitors should be placed as close to the AC7t1500 device as possible. Typically, the top layer does not have space to accommodate caps for all device power supplies. In that case, adding decoupling caps on the top layer can be difficult. For placing decoupling capacitors in such cases, the area underneath the AC7t1500 device package on the bottom layer can be used. The 0402 size of caps fit well on the back side of a PCB, providing a vertical connection to the AC7t1500 BGA pins. In the figure below, the circular pads are BGA pads on the top layer, and the square pads are decoupling capacitor pads on bottom layer.
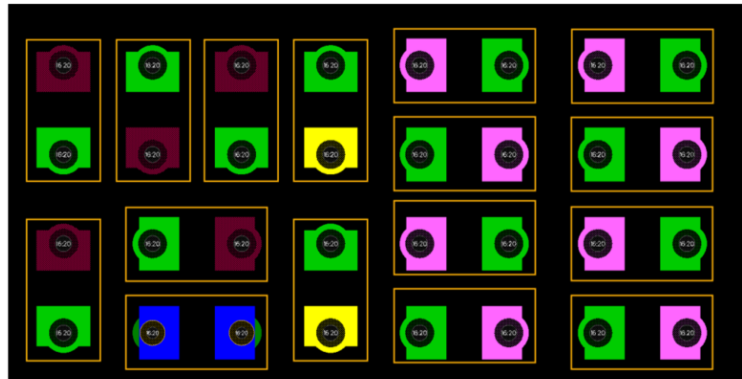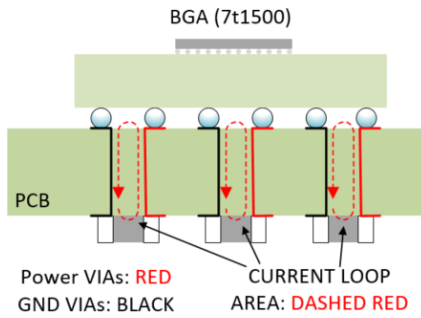
Figure 48: *Placement of Decoupling Capacitors Using Via-in-Pad*

- The connection from voltage regulator to the AC7t1500 FPGA should be as broad and unrestricted as possible. Thicker copper layers are better for the PDN as they have lower IR drop and loop inductance. The figure below shows the wrong and right ways to route the broad plane (highlighted in red) that connects the VRM and the FPGA.
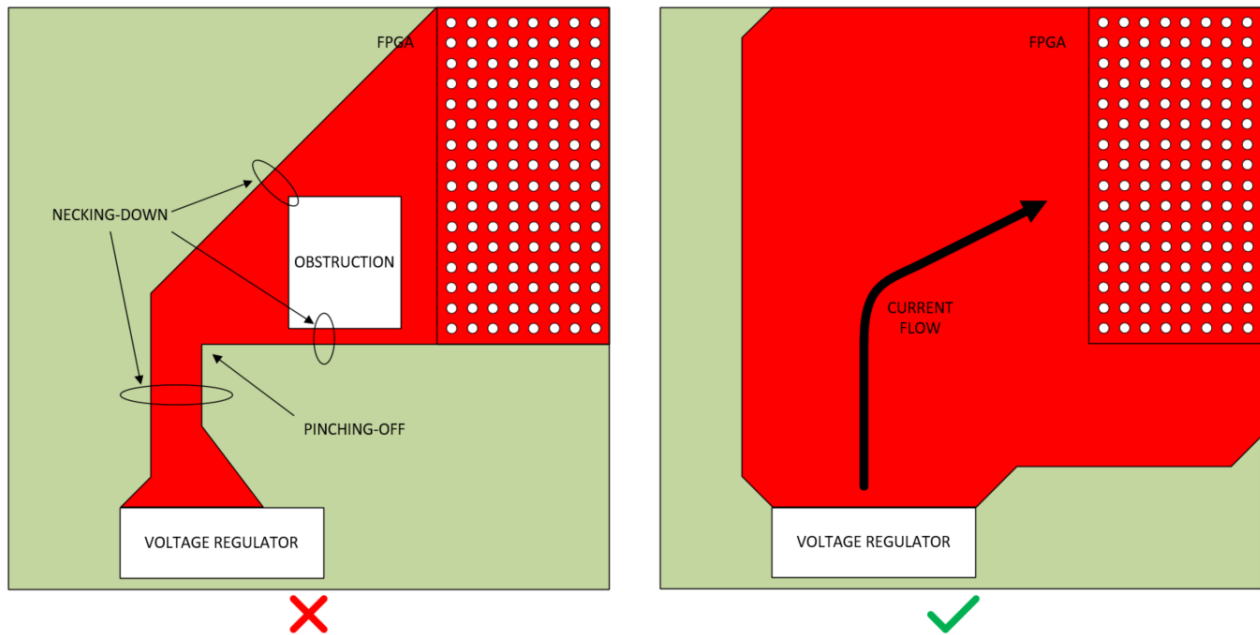
Figure 49: *Routing Power Planes from the VRM to the BGA*

- The pins, `Core_VDD_Sense` and `Core_VSS_sense`, can be used to sense the voltage of $COREV_{DD}$ at the die bump. These pins must be routed together similar to a differential pair, but there is no impedance requirement for this pair.

Chapter - 9: PLL Power Filtering

The Speedster7t AC7t1500 FPGA has very accurate on-board clock synthesizers fed by a PLL. Deterministic Jitter might occur at the output of the PLL due to power supply noise. This jitter depends on many factors, including PVT conditions, the operating point of the PLL, and the magnitude and frequency of the noise. To achieve a reasonable level of long-term jitter, it is vital to provide an analog-grade power supply, with as little noise as possible. There are two primary sources of noise in an electronic system: the voltage regulator module (VRM) and the switching of the digital circuitry.

