Speedster7t Board Design Guide (UG101)

Speedster FPGAs

Achronix
Data Acceleration

Copyrights, Trademarks and Disclaimers

Notice of Disclaimer

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

Speedster®7t FPGAs include 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 to help designers' get the most out of the Speedster7t FPGAs.

This guide is broken down by system components which include the Ethernet, PCIe, GDDR6, DDR4/5 memory interfaces. In addition, this document also details board design requirements for GPIOs, CLKIOs and power regulation and distribution.

- Ethernet (28G, 56G, 112 Gbps)
- PCIe5 (32 Gbps)
- GDDR6 (16 Gbps)
- DDR4 (3200 Mbps)
- DDR5 (5600 Mbps)
- GPIO and CLKIO
- Power and PDN

For more details on the Speedster7t FPGA architecture, see the ***Speedster7t FPGA Datasheet*** **(DS015)**[1].

The diagrams below illustrates the separate interfaces discussed in this guide

1 https://www.achronix.com/documentation/speedster7t-fpga-datasheet-ds015

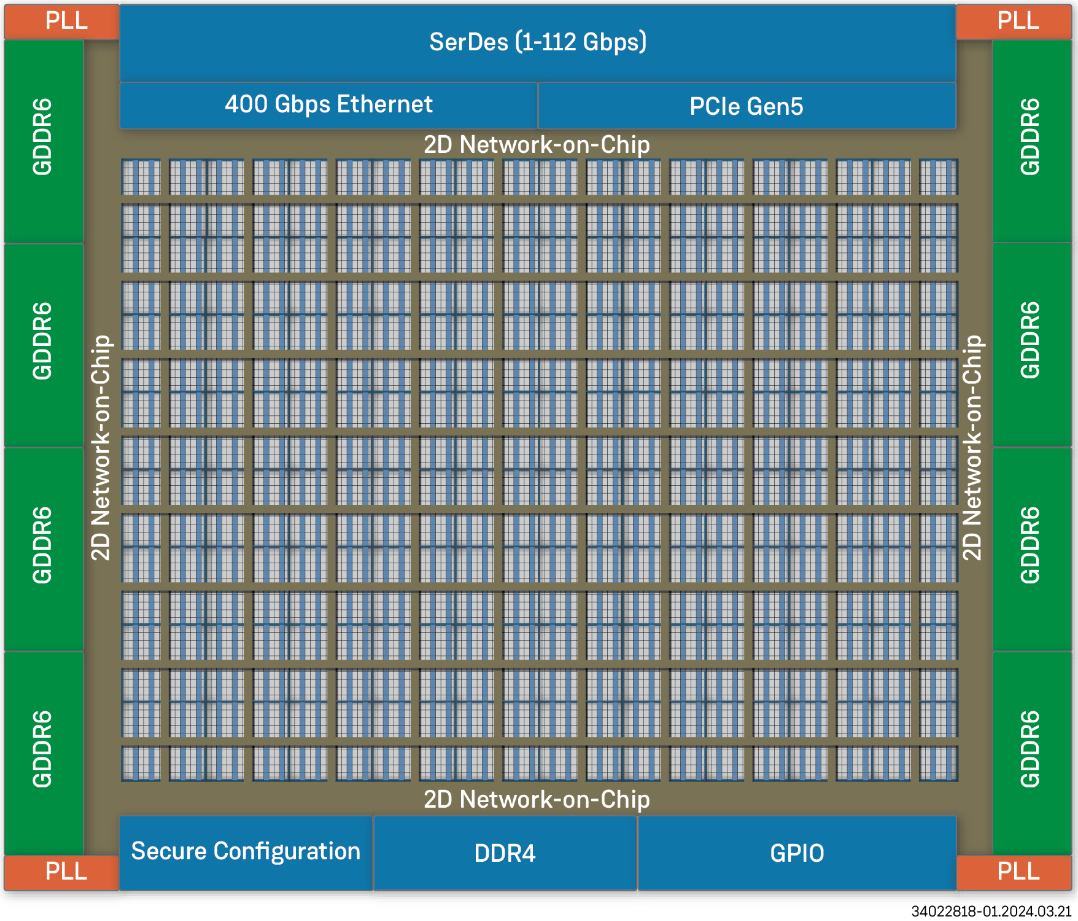

Figure 1 • Speedster 7t11400/1500 FPGA Block Diagram

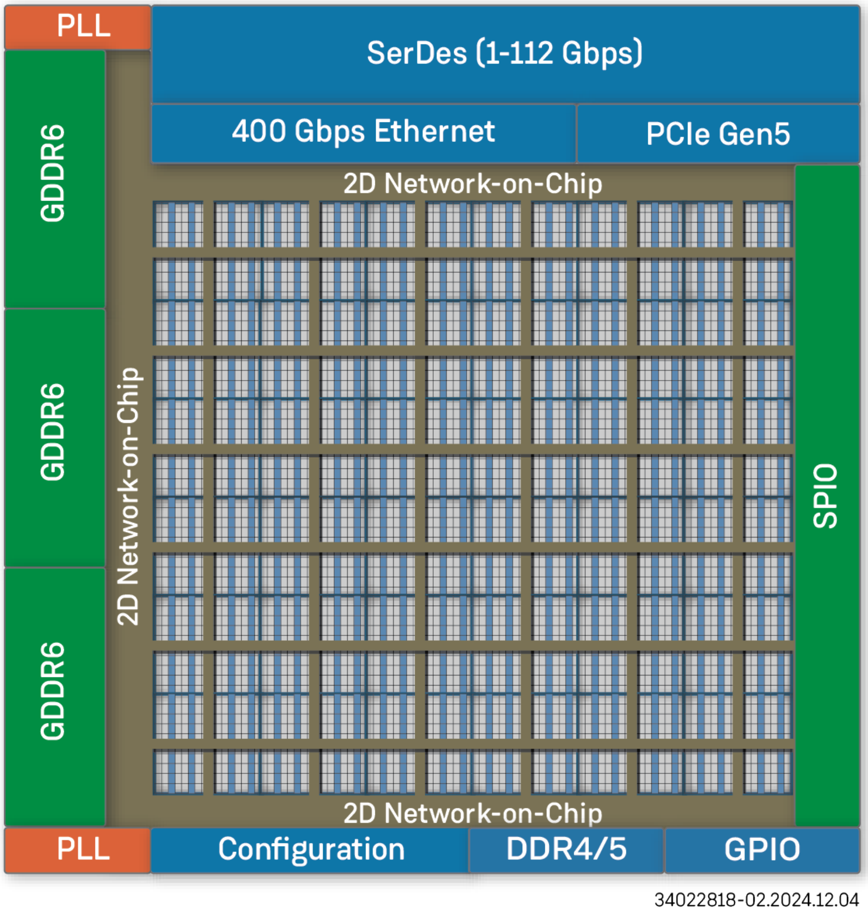

Figure 2 • Speedster 7t700/800 FPGA Block Diagram

Chapter 2 : PCB General Considerations

Extensive efforts have been made across the industry to enable successful high-speed circuit implementations while balancing performance with material cost considerations. This chapter does not attempt to provide an exhaustive compilation of all existing knowledge. Instead, it highlights key topics and design principles that are widely recognized within the industry as best practices for high-speed PCB design.

Board Construction - PCB Stack-up Considerations

When designing and evaluating stack-up configurations for a specified layer count, it is essential to engage with a qualified PCB fabricator experienced in low-loss materials. Such fabricators can offer a range of suitable material options tailored to high-speed signal integrity requirements. Based on the product's electrical and mechanical constraints, advanced fabrication techniques — such as via-in-pad, via back-drilling, and enhanced impedance control — may be necessary to meet performance targets.

A pre-layout signal integrity analysis is strongly recommended to determine the optimal stack-up strategy. This analysis informs the selection of materials and fabrication techniques, ensuring alignment with design requirements. At multi-gigabit speeds, precise layer-to-layer impedance control becomes critical to maintaining signal fidelity, as it minimizes reflections caused by discontinuities in the transmission path.

The table below illustrates a segment of a representative stack-up.

Note

The table below represents a partial example provided for reference only; actual stack-ups vary based on design and application-specific needs. Typically, stack-ups are constructed symmetrically to maintain mechanical stability. In this example, two layers are designated for high-speed signal routing.

For SerDes channels and other high-speed interfaces, it is recommended to use dielectric materials with a low dissipation factor (also known as the loss tangent) to reduce insertion loss and maintain signal quality over longer traces.

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.

- IR drop and ripple noise of the FPGA and other components on the PCB
- Mechanical requirements, especially thickness, but also board warpage, mounting and thermal requirement

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

Distance between reference layers and signal layer may also impact crosstalk. Crosstalk minimization is accomplished by ensuring the impedance is defined mainly by the distance to the reference plane. 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 pre-layout analysis.

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) 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. For high-speed SerDes, in most of the cases, stripline routing is better.

PCB Via Design

For PCB BGA pad design, consider via-in-pad (VIP) and multi-layer stacking of vias to minimize 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.
- Back-drilling plated through-hole (PTH) vias to minimize the stub. Be careful to obtain maximum stub length information from the fab vendor.

Via 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 in an inner layer of the PCB to an outer layer. Since the via goes through the PCB, one end of the via sees an open connection.

In the figure below, the signal via, once it reaches the inner routing layer, is split into two components, A and B. A travels to the unterminated end of the via and reflects back. Depending on the frequency and the length of the stub, the reflected signal A can degrade or improve the signal B. This structure creates a resonance observed in return loss of the signal trace. Maxima of this resonance (i.e., maximum degradation) is seen at the frequency at which stub length is equal to one quarter of the wavelength of the signal.

Via stubs can be mitigated in one of two ways:

- **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.
- **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.

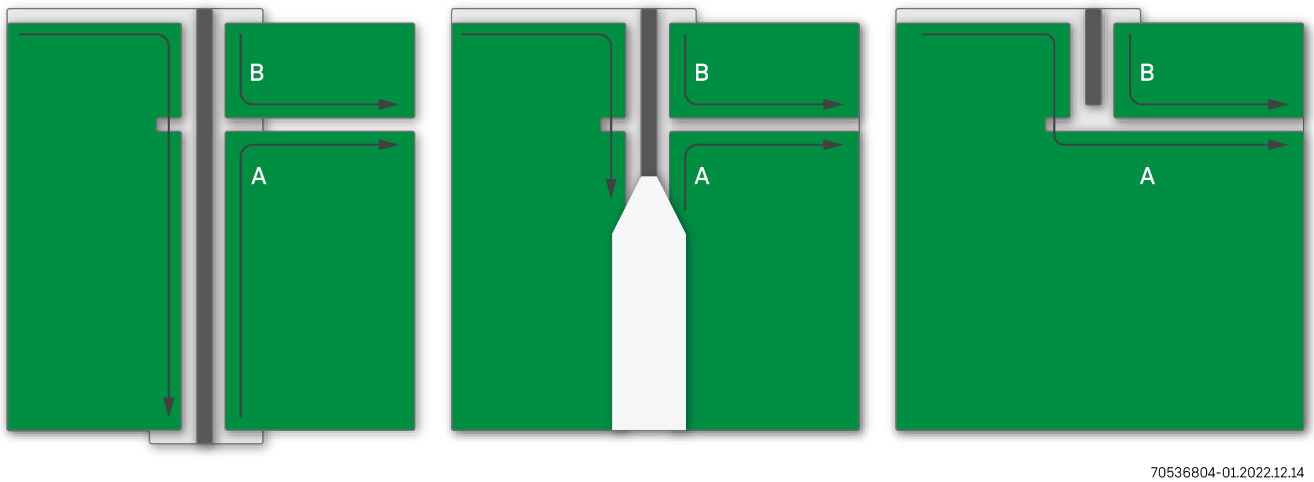

Figure 3 • Via Stub Resonance and Mitigation by Back-drilling and Laser Vias

Power Distribution

The power distribution network (PDN) for Speedster7t FPGAs benefits greatly from power-ground plane pairs integrated into the stack-up. These power-ground 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. While PDN design is covered in the chapter, PCB Power and PDN Design (page 69), it bears repeating that the stack-up must take into account the projected current requirements and provide enough copper (thicker copper layers are better for power than thin copper) to deliver that current with minimal voltage drop. These requirements have an impact on:

- Power efficiency
- Voltage regulator stability
- Copper heating effects

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-100G and Megtron-7 are some of the popular choices for low-loss materials for routing high-speed SerDes signals. Contact your PCB manufacturer to find out the best solution based on project requirements.

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. 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.

Following are some possible solutions to the fiber weave problem:

- Rotate either the design or the PCB panel by a small angle (e.g 10 degrees).
- 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.

Component Placement

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 connector. Four quads, N0-N3, are connected to a ×16 PCI Express slot and the last four quads (N6 and N7) are routed to a secondary PCIE slot. GDDR6 DRAMs are placed on the east and west side to align with the physical location of GDDR6 PHY on the FPGA. The two DDR4 DIMMs are on the south side. This arrangement takes advantage of the natural signal flow from the FPGA.

> **Note**
>
> The block diagram below does *not* show all FPGA blocks and the corresponding connections.

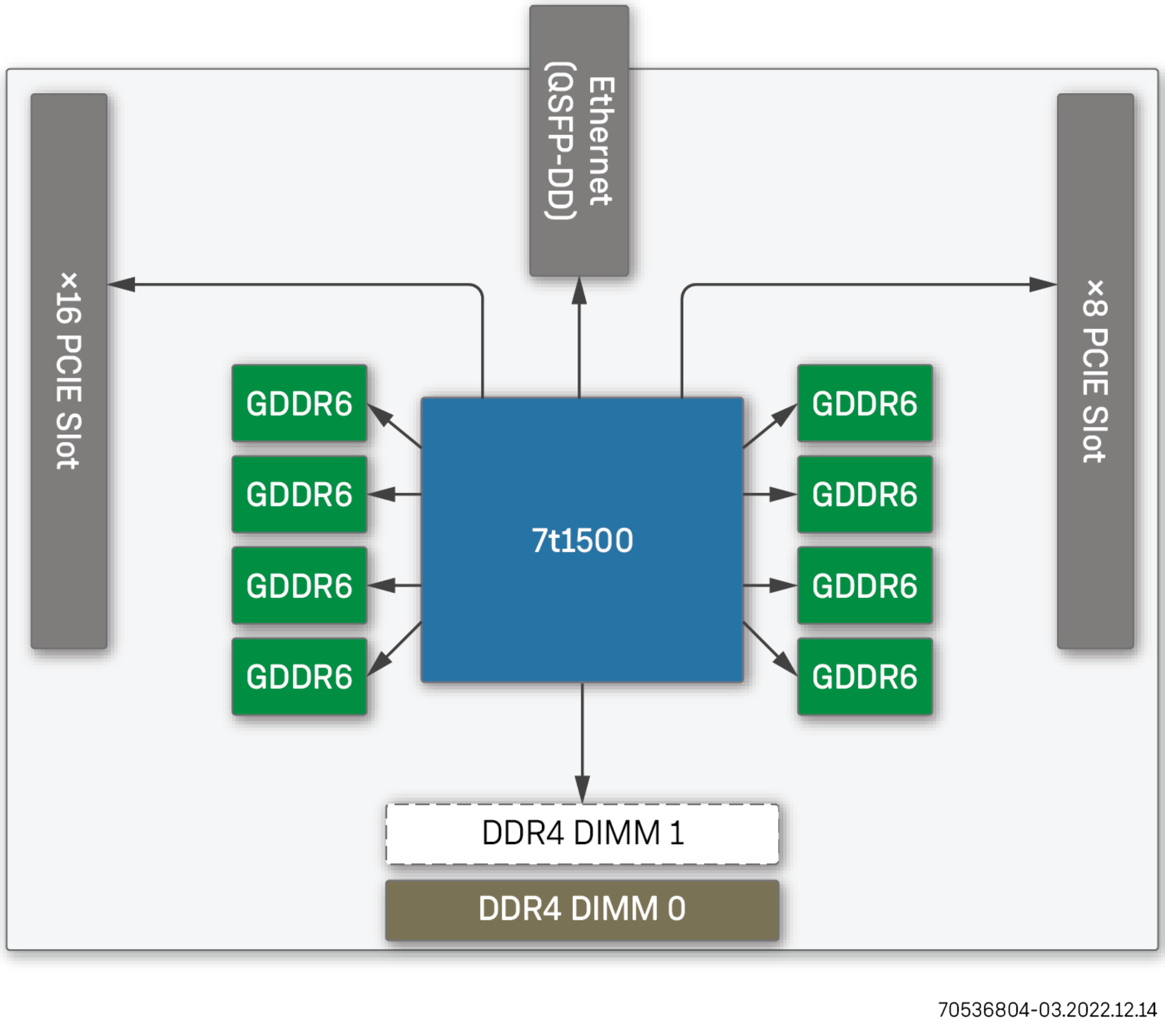

Figure 4 • Component Placement with One QSFP-DD Ethernet Connector, Two PCIE Connectors, and Two DDR4 DIMMs

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.

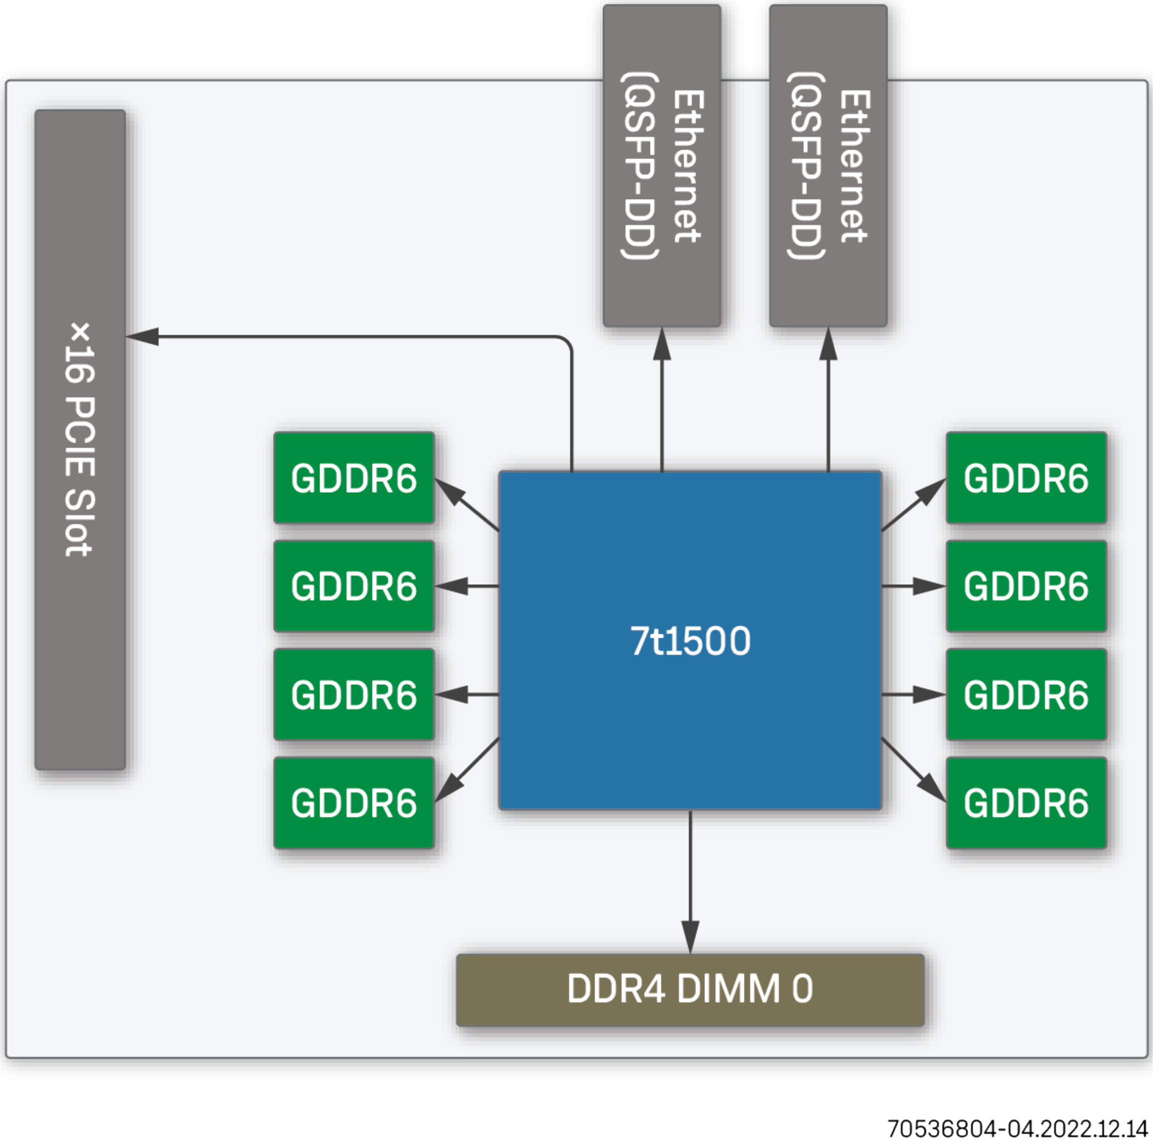

Figure 5 • Component Placement with Two QSFP-DD Ethernet Connectors, One PCIE Connectors, and One DDR4 SO-DIMM

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.

Library Resource

Contact Achronix for symbol, footprint and 3D models. Also, refer to the IPC-7351 standard for footprints.

Chapter 3 : High-Speed SerDes Interfaces

The Speedster7t SerDes Interface is compliant with many serial interfaces, with data rates up to 56 Gbps NRZ and 112 Gbps PAM4. While the SerDes PHY 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 Speedster7t 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.

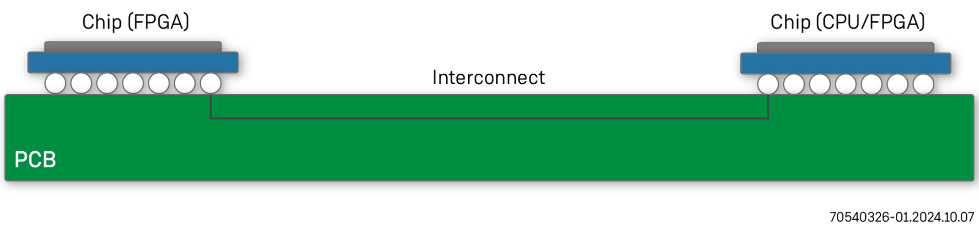

Figure 6 • Chip-to-Chip Topology

Chip-to-Module Topology

The chip-to-module topology supports plugable optical or copper cable modules for high-bandwidth switch applications. The interconnect between the Speedster 7t FPGA and module is a relatively short-reach channel over a PCB consisting of traces, vias and a connector.

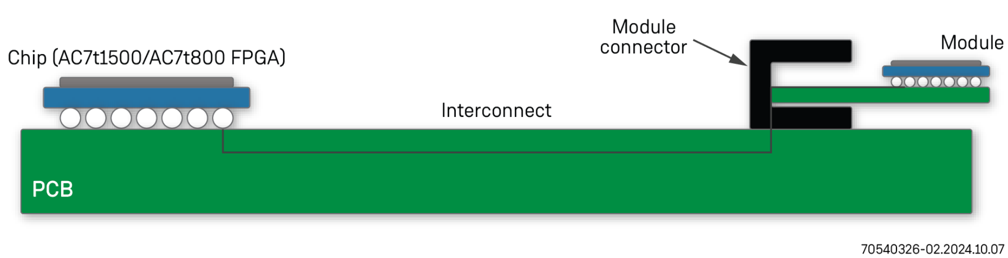

Figure 7 • Chip -to-Module Topology

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 Speedster7t 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.

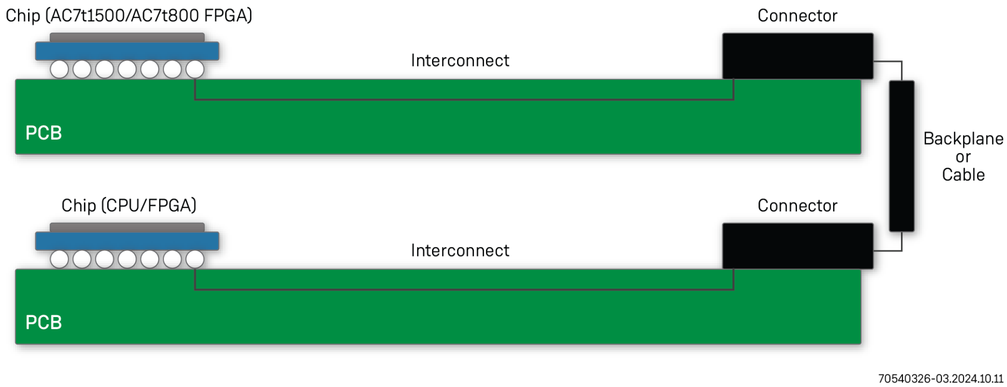

Figure 8 • Chip-to-chip Over Connector or Cable Topology

Signal Integrity Specification

Understanding the different high-speed SerDes standards (PCIe, Ethernet etc.) 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)[2] for detailed list of SerDes standards supported by Speedster7t FPGAs.

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:

2 https://www.achronix.com/documentation/speedster7t-serdes-user-guide-ug099

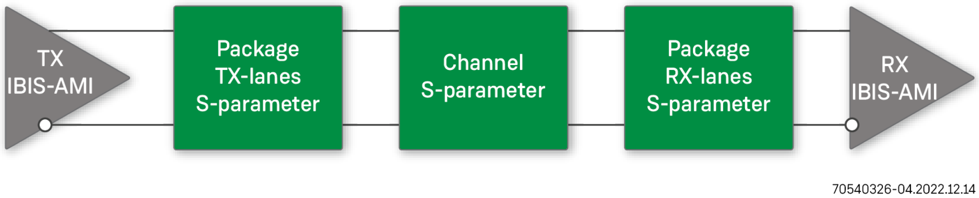

Figure 9 • Sign-off Simulation Topology

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

- **Transmitter TX IBIS-AMI behavioral models for the endpoint** – For Achronix FPGA, these models are available from Achronix upon request (support account required).
- **Package TX-lanes s-parameter models for the endpoint** – For Achronix 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** – For Achronix 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** – For Achronix 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 the Speedster7t FPGAs 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 Speedster7t FPGAs 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 (page 8)" in the chapter, "PCB General Considerations (page 4)".

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.

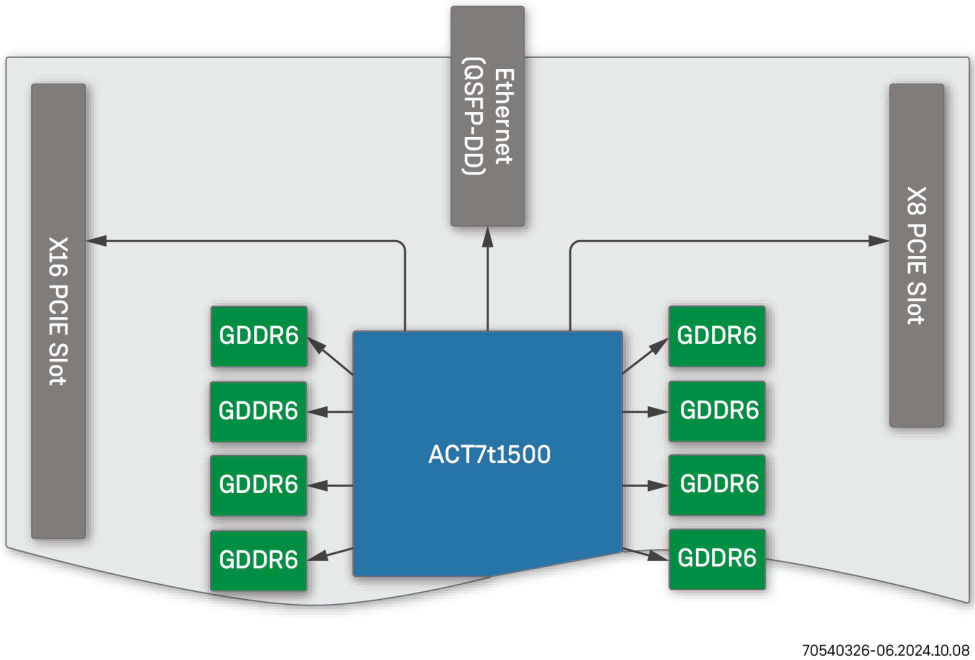

Figure 10 • Example Placement of High-Speed SerDes Connectors (PCIe and Ethernet) for 7t1400/1500

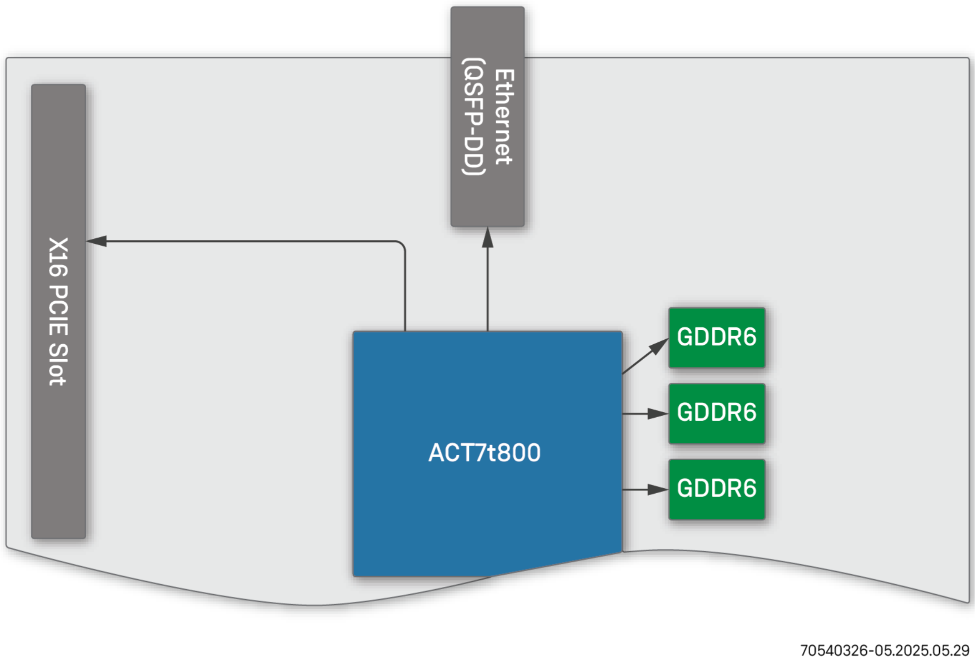

Figure 11 • Example Placement of High-Speed SerDes Connectors (PCIe and Ethernet) for 7t700/800

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 signal loss due to reflections.

> **Note**
>
> The SerDes interface on the 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 (page 10)" in the chapter, "PCB General Considerations (page 4)".

Specific to the SerDes interfaces:

- Consider routing the transmitter and receiver signals on separate routing layers. This separation makes signal breakout and routing easier and reduces near-end crosstalk (NEXT).
  - Transmitter signals are the outermost pins, so its easier to route them on an upper layer.
  - Receiver signals can be brought out of the BGA via array on the next available lower routing layer. Evaluate for crosstalk of signals routed in parallel for a long distance.

- The Speedster7t FPGAs 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, "**PCB Pin Mapping on the SerDes Interface** (page 22)" 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. Ground vias accompanying the signal vias helps to control via inductance.
- Stacking of vias on multiple layers help to avoid dog-bone structures and hence, helps to minimize impedance discontinuity.
- 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 half of the wavelength of the signal's Nyquist frequency 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. However, back-drilling of vias is not foolproof (details covered in the chapter, "**PCB General Considerations** (page 4)").

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 via-in-pad allows for much simpler routing by soldering directly over the via. The antipads 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 accompanying the signal as it transits from 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 adjacent signal differential 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.

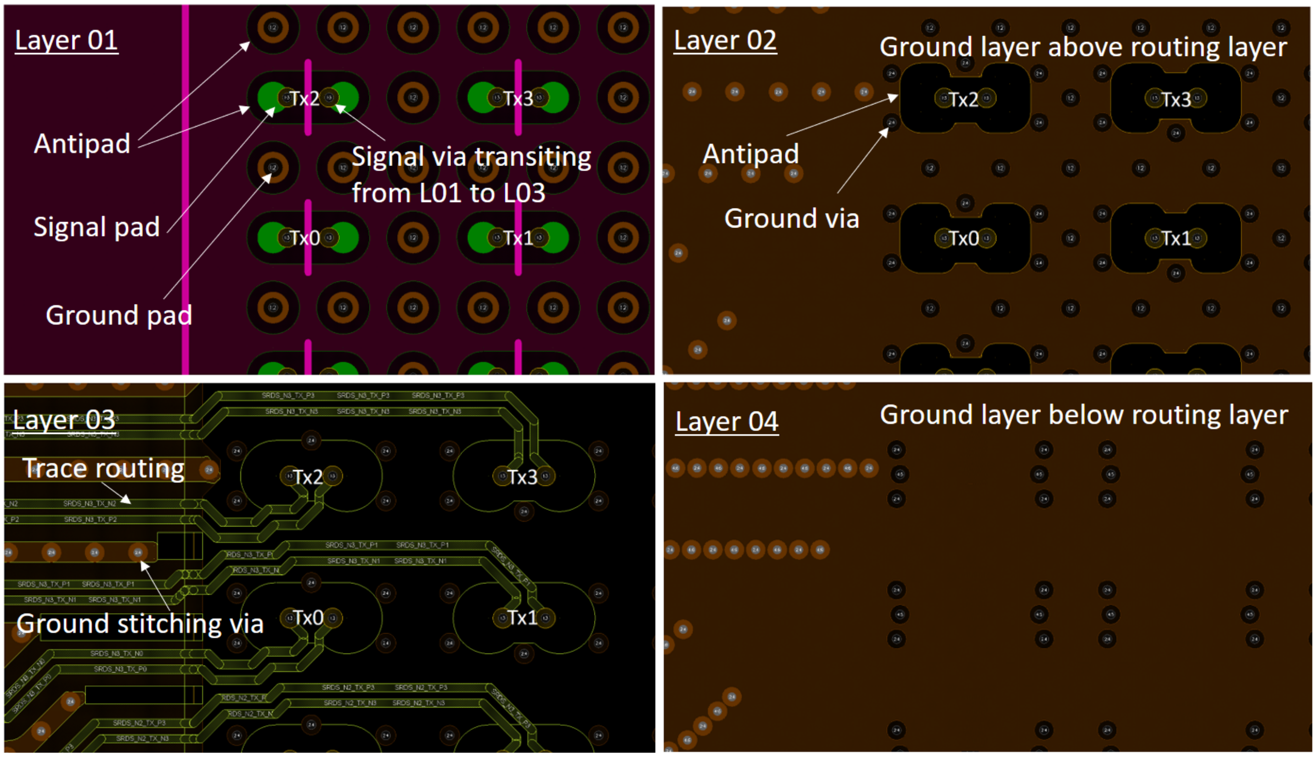

Figure 12 • FPGA BGA Pad Breakout to Layer 3

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 required number of layers, under the capacitors, to be voided. 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 to layer 01.

- 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 03 to layer 02. 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.

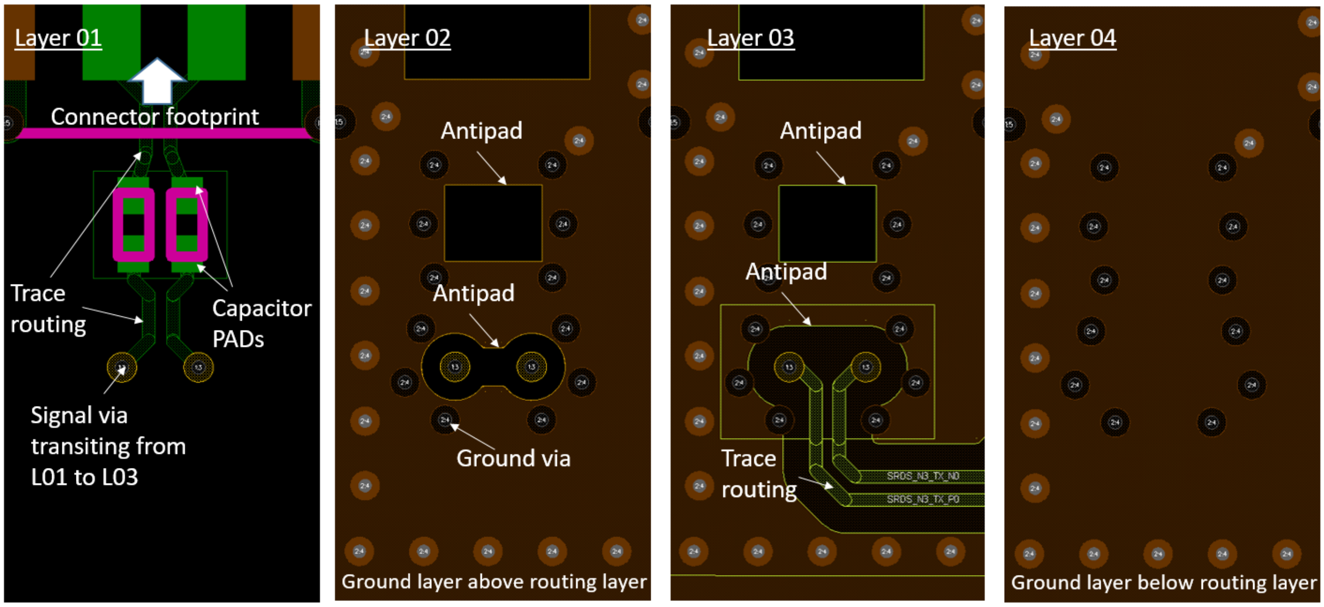

Figure 13 • Example Structure of DC Blocking Capacitor at a PCIe Connector

Connector Footprint

For connector footprint optimizations, following are the two acceptable ways of optimizations:

- 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. Simulations should include the connector structure as well as the PCB to capture the connector to PCB transition. For this kind of simulations, a step model for the connector is required. Simulations using a full-wave solver are to be used in evaluating how many layers underneath the connector footprint are to be cut out.

In case the above options are not available, some common practices can be used but do not guarantee optimal performance. 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. 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 through these vias. 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.