These PLL inputs are designated as PLL_VDDA and have an accompanying PLL_VSSA, as outlined in the table below:

Table 15: *AC7t1500 FPGA PLL Power Pins*

Supply	Pin	VSS	Pin
ENOC_NW_PLL_VDDA	Y18	ENOC_NW_PLL_VSSA	W18
ENOC_N_PLL_VDDA	V17	ENOC_N_PLL_VSSA	V16
ENOC_NE_PLL_VDDA	V35	ENOC_NE_PLL_VSSA	V36
ENOC_SW_PLL_VDDA	BA19	ENOC_SW_PLL_VSSA	BA18
ENOC_S_PLL_VDDA	BA22	ENOC_S_PLL_VSSA	BA21
ENOC_SE_PLL_VDDA	AB36	ENOC_SE_PLL_VSSA	AB35
GCG_NW_PLL_VDDA	W19	GCG_NW_PLL_VSSA	V19
GCG_SW_PLL_VDDA	BA17	GCG_SW_PLL_VSSA	BA16
GCG_NE_PLL_VDDA	W36	GCG_NE_PLL_VSSA	W35
GCG_SE_PLL_VDDA	AD36	GCG_SE_PLL_VSSA	AC36

Noise Control Methods for Analog Supplies

Linear Regulator

Noise from the VRM can be eliminated most directly by using a linear rather than a switching supply (with the attendant switching ripple). The input to the linear regulator may be a switching supply; therefore, it is also critical to select a regulator with a high power supply rejection ratio (PSRR). It is also important to lay the regulator out according to the manufacturer's directions, selecting the correct capacitors to ensure stability and placing them as close as possible to the related pins.

PLL Supply Filtering

Another method for delivering clean power to an analog circuit is through filtering. An L-C filter, composed of a ferrite inductor followed by a capacitor, is typically used. The specific ferrite required is highly dependent on the operating environment and what frequency content is expected at the PLL_VDD input. If a switching supply is used, then the ripple frequency of the regulator guides selection of the cutoff frequency.

A filter circuit and a linear supply can be used together to provide a very effective low-noise supply.

Analog Supply Decoupling

The PLL circuit does cause local switching noise, for which a decoupling capacitor is required across the PLL_VDDA and the PLL_VSSA pins as shown below. This decoupling function requires the highest-value, high-frequency capacitor available in a small package. This capacitor often turns out to be 100 nF in an 0402-sized package.
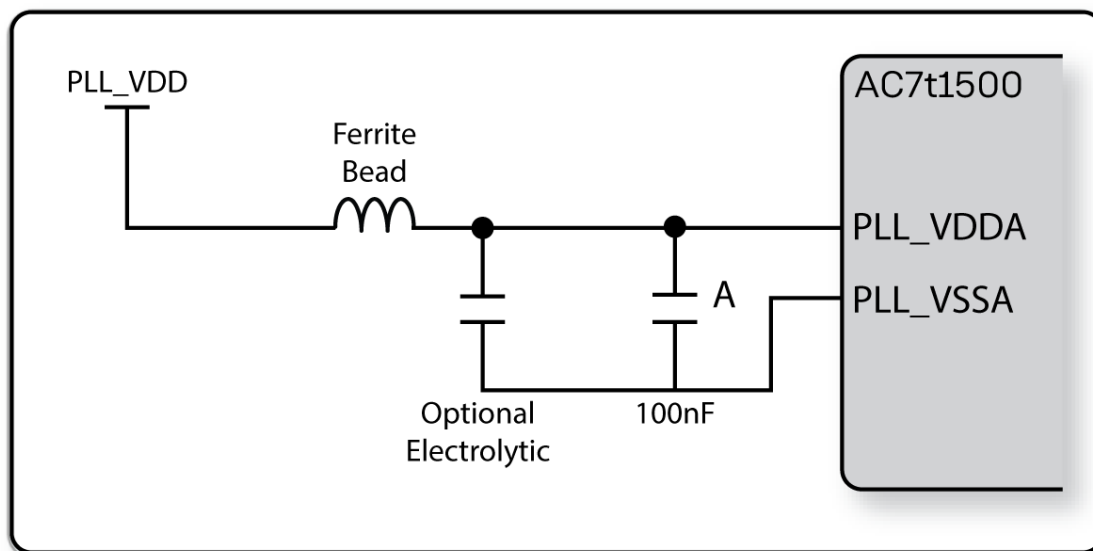
70542647-01.2021.11.08

Figure 50: *PLL Power Filter*

> **Note**
>
> The VSS side of the 100 nF capacitor (noted as "A") is *not* connected to the PCB ground, but is instead only connected to the PL_VSSA pin. The VSSA pin is connected to system ground in the silicon to provide the lowest possible inductance between the two nets. Adding a connection to PCB ground outside the AC7t1500 FPGA package would create a ground loop that would be susceptible to picking up EM radiation.
>
> It is vital that capacitor "A" be placed as close as possible to the VDDA and VSSA pins in order to minimize the loop area of the capacitor and connection.

Low-Frequency Applications

In applications with a low PLL reference frequency and an environment with significant low-frequency components, it is often beneficial to add a large-value capacitor which fits nicely on the board (often 22 µF).

Revision History

Version	Date	Description
1.0	09 Nov 2021	• Initial Achronix Release.