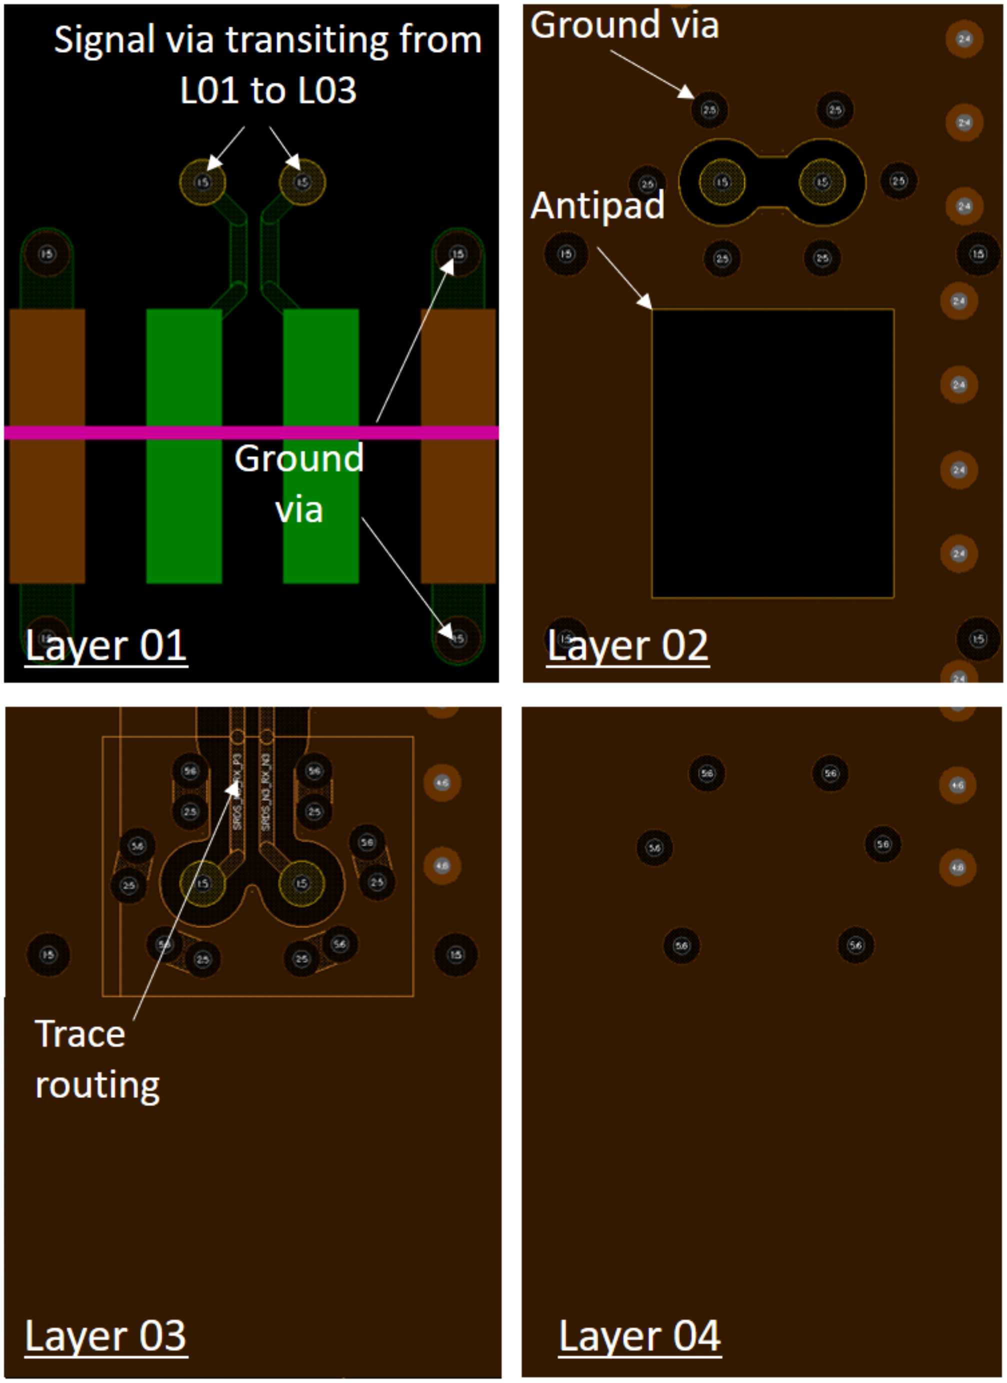

Figure 14 • Example Structure for Transition Via to Connector Pins from Layer 03

Chapter 4 : SerDes Pin Mapping for Ethernet Connectivity

The PCS interface within the Speedster7t devices 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)**[3] for details on lane mapping).

SerDes Connectivity – Linear Mapping for 7t1400/1500

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.

3 https://www.achronix.com/documentation/speedster7t-ethernet-user-guide-ug097

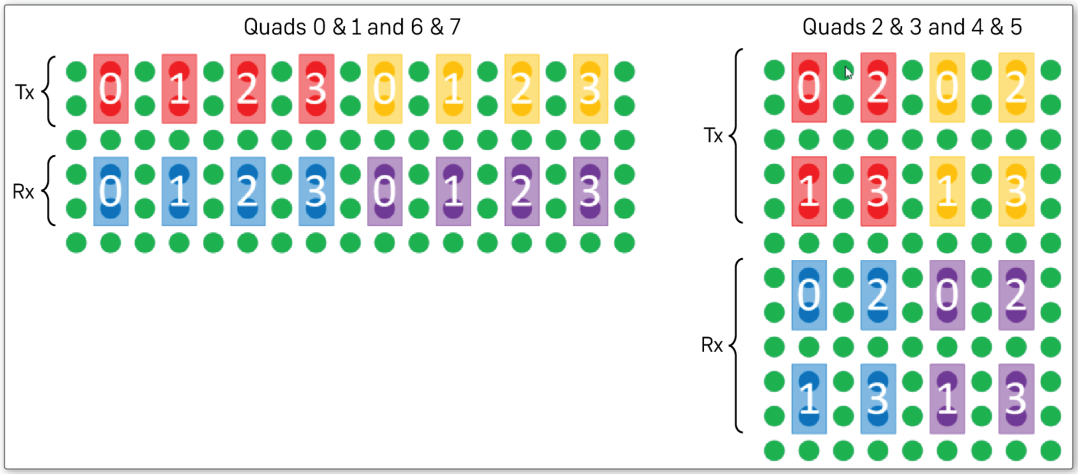

Figure 15 • Linear Pin Mapping (7t1400/1500)

SerDes Connectivity – Linear Mapping for 7t700/800

Linear mapping is generally most useful for PCIE interfaces on Quads 0, 1, 2 & 3. 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.

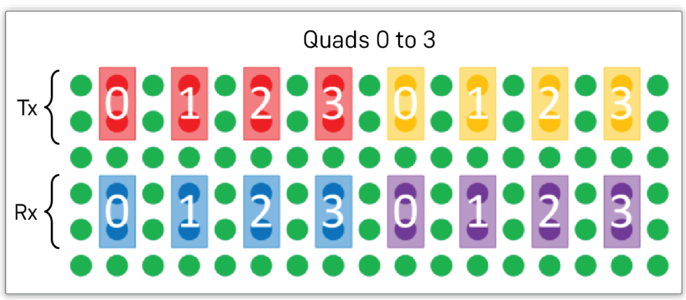

Figure 16 • Linear Pin Mapping (7t700/800)

SerDes Connectivity – QSFP-28 Mapping for 7t1400/1500

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,

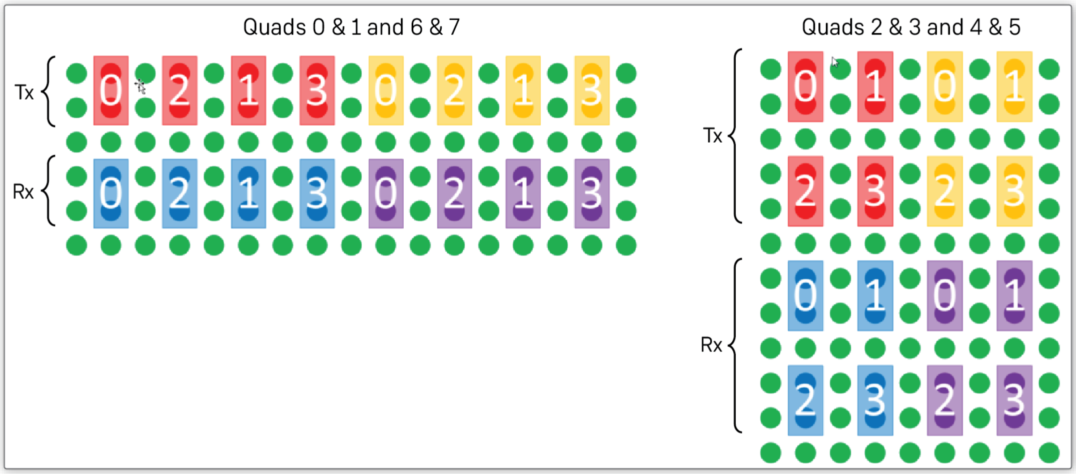

70540707-03.2025.05.29

Figure 17 • QSFP-28 Mapping (7t1400/1500)

SerDes Connectivity – QSFP-28 Mapping for 7t700/800

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,

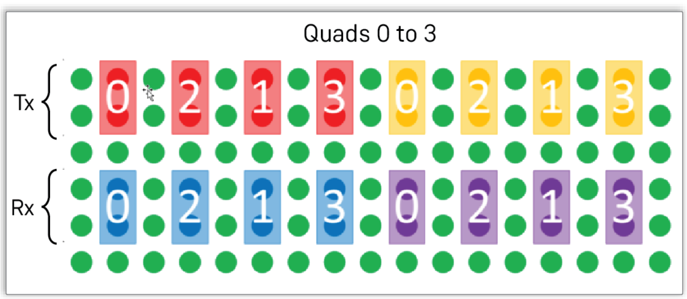

70540707-04.2025.05.29

Figure 18 • QSFP-28 Mapping (7t700/800)

SerDes Connectivity – QSFP-DD Mapping for 7t1400/1500

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.

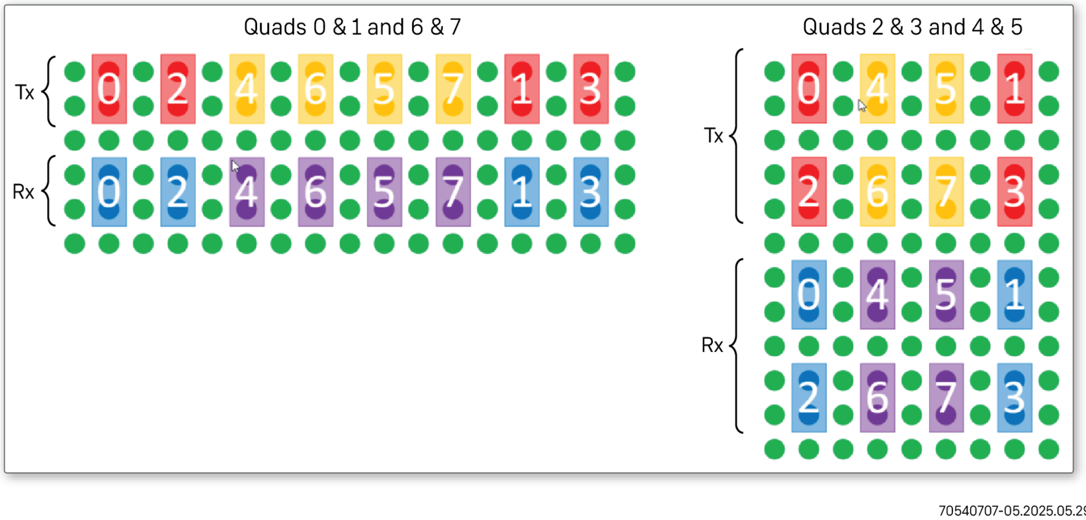

Figure 19 • QSFP-DD Mapping (7t1400/1500)

SerDes Connectivity – QSFP-DD Mapping for 7t700/800

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.

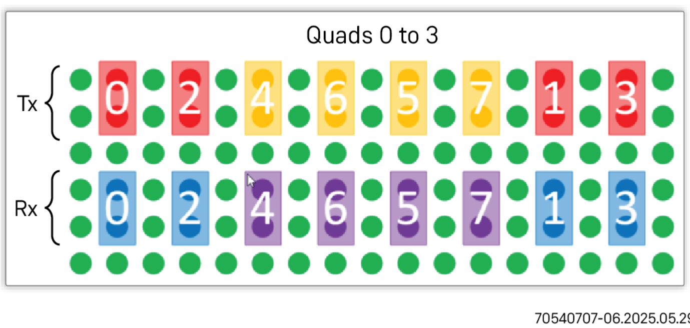

Figure 20 • QSFP-DD Mapping (7t700/800)

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:

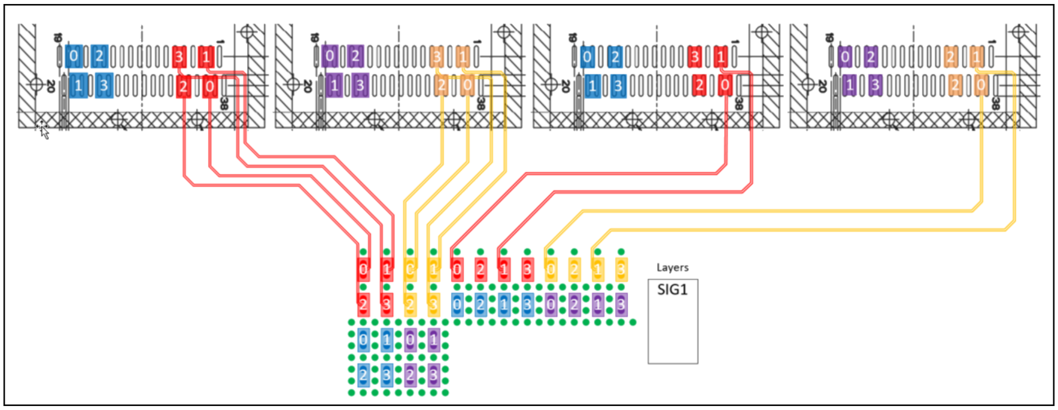

Figure 21 • Routing to QSFP-28 Connectors – Layer 1

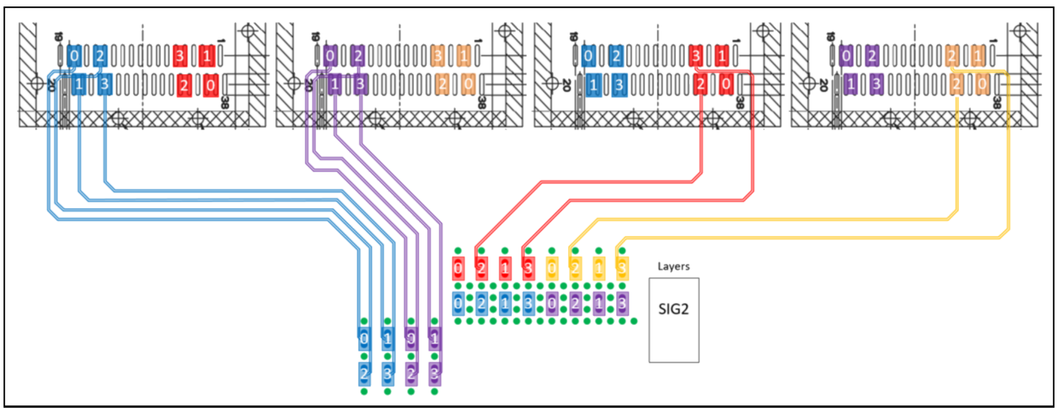

Figure 22 • Routing to QSFP-28 Connectors – Layer 2

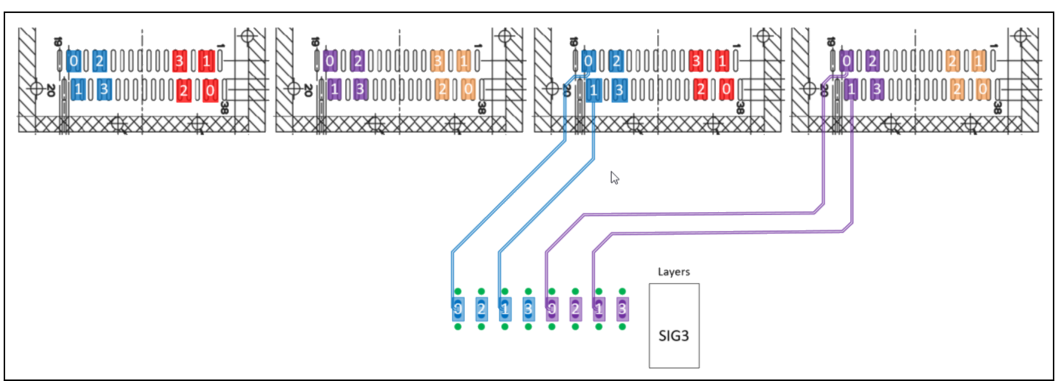

Figure 23 • Routing to QSFP-28 Connectors – Layer 3

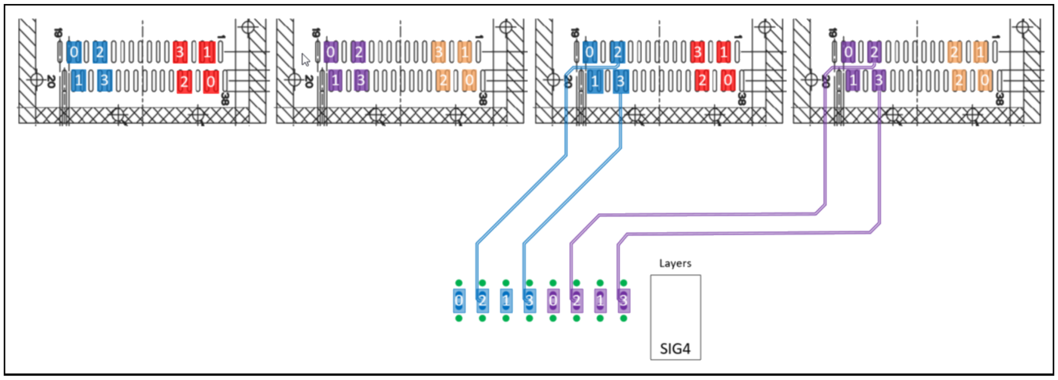

Figure 24 • Routing to QSFP-28 Connectors – Layer 4

Mapping Example – QSFP-DD

The QSFP-DD mapping works 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 are still required.

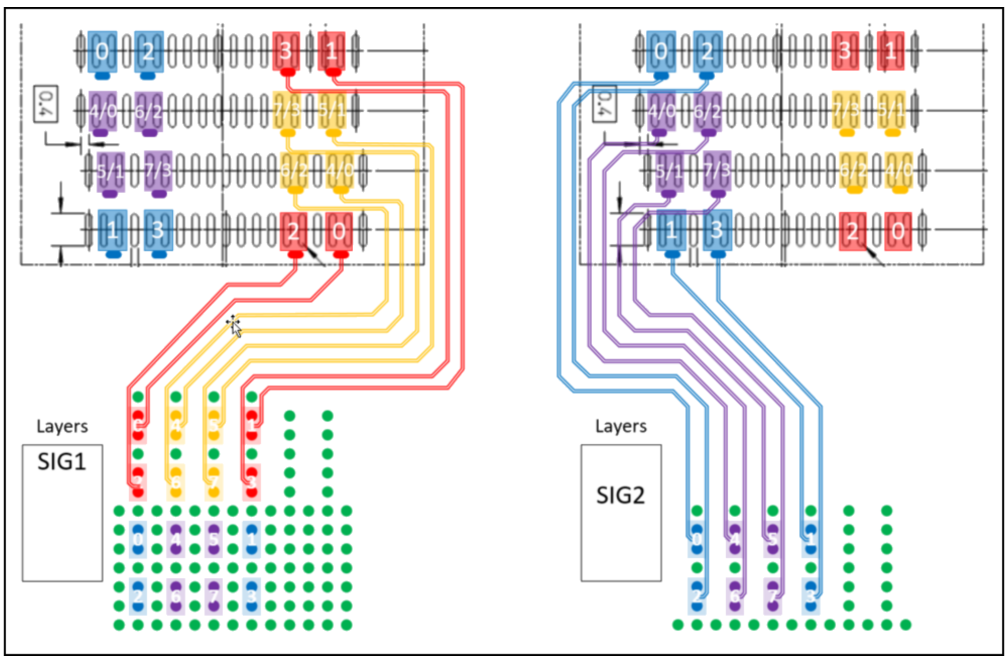

Figure 25 • Routing to QSFP-DD Connectors – Layer 1 & 2

Sample layouts can be made available for specific cases. Contact Achronix support for details.

Chapter 5 : 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 Speedster7t FPGAs:

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

By-16 Configuration

Data connection is point to point for the ×16 mode. Both channels act 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 FPGA and the other being a DRAM on the other endpoint.

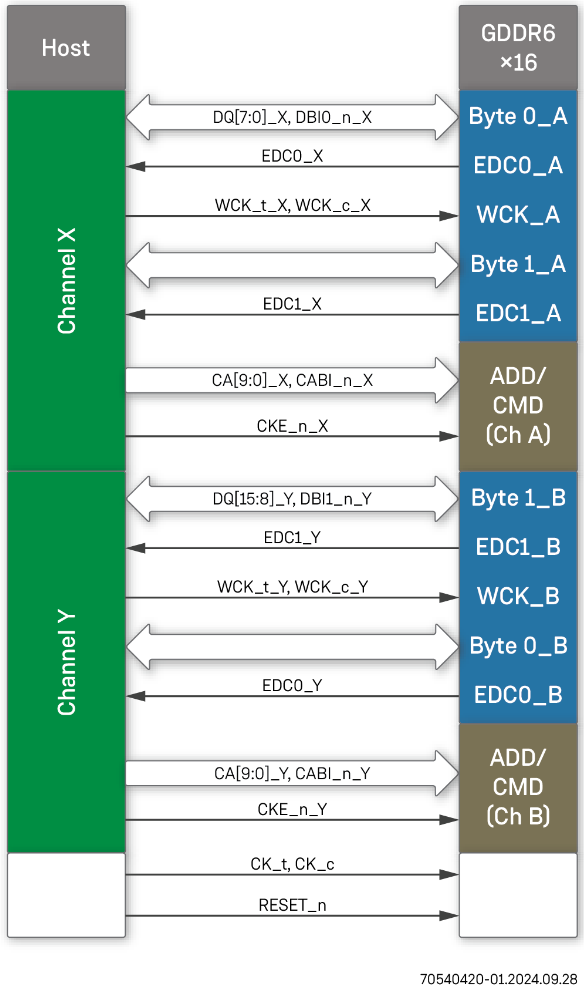

Figure 26 • System View of a ×16 Configuration

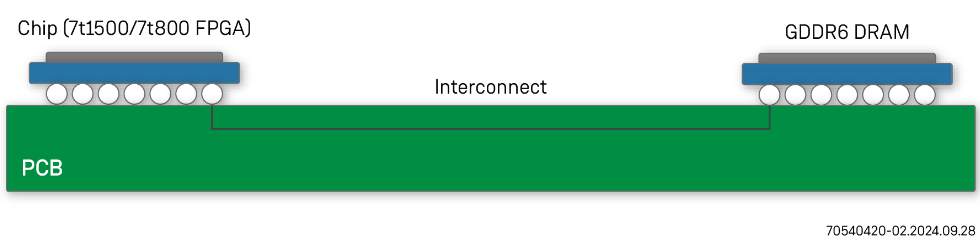

Figure 27 • Data Connection in ×16 Mode

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.

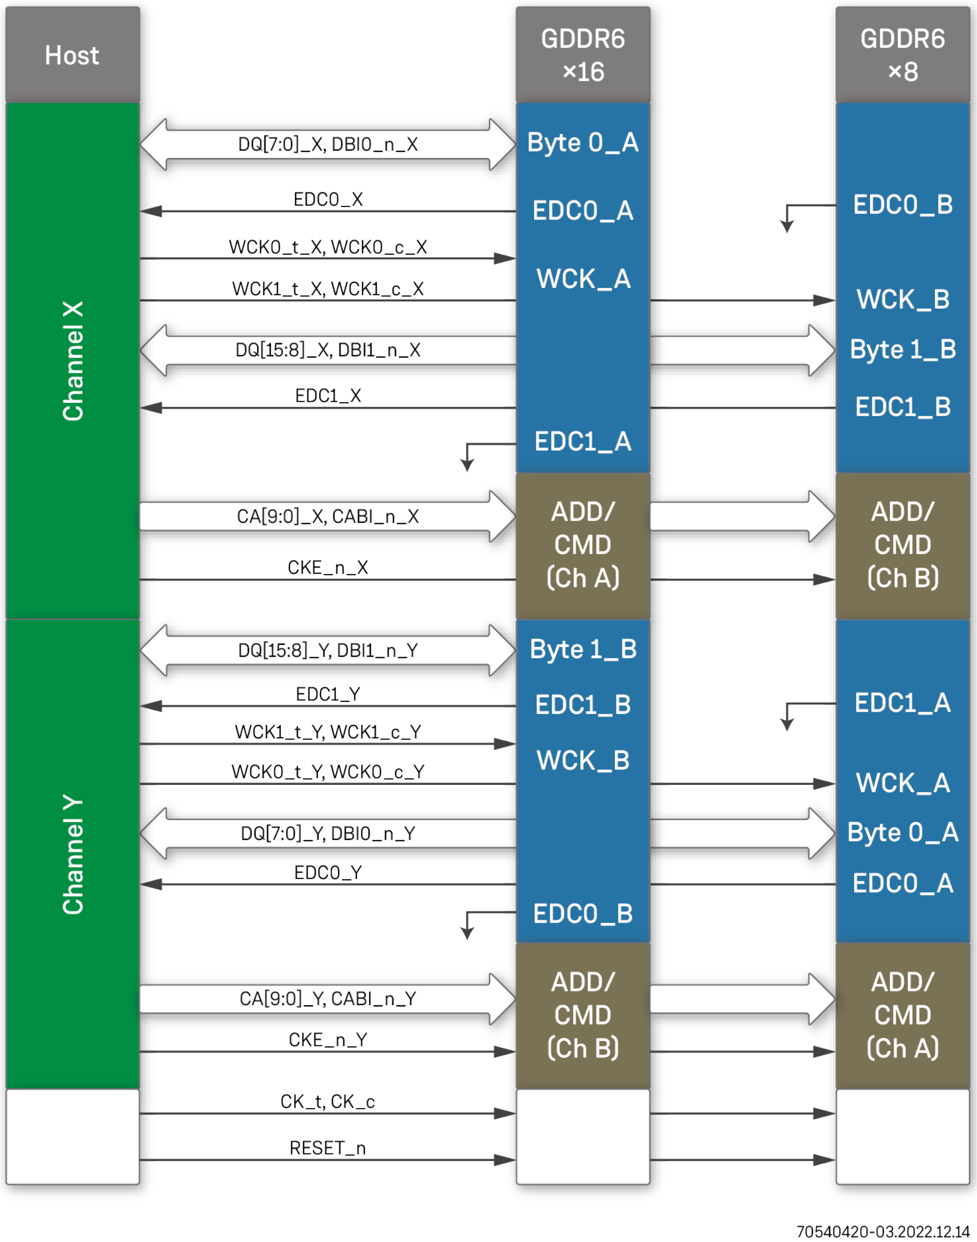

Figure 28 • System View of a ×8 Clamshell (Dual-Memory) Configuration

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.

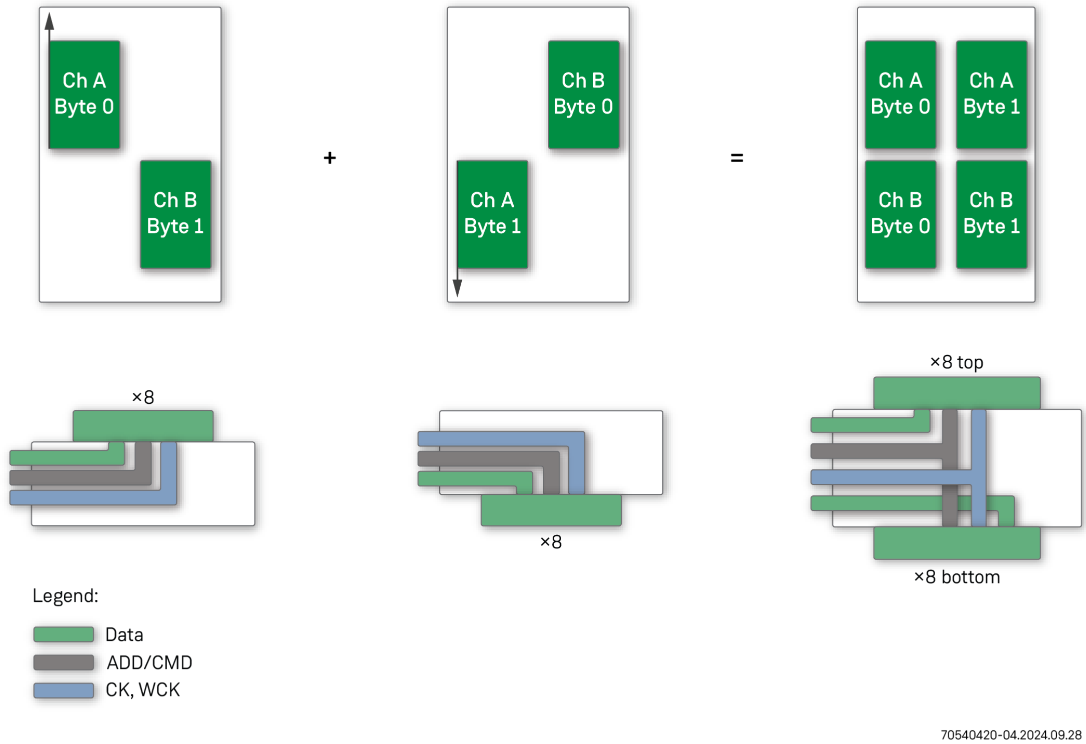

Figure 29 • Layer Assignments for Clamshell Mode

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.

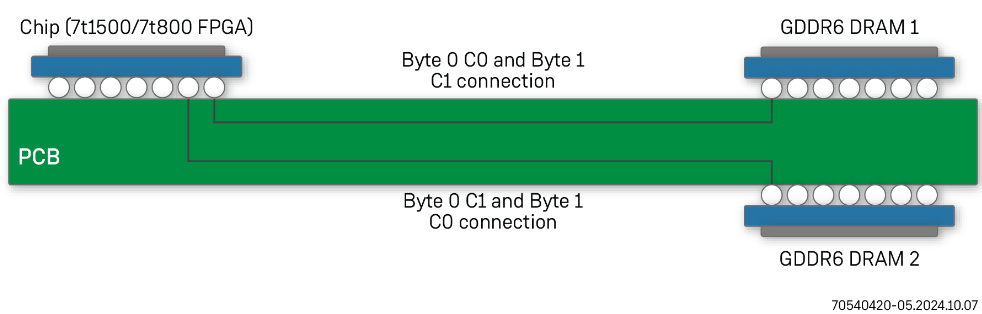

Figure 30 • Data Connection in ×8 Mode

in ×8 configuration follows 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.

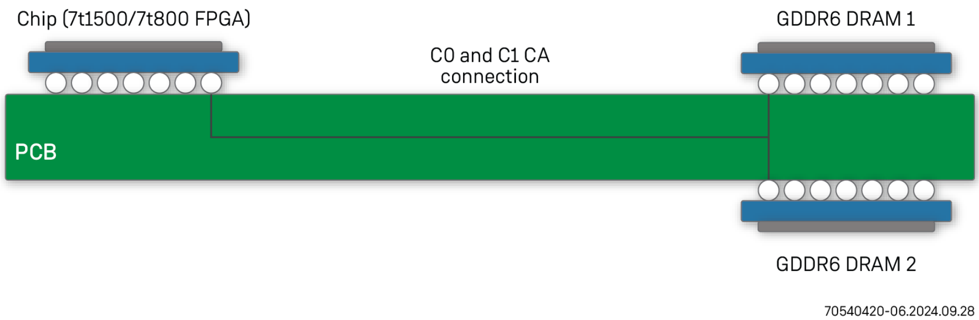

Figure 31 • CA Connection in ×8 Mode

GDDR6 Mapping from ACE for 7t1400/1500

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). Within ACE, the interfaces are assigned sequential numbers. At the package level (physical label), 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

ACE Label	Physical Label		ACE Label	Physical 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**.

GDDR6 Mapping from ACE for 7t700/800

In 7t800, the three GDDR6 interfaces are referenced in ACE and the physical domain (schematic/layout) in the same way.

Table 4 • GDDR6 Mapping to External References

ACE Label	Physical Label
GDDR6_0	GDDR6_0
GDDR6_1	GDDR6_1
GDDR6_2	GDDR6_2

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 5 • 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 @ 2 GHz
	CLK_P/N < −20 dB @ 2 GHz
DQ-DQ power sum near-end crosstalk[(1)] (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)
DQ-CA/CA-DQ power sum crosstalk[(1)] (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)	< -30 dB @ 4GHz < -25 dB @ 8 GHz

Parameter	Specification
CA-CA power sum crosstalk[1] (cumulative sum of crosstalk from all other CA signals on a victim CA pin)	< -25 dB @ 2 GHz < -21 dB @ 4 GHz (desirable)
EDC-DQ individual crosstalk[1] (individual crosstalk between any EDC pin and any DQ pin)	< -33 dB @ 4 GHz and < -30 dB at 8 GHz
EDC-CA individual crosstalk[1] (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[1] (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

1. Both near-end crosstalk (NEXT) and far-end crosstalk (FEXT) must meet the crosstalk specs provided separately in the above table.
2. Impedance recommendations are based on package impedance 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 Speedster7t FPGA 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 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 6 • GDDR6 Length Matching Targets

Signal to Target	Match Requirement	Suggested Target1(Conservative)[2]	Suggested Target2 (Aggressive)[3]
DQ [[7:0]] to WCK0 DQ [[15:8]] to WCK1	±1 UI of DQ = ±62.5 ps	±0.8 UI of DQ = ±50 ps	±0.95 UI of DQ = ±59.375 ps
CA[[9:0]] to SD_CLK	±0.5 UI of DQ = ±31.25 ps	±0.5 UI of DQ = ±31.25 ps	±0.5 UI of DQ = ±31.25 ps
WCK[[1:0]] to SD_CLK [4]	±2 UI of DQ = ±125 ps	±1UI of DQ = ±62.5 ps	±2UI of DQ = ±125 ps

Table Note

1. Achronix has been successful in achieving these matching targets within its own designs. They serve as a useful but flexible design target for designers.
2. Suggested target 1 can be overridden with suggested target 2 if required. The suggested target 1 is just created with some extra margin.
3. Suggested target 2 has been created based on feedback for GDDR6 IP. This target ensures that the complete silicon limit is not being utilized (i.e., providing some extra margin).
4. GDDR6 IP does not have a specification for this skew but the recommendation is to keep skew within 4 UI of DQ.

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 x16 and ×8 (clamshell mode), the data byte connection is always point to point and hence, the simulation setup is similar.

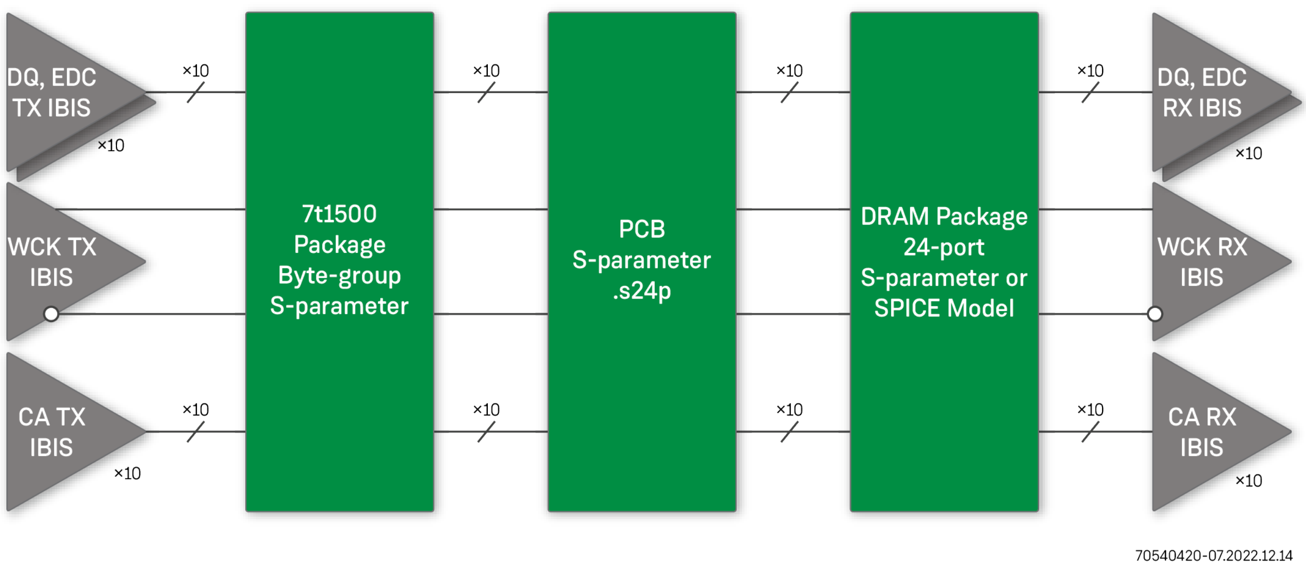

Figure 32 • GDDR6 DQ Write Simulation Setup

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 Speedster7t FPGAs GDDR6 DQ/DBI/EDC and WCK I/O. These models are available from Achronix upon request (support account required).
- **Package S-parameter** – 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 Speedster7t 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. Achronix can provide full package S-parameter .s148p

for the 7t1400/1500 and .s152 touchstone file for 7t700/800. Recommended is to simulate all the data, command address, and clock signals together for tape out, in order to get better idea on eye-margin.

- **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 7t1400/1500 device (excluding 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 x16 and ×8 (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.

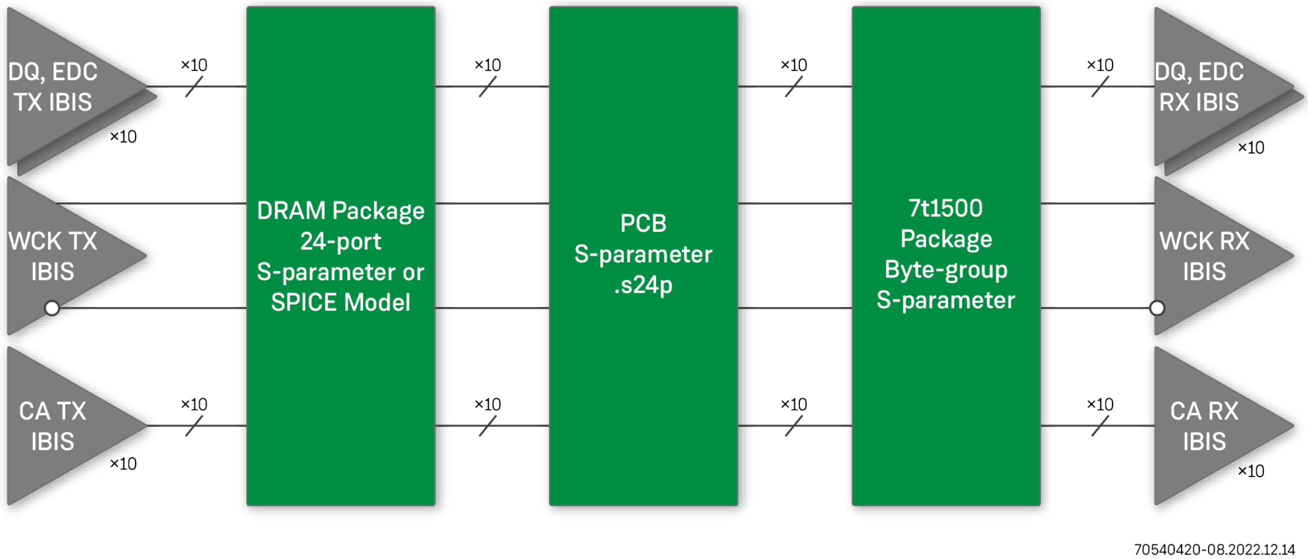

Figure 33 • GDDR6 Read Simulation Setup

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 Speedster7t FPGA 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.

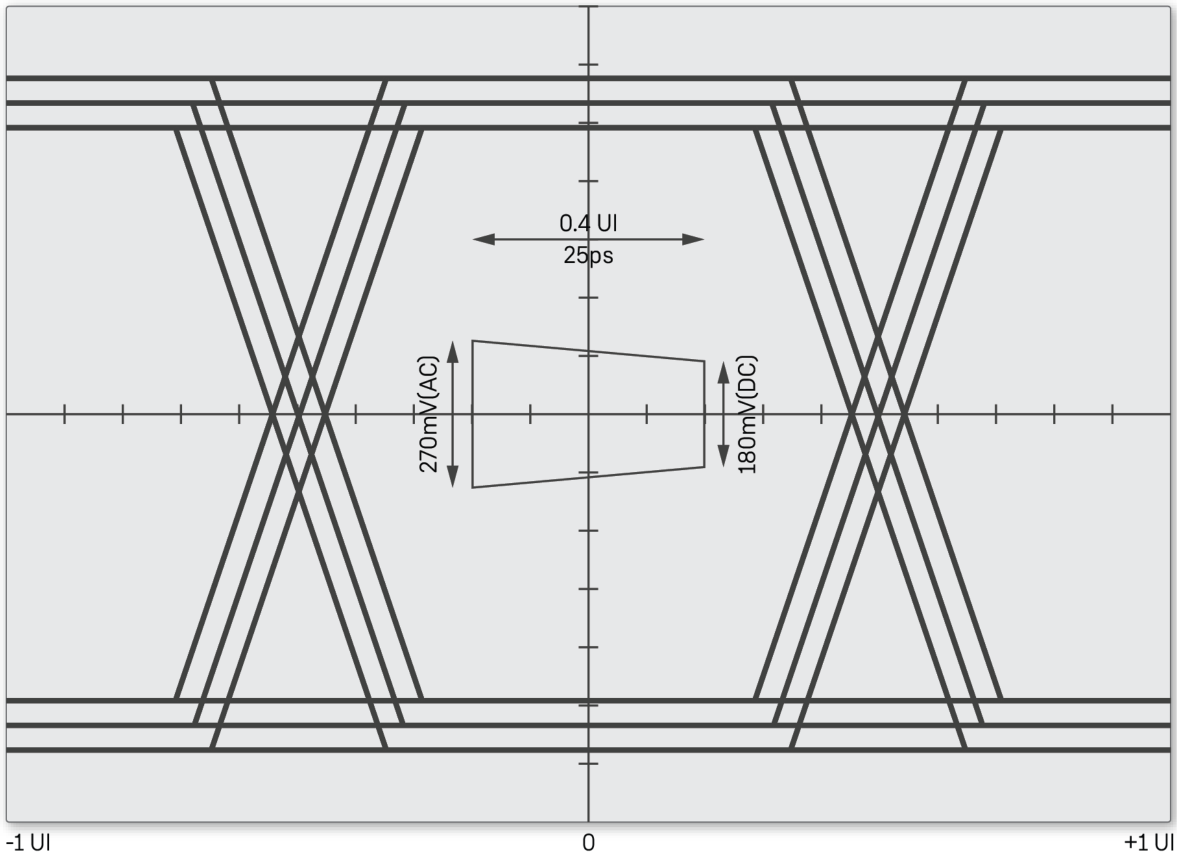

Figure 34 • Read-Eye Mask

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.

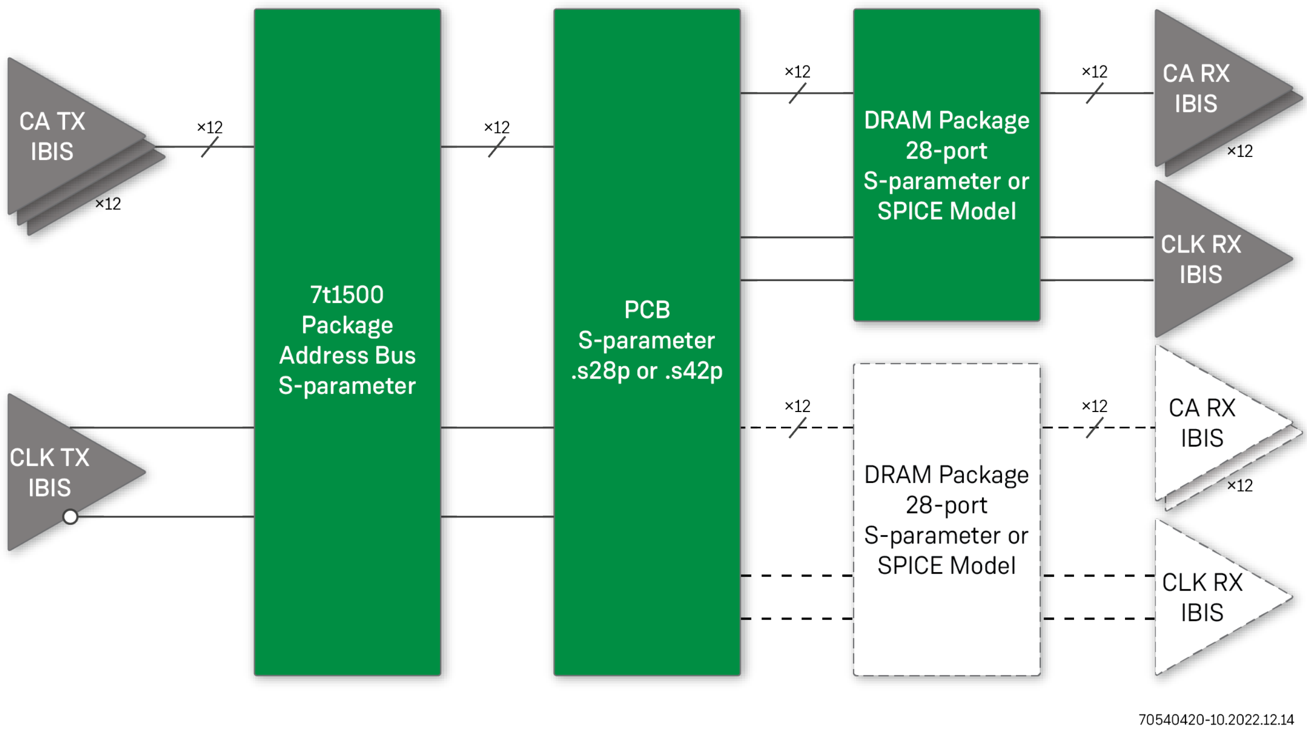

Figure 35 • GDDR6 Command-Address Bus Simulation Setup

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 7t1400/1500's GDDR6 CA and CLK pair signals. These models are available from Achronix upon request.
- **Package S-parameter** – 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 Speedster7t 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. The S-parameter model for only CAs for 7t700/800 would be .s30p as there signals are CA0-CA10, CKE_N, CABI_N and CLK_P/N. For 7t700/800, the full-package S-parameter model (.s152p) can be provided by Achronix upon request which includes EM extraction of all nets (DQs, DBIs, WCKs, EDCs, CAs, CKEs, CABIs, CLK). .s30p for only CAs can be extracted through .s152p touchstone file by terminating other ports.
- **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 transmit 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 FPGA receiver
- Eye height requirement = 360 mV (AC) and 270 mV (DC)
- Eye width requirement = 125 ps (0.5 UI of CA @ 4 Gbps)

Note

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.

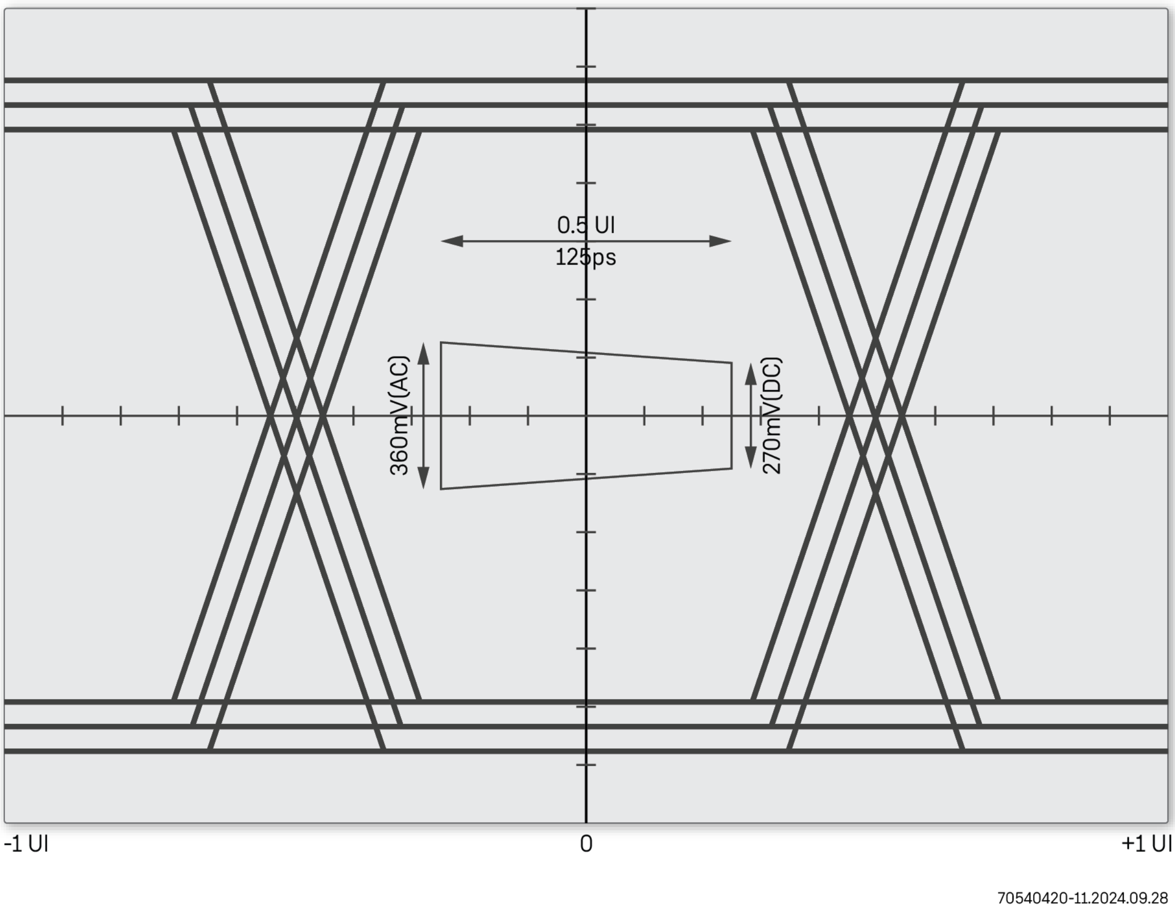

Figure 36 • CA Eye Mask

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** (page 0)", for general direction on constructing the PCB stack-up

Design the stack-up to keep via stubs to a minimum (less than ¼ 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, back-drill 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 is 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.

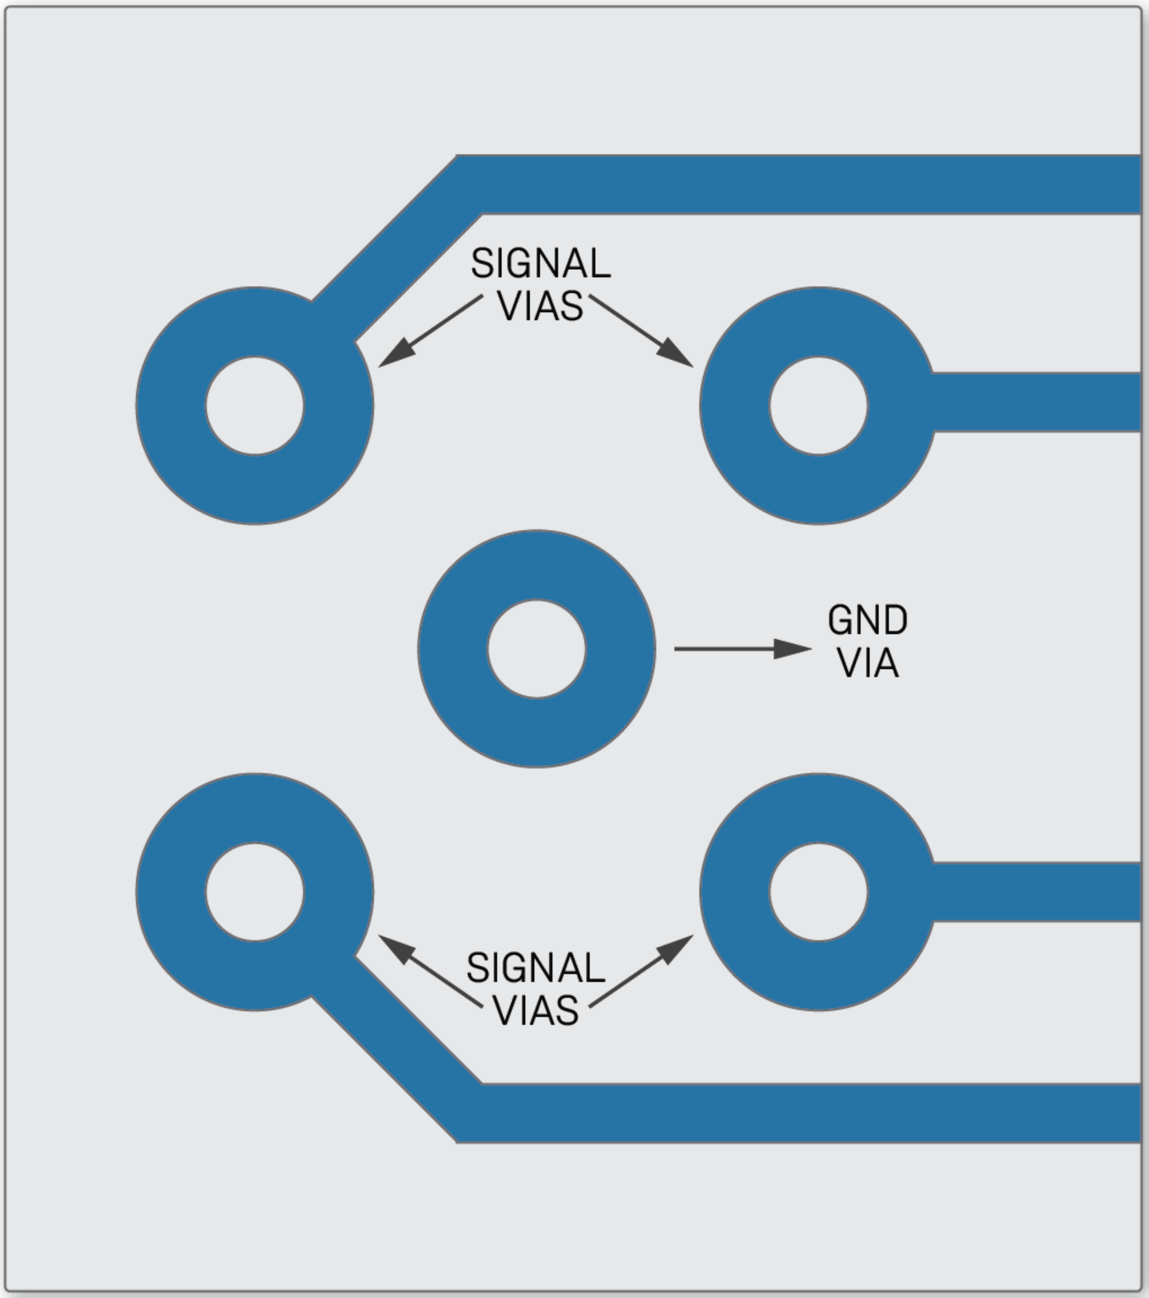

Figure 37 • Via Crosstalk Optimization

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 memory as it is not designed to support this strategy.*

For general routing guidelines, see the section, "**Routing Guidelines** (page 10)" in the chapter, "**PCB General Considerations** (page 4)".

Chapter 6 : DDR4 Interfaces

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 Mbps. This high frequency of operation requires the package and PCB design to be optimized for minimal losses and minimal crosstalk.

This chapter provide 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 may 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:

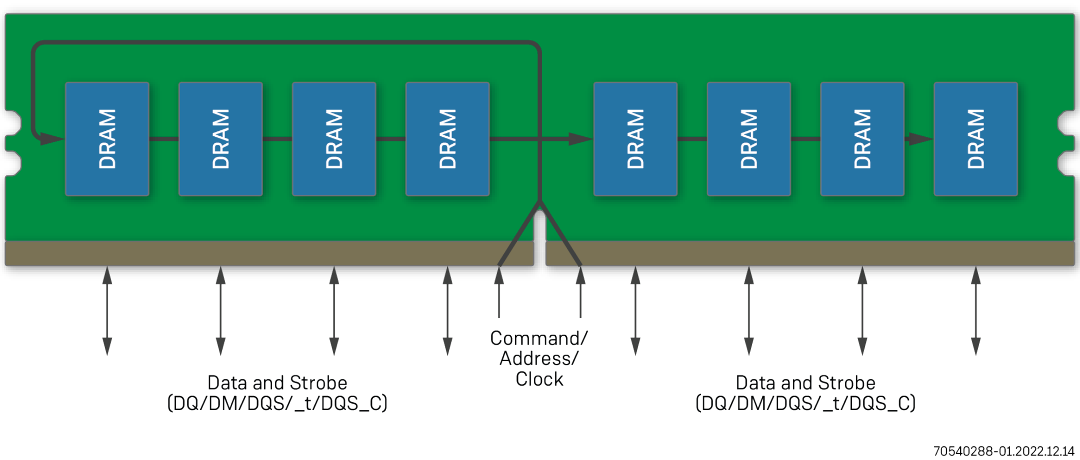

Figure 38 • Fly-By Routing Used in Unbuffered DIMMs (Single Rank Shown)

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.

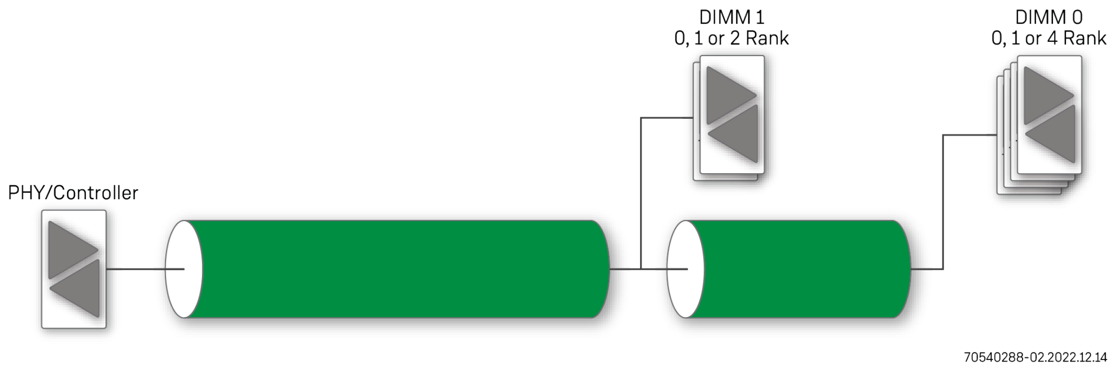

Figure 39 • DQ/DQS/DM Point to Point

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 a 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.

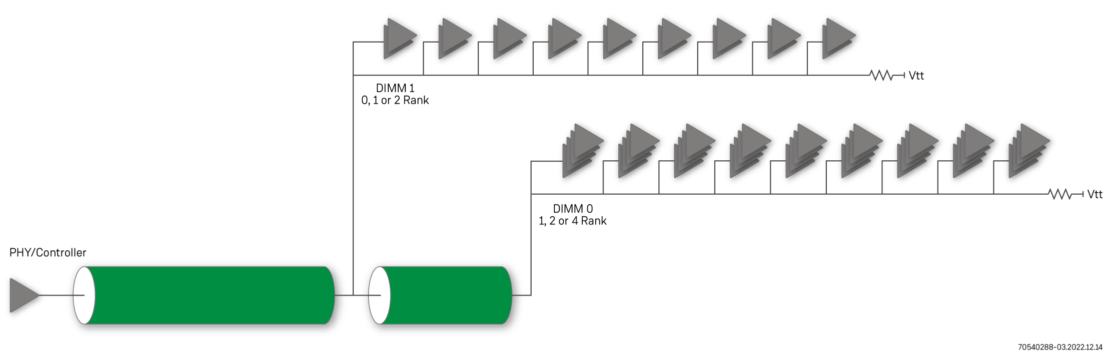

Figure 40 • Unbuffered DIMMs Implemented in Two-Slot Daisy-chain Configuration

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.

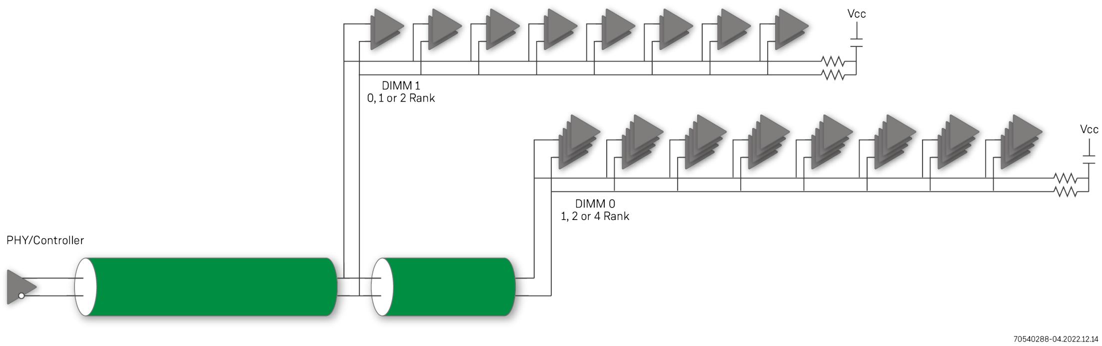

Figure 41 • Daisy-Chain (Fly-By) Topology for Differential Clock in a UDIMM

Signal Integrity Specification

Signal Integrity PCB Design Specification Guidelines

Table 7 • Signal Integrity Specifications for 3200 Mbps 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 3200Mbps	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 JESD79-4D 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 Mbps single-rank operation (from the JEDEC specification), the key specifications are:

- The data UI is 1/3200 ×10 ^{6} = 312.5 ps. The command/address (CA) UI is twice that, resulting in 625 ps.
- JEDEC specify 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. The Achronix controller provides a DLL wherein additional recommended skew of less than 20 ps, and the available deskew limit is less than 200ps.
- t_{DQS2DQ} is the receive mask DQS-to-DQ offset,and is specified at ±0.220 UI, or ±68.8 ps. The Achronix controller provides a DLL wherein additional recommended skew of ±10 ps, and the available deskew limit is ±100 ps.
- 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 -1 UI to 7.97 UI. The recommended skew is: CK position ±75 ps.
- t_{IS} and t_{IH} are the CA setup and hold times (refer to CA mask calculation for the values of t_{IS} and t_{IH}). The controller groups the CA and clock signals into bit groups of four signals each, which can be delayed independently. This group is called AC4X. Refer to the table "AC4X Bit Groups for Address Matching (page 49)" for details on the CA signals which are in different groups. It is required that the signals within an AC4X group have delay less than 5 ps. The CLK to AC4X group skew range is ±25 ps.
- Intra-pair skew between the true and complement legs of a differential pair should be kept below 3 ps to preserve monotonicity of the DQS/DQS# or CK/CK# differential signals.

Table 8 • Deskew Range Availability and Recommended PCB Routing Skew Limits for 7t800 and 7t1500 Devices for DDR4 Operation

Constraints	Available Deskew Range	Recommended Routed Skew Limits	Notes
DQ to DQ arrival time mismatch within a byte or nibble	<200 ps	< 20 psec	It is highly recommended that flight times across the data lanes be tightly matched to preserve operating margin and reduce vertical eye collapse that can occur from crosstalk between unaligned DQ signals.
DQ to DQS domain	DQS position ±100 ps	DQS position ±10 ps	
AC4X group (†) to CK/ CK#	0 to +1 UI	CK position ±25 ps	PHY can be programmed to add delay to 4-bit groups of AC signals to address potential skew violations.
DQS to CK domain	-1.0 to +7.97 clock cycles	CK edge position ±75 ps (without write leveling)	Write Leveling training can compensate for delay differences that extend over multiple clock cycles.

Table Note

† Refer to the AC4X bit group tables below for information on CA signals in different AC4X groups.

Table 9 • AC4X Bit Groups for Address Matching (AC7t800)

Group No	Signals
1	DDR45_0_CKE0_CSA0, DDR45_0_CKE1_CSA1, DDR45_0_CKE2_CSA2, DDR45_0_CKE3_CSA3
2	DDR45_0_BG0_CA_A0, DDR45_0_MALERT_N_MALERT_N, DDR45_0_BG1_CA_A2, DDR45_0_ACT_N_NULL
3	DDR45_0_A9_CA_A1, DDR45_0_A12_NULL, DDR45_0_A11_CA_A3, DDR45_0_A7_NULL
4	DDR45_0_A8_CA_A4, DDR45_0_A6_NULL, DDR45_0_A5_CA_A6, DDR45_0_A4_NULL
5	DDR45_0_A3_CA_A5, DDR45_0_A2_NULL, DDR45_0_A1_PAR_A, DDR45_0_BA1_NULL
6	DDR45_0_PAR_CA_B0, DDR45_0_A13_NULL, DDR45_0_BA0_CA_B2, DDR45_0_A10_NULL
7	DDR45_0_A0_CA_B1, DDR45_0_CAS_N_NULL, DDR45_0_WE_N_CA_B3, DDR45_0_RAS_N_NULL
8	DDR45_0_C0_CA_B4, DDR45_0_C1_RSPN_A0, DDR45_0_C2_CA_B6, DDR45_0_A17_RSPN_A1
9	DDR45_0_CS_N0_CA_B5, DDR45_0_CS_N1_RSPN_B0, DDR45_0_CS_N2_PAR_B, DDR45_0_CS_N3_RSPN_B1
10	DDR45_0_ODT0_CSB0, DDR45_0_ODT1_CSB1, DDR45_0_ODT2_CSB2, DDR45_0_ODT3_CSB3

Table 10 • AC4X Bit Groups for Address Matching (AC7t1500)

Group No.	Signals
1	DDR4_S0_CKE_0, DDR4_S0_CKE_1, DDR4_S0_CKE_2, DDR4_S0_CKE_3
2	DDR4_S0_BG_0, DDR4_S0_BG_1, DDR4_S0_ACT_N, DDR4_S0_A_9
3	DDR4_S0_A_12, DDR4_S0_A_11, DDR4_S0_A_7, DDR4_S0_A_8
4	DDR4_S0_A_6, DDR4_S0_A_5, DDR4_S0_A_4, DDR4_S0_A_3
5	DDR4_S0_CK_P_0, DDR4_S0_CK_N_0, DDR4_S0_CK_P_2, DDR4_S0_CK_N_2
6	DDR4_S0_CK_P_1, DDR4_S0_CK_N_1, DDR4_S0_CK_P_3, DDR4_S0_CK_N_3
7	DDR4_S0_A_2, DDR4_S0_A_1, DDR4_S0_BA_1, DDR4_S0_PAR
8	DDR4_S0_A_13, DDR4_S0_BA_0, DDR4_S0_A_10, DDR4_S0_A_0
9	DDR4_S0_CAS_N, DDR4_S0_WE_N, DDR4_S0_RAS_N, DDR4_S0_CID_0
10	DDR4_S0_CID_1, DDR4_S0_CID_2, DDR4_S0_A_17, DDR4_S0_CS_N_0

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) IBIS behavioral models** – These models are specific to the ACt800 and ACt1500 devices and are available from Achronix upon request (support account required).
- **Package S-parameter models** – These models cover the AC7t 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 S-parameter models cover the PCB parameters from package BGA bumps up to until memory BGA bumps.
- **Connector S-parameter** – This model covers the DIMM/SO-DIMM connector, if present, and must be provided by the connector manufacturer.
- **HSPICE models –** These are the DIMM card models from the memory vendor.
- **Receiver (RX) IBIS behavioral 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.

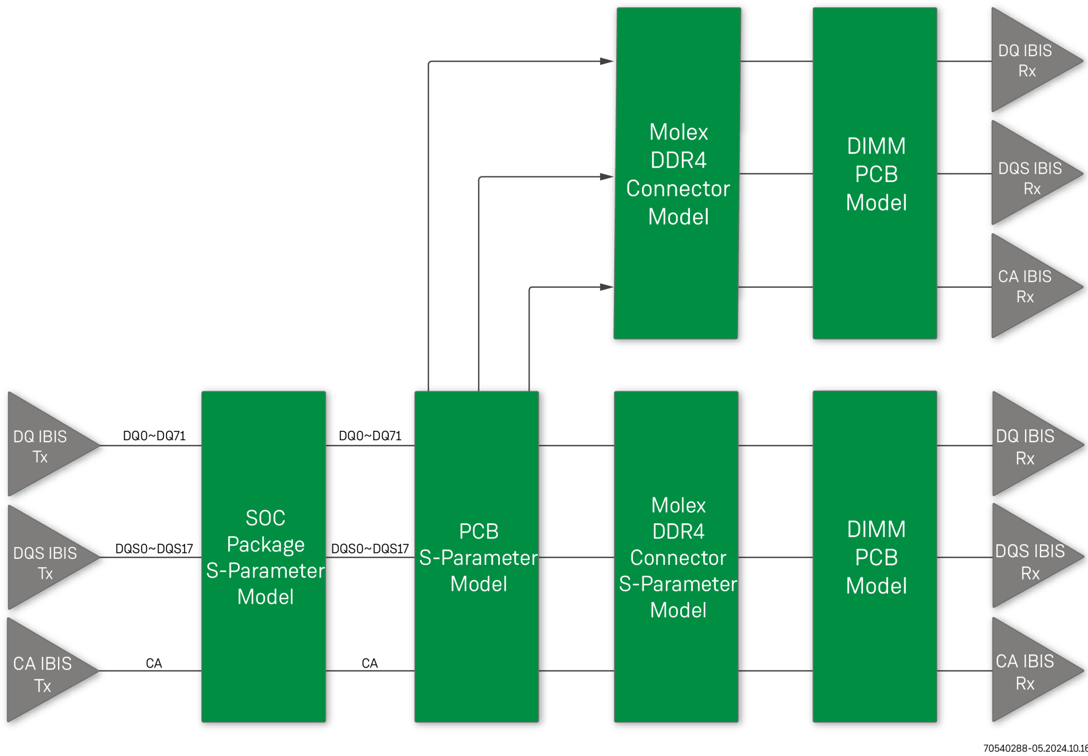

Figure Note

Solid lines represent singe-module configuration. Dashed lines represent an optional second module.

Figure 42 • DDR4 Write Simulation Setup

Signal Integrity Specification for Data Write Cycle

The eye mask shown below is taken from JEDEC specification (JESD79-4D). Refer to latest JEDEC specification for the eye mask requirement at the DRAM component.

Key DDR4 eye diagram terms are:

- **V_{CENT} Value** – Defined as the midpoint between the largest DQ reference voltage and the smallest DQ reference voltage level, computed using voltage training.
- **Unit interval (UI)** – 312.5 ps for 3.2 Gbps operation of DDR4
- **Receive mask voltage (VdivW)** – JESD79-4D calls out greater than 110 mV; refer to JEDEC document.
- **Receive timing Window (TdivW)** – ESD79-4D calls out greater than 0.23 UI = 71.875 ps; refer to JEDEC document.

Jitter Analysis (AC7t1500)

Table 11 • Transmitter Jitter Budget Table (Write) - (AC7t1500)

Component		Setup (ps)	Hold (ps)	Notes
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 Transmitter Components	Linear Summation	37.5	42.2	
	RSS Summation	23.4	27.6	
DRAM Receiver Mask		35.9	35.9	JEDEC mask requirement: 0.23 UI at 3200 Mbps
Interconnect allowance	Linear Summation	73.4	78.1	
	RSS Summation	59.3	63.5	

Write Cycle Mask Calculation (AC7t1500)

- Speed Grade: 3.2 Gbps
- UI (Unit interval) = 312.5 ps
- VdivW (receive mask voltage): JEDEC-4D calls out 110 mV; refer to the latest JEDEC document.
- TdivW (receive timing window): JEDEC-4D calls out >0.23 UI; refer to the latest JEDEC document.
- Setup and hold timing

- Setup timing = TdivW/2
- Hold timing = TdivW/2

- Transmit jitter linear sum
  - Setup jitter = 37.5 ps
  - Hold jitter = 42.2 ps
- Transmit jitter RSS sum
  - Setup jitter = 23.4 ps
  - Hold jitter = 27.6 ps
- Receive mask including transmit jitter (linear sum)
  - Setup timing = (TdivW/2) + 37.5 ps
  - Hold timing = (TdivW/2) + 42.5 ps
- Receive mask including transmit Jitter (RSS Sum)
  - Setup timing = (TdivW/2) + 23.4 ps
  - Hold timing = (TdivW/2) + 27.6 ps
- Below is an example of the mask coordinates (linear summation, TdivW = 0.23 UI) that can be used in Keysight ADS tool:
  - 1
  - 4
  - 0.26512 UI, V_{CENT} + 55 mV
  - 0.74992 UI, V_{CENT} + 55 mV
  - 0.26512 UI, V_{CENT} – 55mV
  - 0.74992 UI, V_{CENT} – 55mV

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 models are specific to the DRAM vendor and are available from Achronix upon request (support account required).
- **Package S-parameter models** – These models cover the AC7t 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 S-parameter models cover the PCB parameters from package BGA bumps up to until memory BGA bumps.
- **HSPICE models –** These are the DIMM card models from the memory vendor.
- **Receiver (RX) IBIS behavioral models** – These models are specific to the AC7t1500 and AC7t800 devices and are available from Achronix upon request.

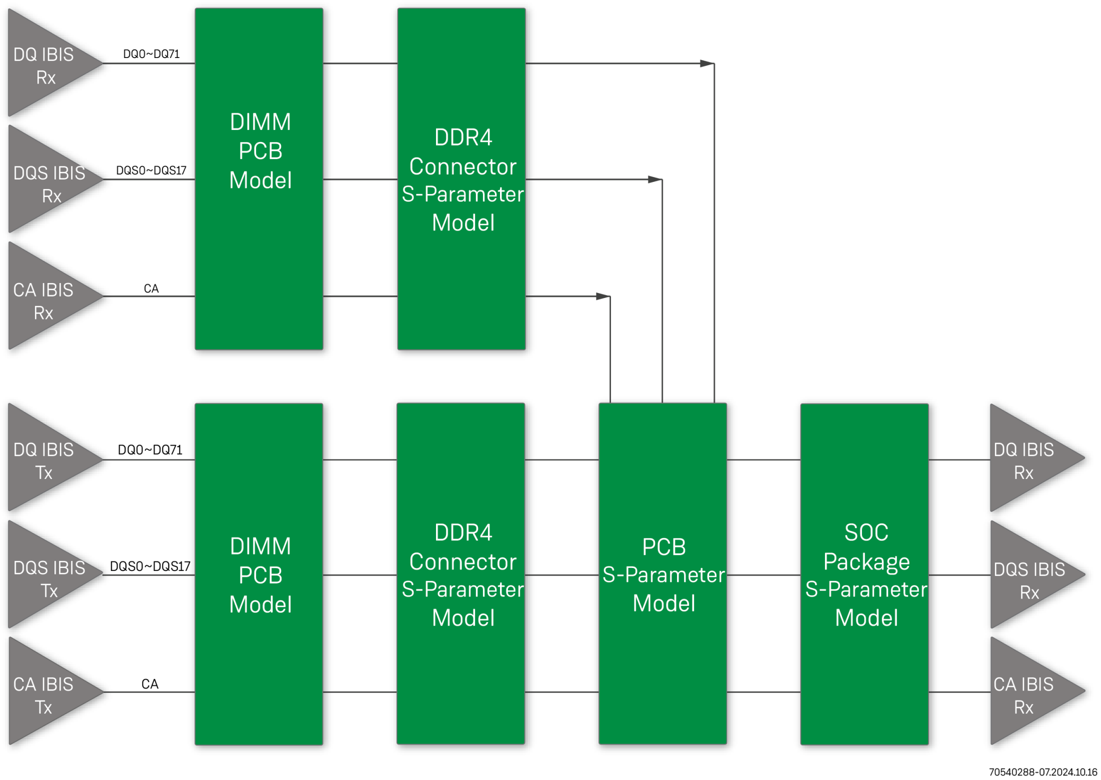

Figure 43 • DDR4 Read Simulation Setup

Jitter Analysis

Table 12 • Read Operation Jitter Budget Table - (AC7t1500)

Component		Setup (ps)	Hold (ps)	Notes
UI width at 3200 Mbps		156.25	156.25	
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.

Component		Setup (ps)	Hold (ps)	Notes
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 setup/hold requirements		7.0	4.0	
Total Receiver Components	Linear Summation	40.7	39.6	
	RSS Summation	21.9	24.5	
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	Linear Summation	59.8	59.8	
	RSS Summation	46.6	46.6	
Interconnect allowance	Linear Summation	55.8	56.9	
	RSS Summation	87.8	85.2	

Signal Integrity Specification for Data Read Cycle

DDR4 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 = 140 mV
- Minimum eye width (setup+hold) requirement (at the AC7t1500 device die pin level) @ 1e-16 BER (Linear summation) = 40.7 + 39.6 + 59.8 + 59.8 = 199.9ps

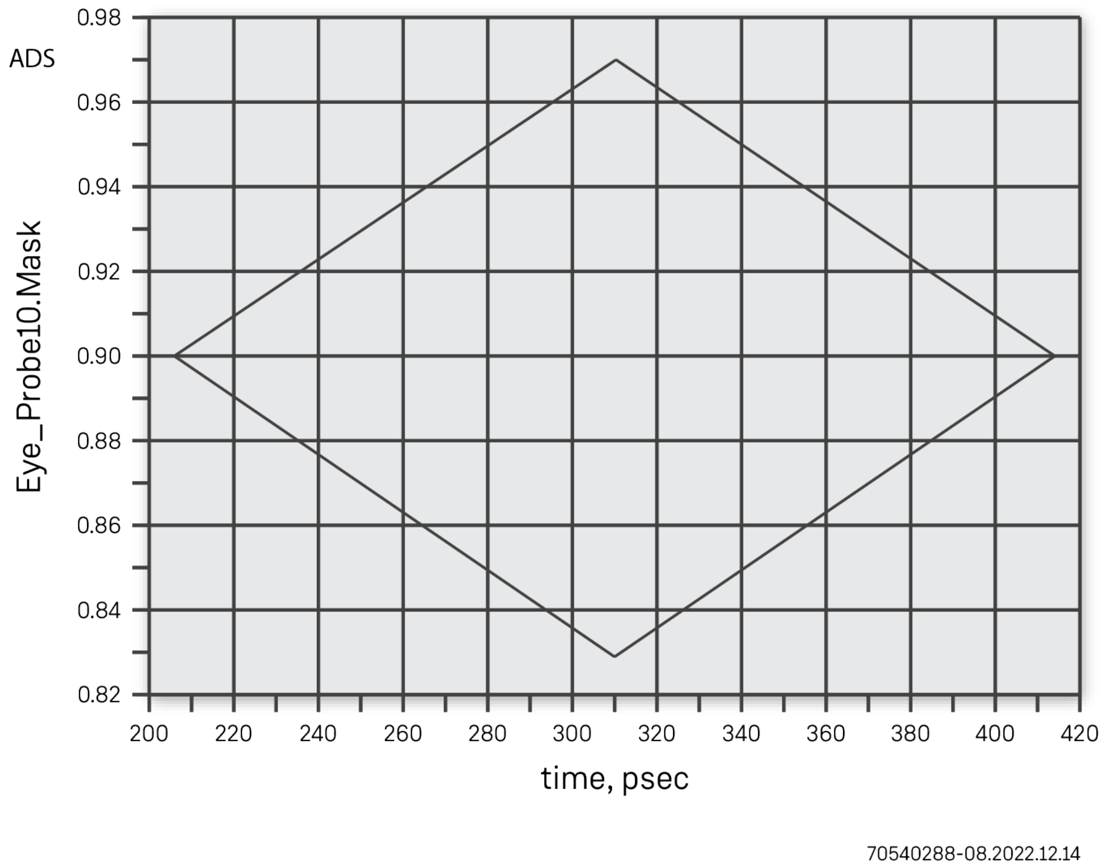

Figure 44 • Read Cycle Eye Mask

Read Cycle Mask Calculation

- Receiver voltage requirement = 140 mV peak-peak (7t1500 device requirement)
- Setup and hold timing:
  - Setup timing = 43.8 ps
  - Hold timing = 43.8 ps
- Total DRAM transmitter jitter (linear summation):
  - Setup jitter = 59.8 ps
  - Hold jitter = 59.8 ps
- Total DRAM transmitter jitter (RSS summation):
  - Setup jitter = 46.6 ps
  - Hold jitter = 46.6 ps
- Receiver jitter (linear summation)
  - Setup jitter = 40.7 ps

- Hold jitter = 39.6 ps
- Receiver jitter (RSS summation)
  - Setup jitter = 21.9 ps
  - Hold jitter = 24.5 ps
- Read mask (linear summation):
  - Setup timing = 59.8 ps + 40.7 ps = 100.5 ps
  - Hold timing = 59.8 ps + 39.6 ps = 99.4 ps
- Read Mask (RSS Summation):
  - Setup timing = 46.6 ps + 21.9 ps = 69.5 ps
  - Hold timing = 46.6 ps + 24.5 ps = 71.1 ps
- Below is an example of the mask coordinates (linear summation) that can be used in Keysight ADS tool:
  - 1
  - 4
  - 0.1784 UI, V_{CENT}
  - 0.5 UI, V_{CENT} + 70 mV
  - 0.81808 UI, V_{CENT}
  - 0.5 UI, V_{CENT} – 70 mV

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) IBIS behavioral models** – These models are specific to the AC7t800 and AC7t1500 devices and are available from Achronix upon request.

Final sign-off criteria must also meet the specifications provided by JEDEC at the DRAM component.

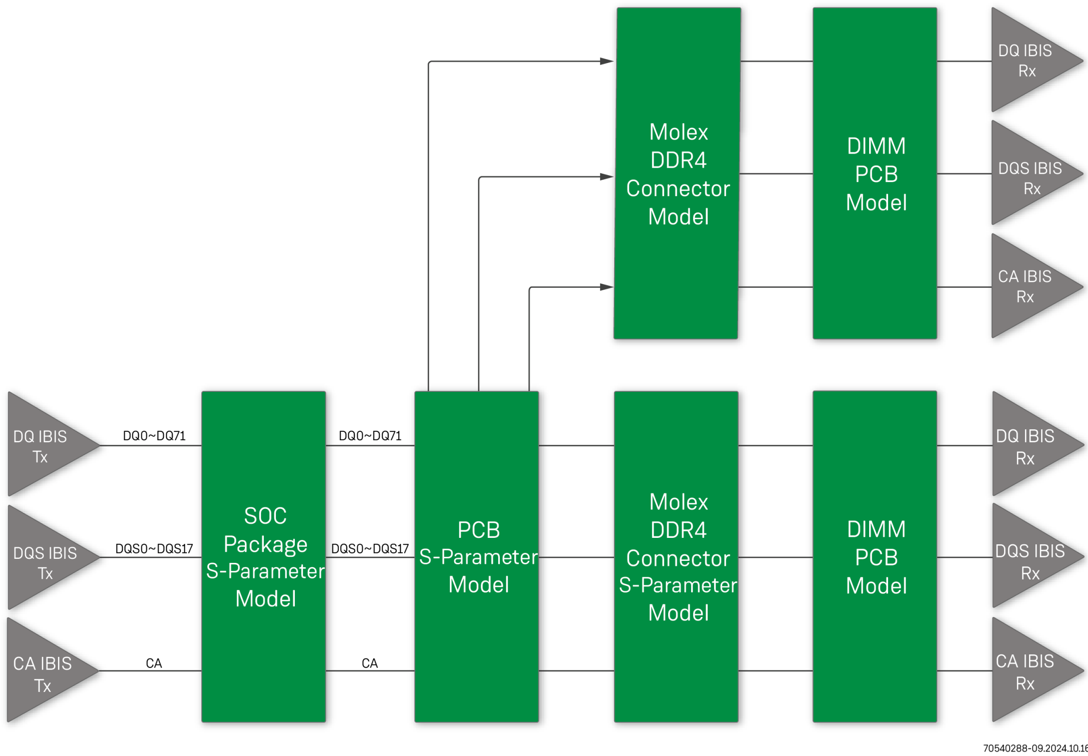

Figure 45 • DDR4 Command Address Write Setup

Jitter Analysis (AC7t1500)

Table 13 • Jitter Budget Tables for CA RDIMM (AC7t1500)

CA Transmit Budget Tables Linear Summation				
Component		**Setup (ps)**	**Hold (ps)**	**Notes**
UI width at 3200 Mbps - 1600 MHz		312.5	312.5	Single data rate - 1T timing
Data-dependent jitter				
	Output rise/fall mismatch	11.3	11.3	Data-dependent jitter, power supply induced jitter, rise/fall delay differences. V_{DDQ} rail ±2.5%, V_{DD} rail ±2%

CA Transmit Budget Tables Linear Summation				
Component		**Setup (ps)**	**Hold (ps)**	**Notes**
	V_{DD} - PSIJ (±2.5%)	18.5	18.5	Jitter induced by noise on the V_{DD} rail
	V_{DDQ} - PSIJ (±5%)	11.8	11.8	Jitter induced by noise on the V_{DDQ} rail
On-die mismatch		25.0	25.0	Process delay variations across CA bus.
PLL jitter		6.6	6.6	RefClk feed-through, internal noise, supply modulation
Total Transmitter Components	Linear Summation	73.2	73.2	
	RSS Summation	59.2	59.2	
RCD receiver requirement		46.9	46.9	JEDEC requirement for RCD. 0.15 × t_{CK}
CK-to-CA routing skew		25.0	25.0	
Interconnect allowance	Linear Summation	167.4	167.4	Interconnect uncertainty from simulation measured at DRAM receiver thresholds.
	RSS Summation	181.4	181.4	

Jitter Analysis (AC7t800)

Table 14 • Jitter Budget Tables for CA RDIMM (AC7t800)

CA/CS Budget		Set Up (ps)	Hold (ps)	
Available budget at 3200 Mbps		312.5	312.5	
Jitter and error sources				Components that erode set up and hold during the CA/CS operation.
SOC transmit				

CA/CS Budget		Set Up (ps)	Hold (ps)	
PLL output		6.3	6.3	Output 2- cycle jitter from PLL including filtered PLL input jitter. RJ=~70% of total jitter, BER=1E-16
PSIJ - V_{DD} and V_{DDQ} domain		21.6	21.6	Power supply induced jitter on core domain for 5% p-p of V_{DD}(nom) at 200 MHz and 10% p-p of V_{DDQ}(nom) at 200 MHz.
Duty cycle error I/O output		1.5	1.5	Static rising- and falling-edge output delay mismatch. Based on optimized V_{REF} level.
Temperature drift DCD impact		1.9	1.9	
On-die process skew		25.0	25.0	Static I/O output delay variations due to process variations across the AC nibble.
Total SOC Jitter	Linear sum	56.3	56.3	
	Modified RSS sum	50.9	50.9	
DRAM receive				
RCD CA receiver mask		42.2	42.2	JEDEC RCD mask requirement: 0.135 × t_{CK}(avg) at 3200 Mbps
Total receive error	Linear sum	42.2	42.2	
	Modified RSS sum	42.2	42.2	
Margin for interconnect	Linear	214.0	214.0	68.5% UI available for interconnect uncertainty characterized at VcIVW_Total.
	Modified RSS	219.4	219.4	70.2% UI available for interconnect uncertainty characterized at VcIVW_Total.

Command Address Mask Calculation

Below is an example for mask calculation using AC7t1500 linear sum jitter values.

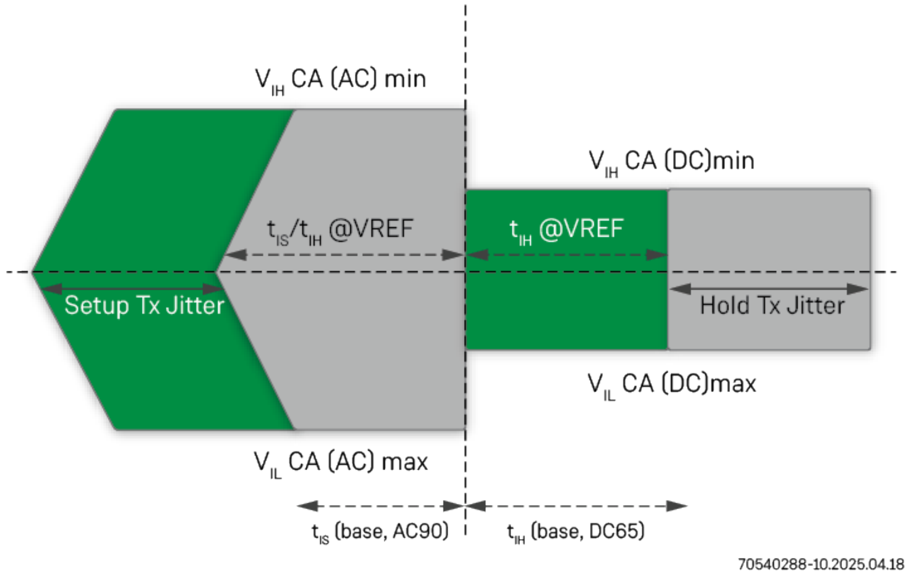

Figure 46 • Command Address Mask

- Below are the calculations for command /address eye mask as per JESD79-4D document.
- Receive voltage requirements from JEDEC:
  - V_{IH} CA (AC)min = 90 mV
  - V_{IL} CA (AC)max = −90 mV
  - V_{IH} CA (DC)min = 65 mV
  - V_{IL} CA (DC)max = −65 mV
- Receive mask timing requirements for CA signal
- From JEDEC Standard, the timing requirement to be met at DRAM is as below:
  - t_{IS} (base, AC 90) = 40 ps
  - t_{IH} (base, DC 65) = 65 ps
  - t_{IS} @ V_{REF} = 130 ps
  - t_{IH} @ V_{REF} = 130ps
- Transmit jitter to be added to the JEDEC timing requirement.
  - SOC setup jitter = 73.2 ps
  - SOC hold jitter = 73.2 ps
- Total receive timing requirement (adding transmit jitter to the DRAM spec requirement) which can be used in signal integrity simulations:
  - t_{IS} @ V_{REF} = 130ps + 73.2ps (setup transmit jitter) = 203.2ps

- t_{IH} @ V_{REF} =130ps + 73.2ps (hold transmit jitter) = 203.2ps
- t_{IS} (base, AC 90) = 40 ps + 73.2ps (setup transmit jitter) = 113.2ps
- t_{IH} (base, DC 65) = 65ps + 73.2ps (hold transmit jitter) = 138.2ps

- 1 UI timing = 625 ps
- Example of receive Keysight ADS mask coordinates for command address signals
  - 1
  - 9
  - 0.3188 UI, V_{CENT} + 90 mV
  - 0.5 UI, V_{CENT} + 90 mV
  - 0.5 UI, V_{CENT} + 65 mV
  - 0.8251 UI, V_{CENT} + 65 mV
  - 0.8251 UI, V_{CENT} – 65 mV
  - 0.5 UI, V_{CENT} – 65 mV
  - 0.5 UI, V_{CENT} – 90 mV
  - 0.3188 UI, V_{CENT} –90 mV
  - 0.1748 UI, V_{CENT}

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.
- 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** (page 0)", 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 AC7t800 and 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 the selection of drive strength and on-die termination to achieve the best signal integrity performance.

For general routing guidelines, see the section, "**Routing Guidelines** (page 10)" in the chapter **PCB General Considerations** (page 4).

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* (UG112)**[4]. Some features which may impact a board designer are included below:

Table 15 • 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 specifications). 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 these standards 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	

4 https://www.achronix.com/documentation/speedster7t-gpio-user-guide-ug112

Table 16 • GPIO Transmit Configurability

Transmit Lane Parameter	Design Specification
LVCMOS Drive Strength	Pin programmable drive strength (2/4/6/8/10/12/14/16 mA) 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 17 • 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 initiator and target) • 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.

GPIO_ZCAL pin requires a 240 ohm ±1% resistor terminated to ground. This pin is used for receive ODT calibration and driver impedance calibration.

For differential transmission lines, the impedance can be programmed to 100Ω. Refer to the ***Speedster7t GPIO User Guide*** **(UG112)**[5] for further details.

GPIO V_{REF}

The V_{REF} feature provides a flexible way to generate a reference voltage for logic level thresholds. V_{REF} can be internally generated or externally provided through V_{REF} pin. Further, 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*** **(UG112)**[6]) 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*** **(UG112)**[7] 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

5 https://www.achronix.com/documentation/speedster7t-gpio-user-guide-ug112
6 https://www.achronix.com/documentation/speedster7t-gpio-user-guide-ug112
7 https://www.achronix.com/documentation/speedster7t-gpio-user-guide-ug112

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)[8] for further details.

Miscellaneous

FTDI Interface for FPGA Programming

Achronix ACE software can program the Speedster FPGAs via a USB cable. To facilitate this option, 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)[9] and the *Bitstream Programming and Debug Interface User Guide* (UG004)[10] for details on FTDI and JTAG programming.

Table 18 • 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 19 • AC7t1500 JTAG Pins

Signal	Pin	Direction	Pull-up/Pull-down
JTAG_TCK	AR27	Input	
JTAG_TDI	AW25	Input	
JTAG_TDO	AU27	OUT	

8 https://www.achronix.com/documentation/speedster7t-clock-and-reset-architecture-user-guide-ug083
9 https://www.achronix.com/documentation/speedster7t-configuration-user-guide-ug094
10 https://www.achronix.com/documentation/bitstream-programming-and-debug-interface-user-guide-ug004

Signal	Pin	Direction	Pull-up/Pull-down
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 is needed, even at "low speeds", to 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.

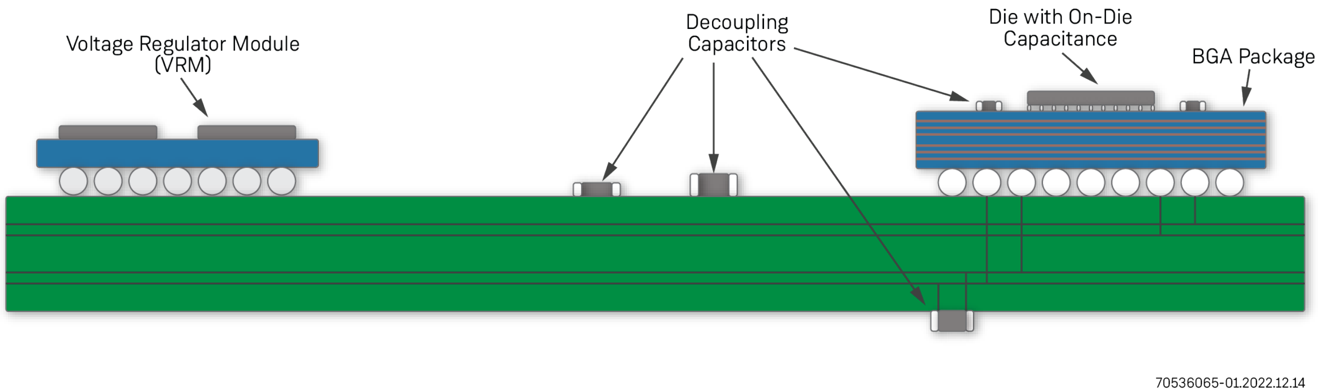

Figure 47 • A Typical Power Distribution Network

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.

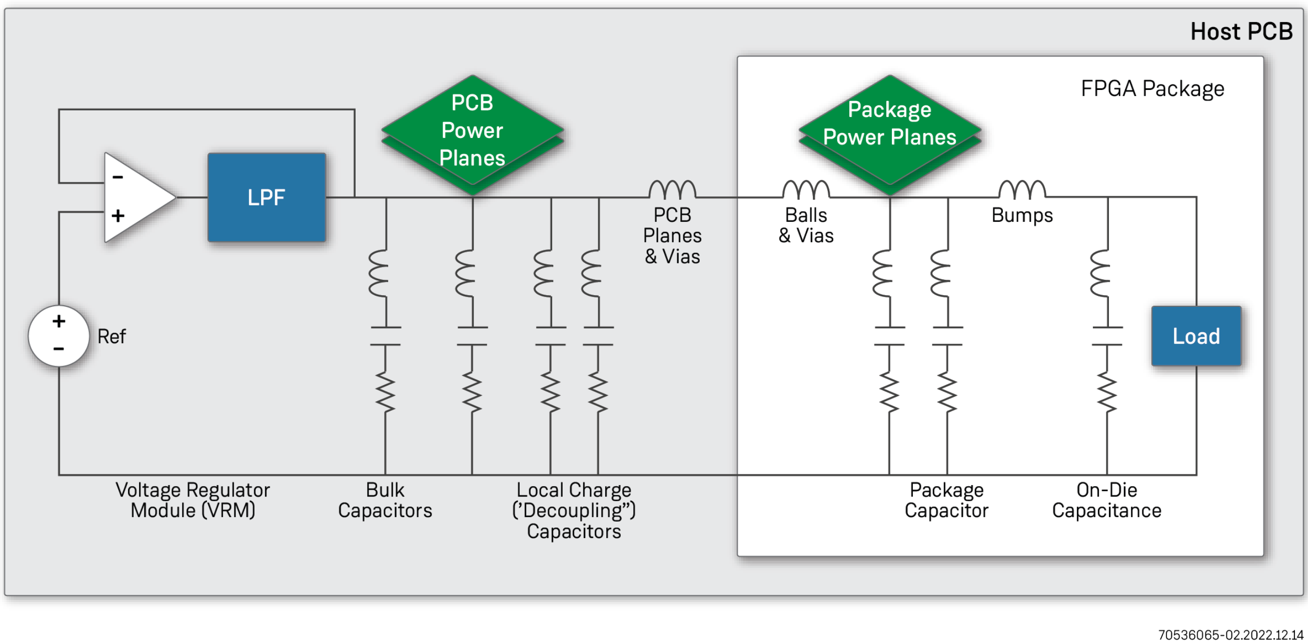

Figure 48 • Typical Power Distribution Network Equivalent Circuit

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 specification 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, the regulator becomes less efficient and can even become less stable or have its effective operating frequency range limited. In the case of the $COREV_{DD}$, a reasonable voltage drop might be 200 mV. This value 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.

For SerDes voltage supplies, for high-speed operation (data rate of 56 Gbps per lane and above), it is recommended to use low drop-out regulators as they have much better ripple noise performance as compared to switching regulators.

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 FPGA 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.

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

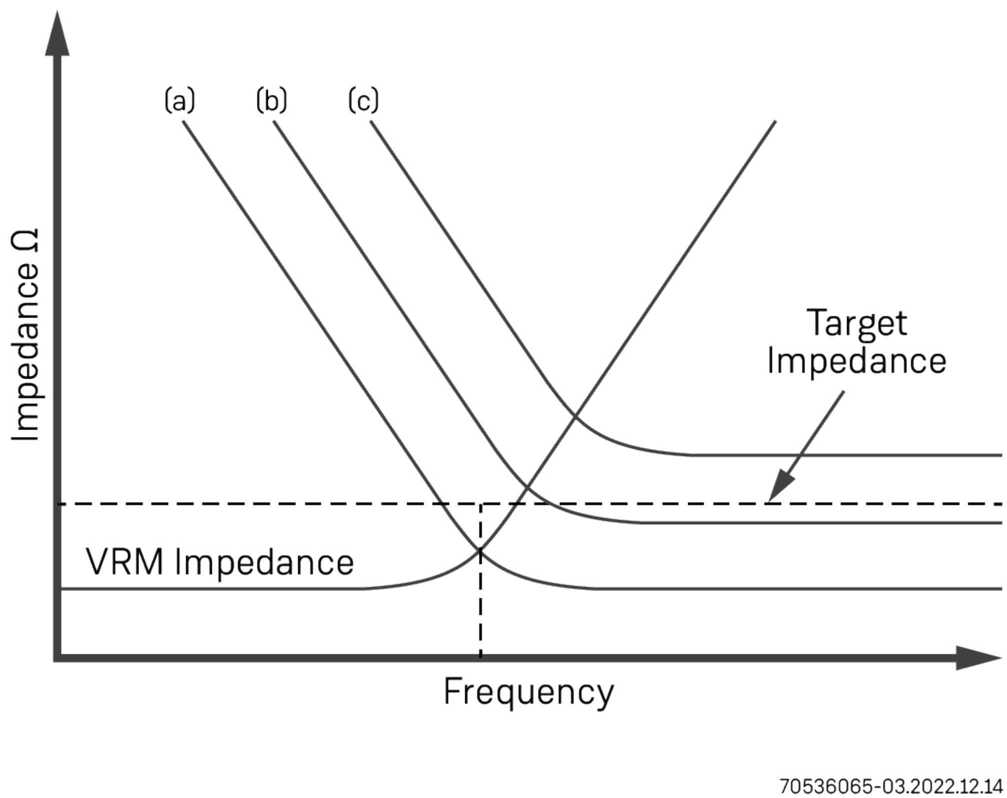

Figure 49 • Selecting Bulk Capacitance Values

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** (page 71). Alternatively, Achronix provides impedance targets for each power supply. Refer to the ***Speedster7t Power User Guide* (UG087)**[11] 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.
- The decoupling capacitors should be placed as close as possible to the FPGA. Refer to **PDN Layout Guidelines** (page 75) 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 plus spreading resistance of the decoupling capacitor, C.

This effect can be seen in the **figure** (see figure 49) 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).

11 https://www.achronix.com/documentation/speedster7t-power-user-guide-ug087

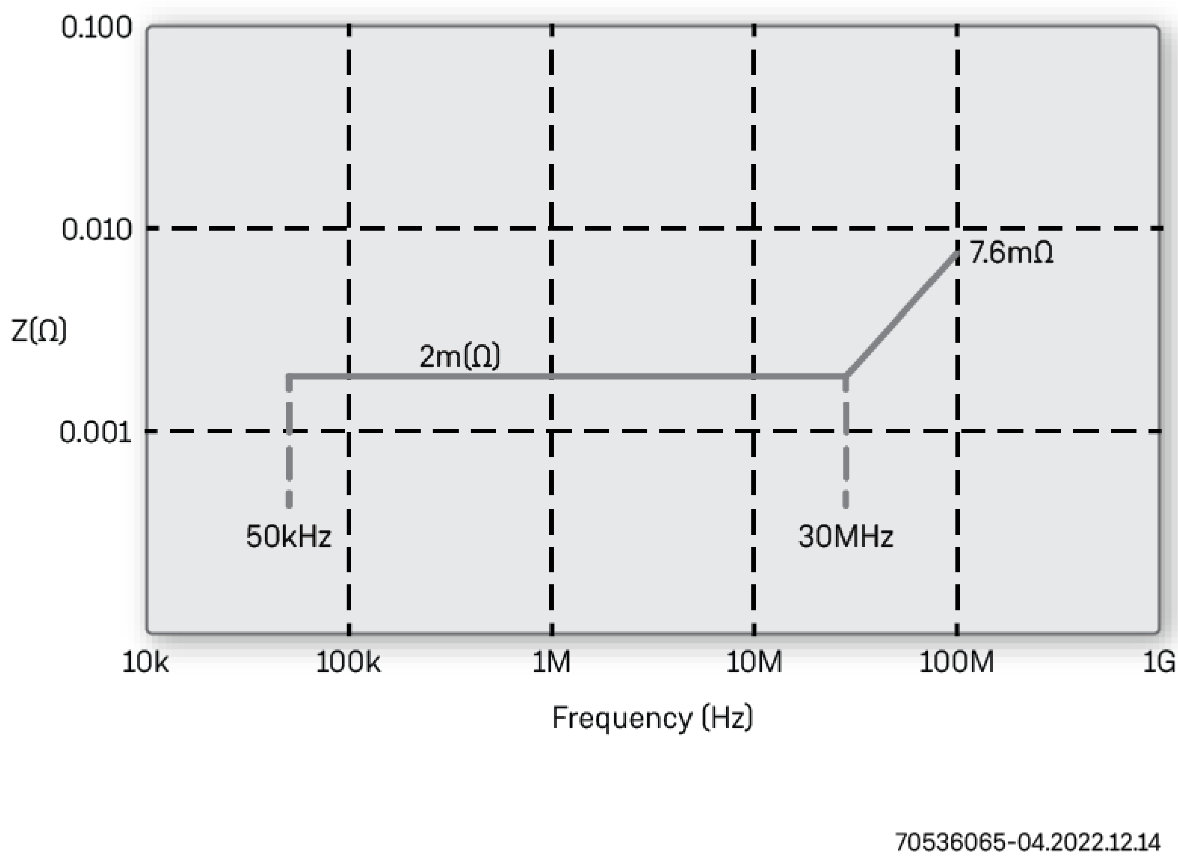

Figure 50 • Impedance Envelope of COREVDD Supply

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 FPGA, 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.

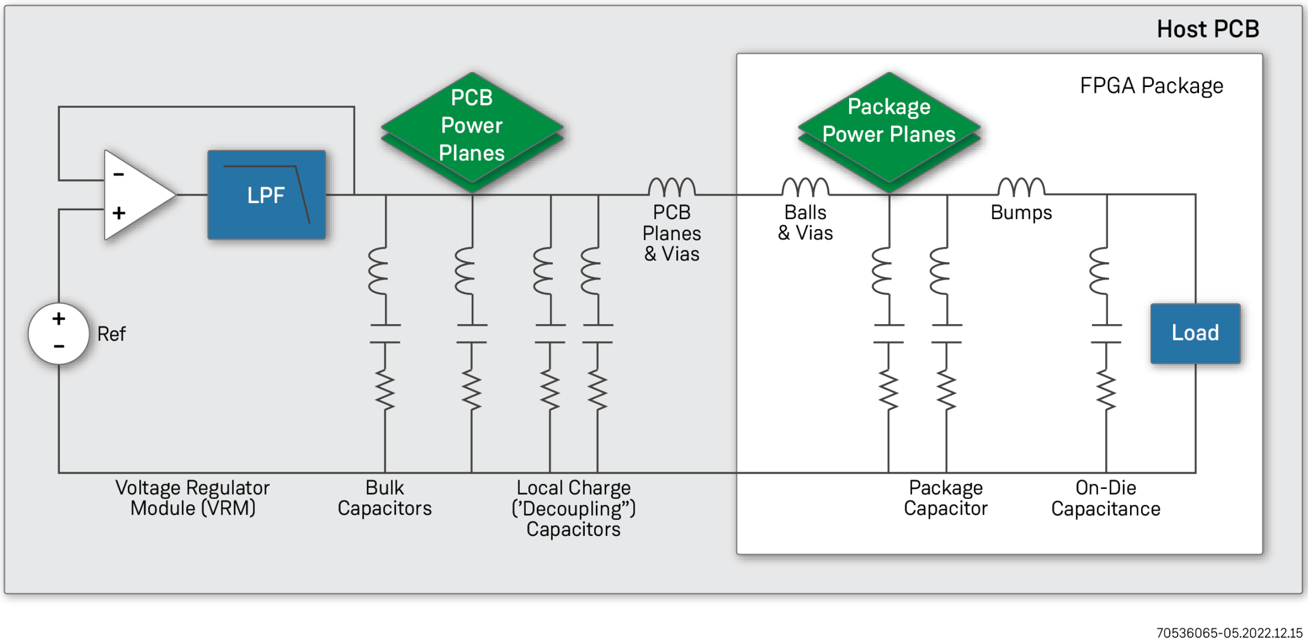

Figure 51 • System-Level Transient Model

> **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 gaps (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, plus 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 FPGA 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 FPGA 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 FPGA 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.

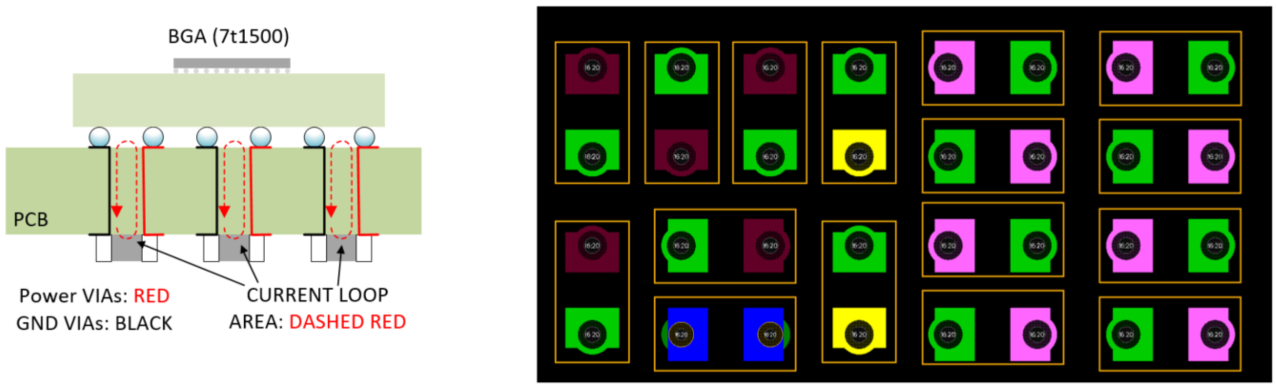

Figure 52 • Placement of Decoupling Capacitors Using Via-in-Pad

- The connection from voltage regulator to the 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.

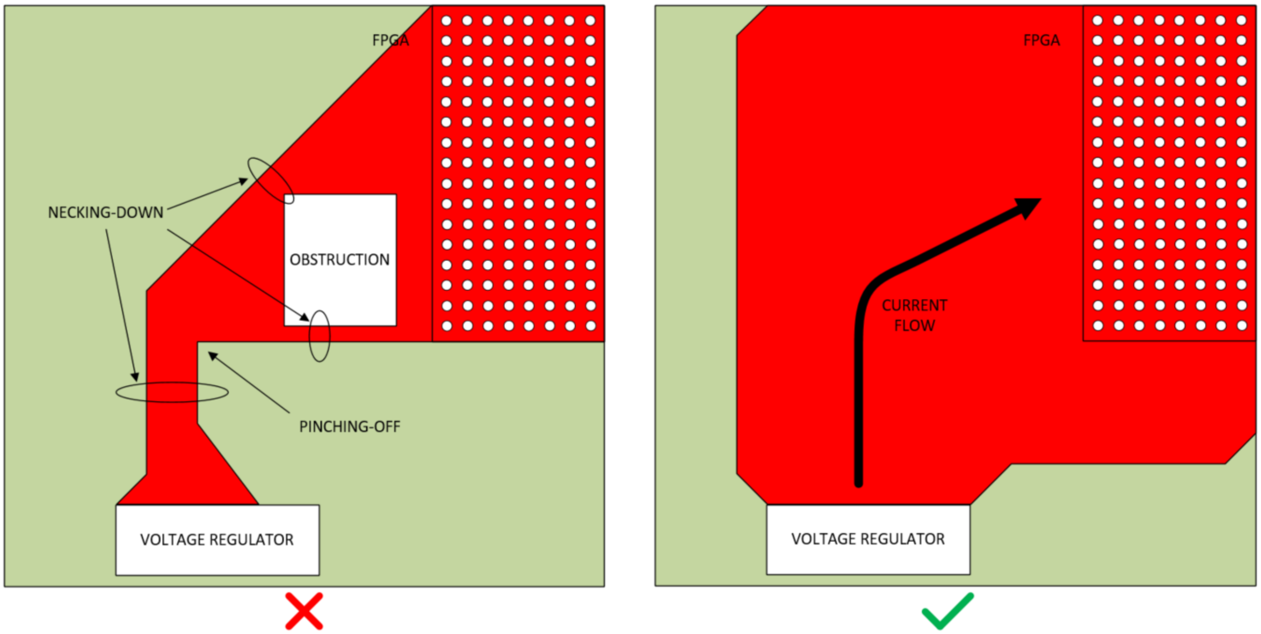

Figure 53 • Routing Power Planes from the VRM to the BGA

- 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

Speedster7t FPGAs have 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 20 • 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

Table 21 • AC7t800 FPGA PLL Power Pins

Supply	Pin	VSS	Pin
ENOC_0_PLL_VDDA	AH18	ENOC_0_PLL_VSSA	AG17

Supply	Pin	VSS	Pin
ENOC_1_PLL_VDDA	AC24	ENOC_1_PLL_VSSA	AC23
ENOC_2_PLL_VDDA	AB24	ENOC_2_PLL_VSSA	AB23
ENOC_3_PLL_VDDA	Y24	ENOC_3_PLL_VSSA	Y23
ENOC_4_PLL_VDDA	P18	ENOC_4_PLL_VSSA	R19
ENOC_5_PLL_VDDA	T10	ENOC_5_PLL_VSSA	R11
GCG_0_PLL_VDDA	AE21	GCG_0_PLL_VSSA	AF22
GCG_1_PLL_VDDA	V24	GCG_1_PLL_VSSDA	V23

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.

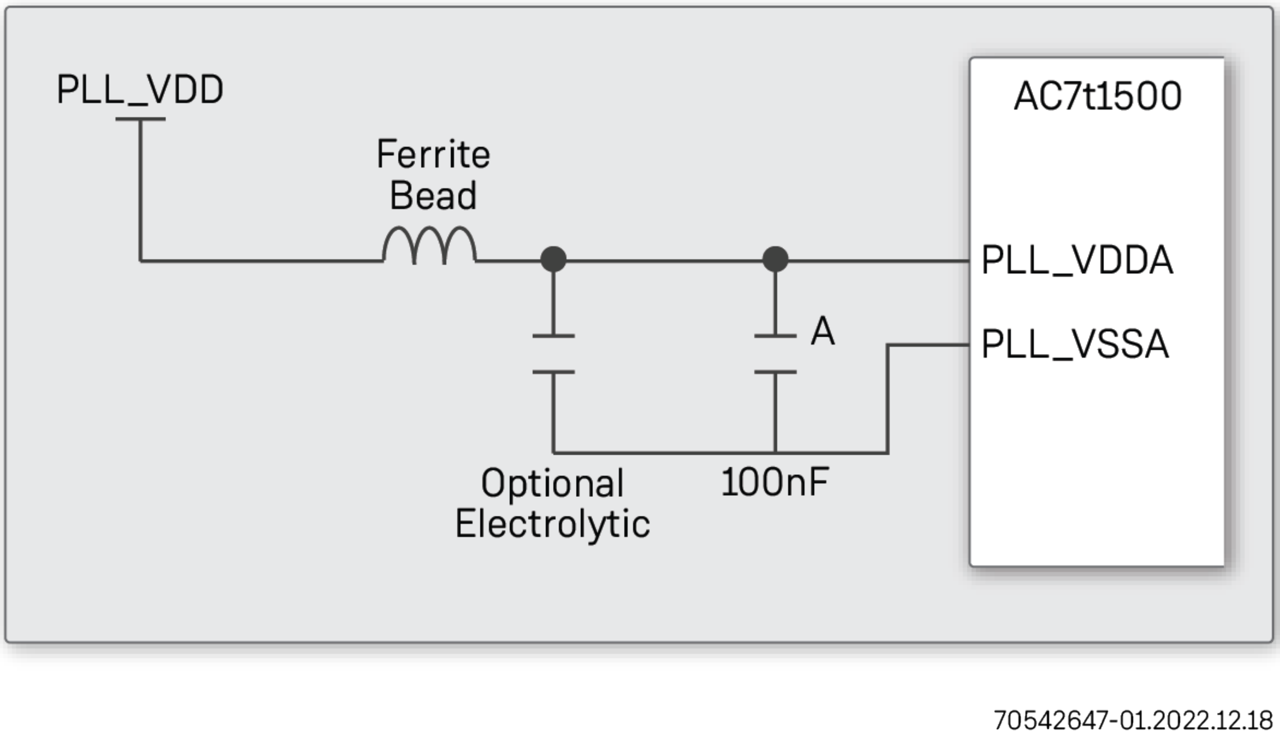

Figure 54 • PLL Power Filter

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 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 V_{DDA} and V_{SSA} 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).

Chapter 10 : Speedster7t PCB System Level Considerations

This chapter covers key considerations for system-level design, focusing on what external components are needed, how they are connected, and the additional balls required on the Speedster7t FPGA. The block diagram below illustrates one example of implementation.

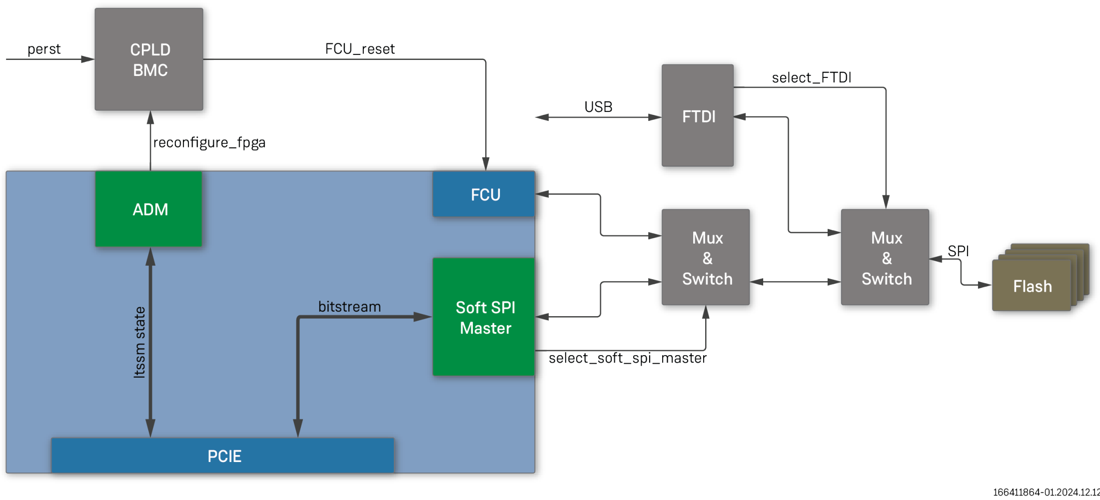

Figure 55 • Block Diagram

See also ***Speedster7t Configuration User Guide*** (UG094)[12] for additional details.

PCIe PERST

Background

Speedster7t FPGAs, when used as a PCIe endpoint, can not perform a PCIe reset upon PERST without also resetting the whole FPGA. In other words, if the PCIe interface needs to be reset upon PERST, the FPGA must be completely reconfigured. This situation requires additional logic on the PCB to trigger a reconfiguration using the bitstream stored in an on-board flash device.

Recommendation

An external component is needed to monitor PERST and trigger a reconfiguration. Such reconfiguration can occur in two different ways. In the first method, the external component toggles FCU reset, which causes the FCU of the FPGA to act as the SPI initiator to read the bitstream from flash memory.

12 https://www.achronix.com/documentation/speedster7t-configuration-user-guide-ug094

In the second method, the external component acts as the SPI initiator to read from flash memory and configures the FPGA over the FCU CPU interface

PCIe Hot Reset

Background

As in the PERST situation described above, when the PCIe interface of the Speedster7t FPGA is reset upon a hot reset, the whole FPGA must also be reconfigured.

Recommendation

The Achronix Device Manger (ADM) can be optionally enabled to monitor the PCIe link status and output a signal to request an external device to initiate the mechanism to reconfigure the FPGA. In the RTL, this "reconfiguration request" output of the ADM is then used to drive the output port of the top-level RTL. As a result, an additional ball needs to be reserved for this "reconfiguration request" signal

Writing Bitstream to the Flash Using FTDI

Background

The FCU can be used effectively as an SPI initiator to read the bitstream from flash memory. A common usage is to configure the FPGA from the image stored in flash memory on power-up using the FCU. However, using it as the SPI initiator to write a bitstream into flash memory is *not* supported. As a result, in order to store a default bitstream into the flash memory before deployment in the field, an alternate method must be employed.

Recommendation

An FTDI device that performs USB-to-SPI conversion is commonly used. This requirement does not affect the available ball count of the FPGA.

Writing Bitstream to the Flash over PCIe

Background

As previously described, using the FCU as an SPI initiator to write is not supported. In the field, a USB connection to write a new bitstream is usually not available. Therefore, writing to flash memory must happen over PCIe.

Recommendation

One approach is to implement a soft SPI controller inside the FPGA fabric. In one method, every design implementation of the FPGA can contain this SPI controller so that the flash memory can be written to at anytime. In another method, a special design implementation with the sole purpose of facilitating writing a new bitstream into flash memory can be loaded only as needed. As a results, additional balls need to be reserved for the SPI interface.

Multiplexing the Various SPI Initiators to the Flash Memory

Background

When additional SPI initiators need to access the flash memory, a way to properly select between them is required. There are typically three SPI initiators to consider.

- FTDI
- SPI controller inside the FPGA (for writing to the flash memory over PCIe)
- FCU (if the FCU is used to read the the flash memory for configuration)

Recommendation

External multiplexers and switches can be used to select the SPI Initiator and route the read data. A package ball needs to be reserved for selecting the soft SPI initiator and FCU. The default states of these control lines upon power-up must be carefully considered.

Chapter 11 : Revision History

Version	Date	Description
1.0	09 Nov 2021	• Initial Achronix Release.
2.0	20 Apr 2026	• Changed document title from *Speedster7t AC7t1500 Board Designers Guide*. • General updates and corrections. **Updated** the following chapters for the addition of support for 7t700, 7t800, and 7t1400 devices: • Chapter, "PCB Pin Mapping on the SerDes Interface (page 22)" • Chapter, "PCB DDR4 Interface (page 45)" **Added**: • Chapter, "PCB PLL Power Filtering (page 78)" • Chapter, "Speedster7t PCB System Level Considerations (page 81)"