

UltraScale Architecture PCB Design

User Guide

UG583 (v1.29) December 23, 2025



Table of Contents

Chapter 1: Power Distribution System in UltraScale Devices.....	6
Introduction to the UltraScale Architecture.....	6
Introduction.....	7
PCB Decoupling Capacitors.....	8
V _{CC_PSDDR_PLL} Supply.....	8
Video Codec Unit (MPSoC EV Devices Only).....	10
Transceiver PCB Routing Guidelines.....	11
Power Supply Consolidation Solutions for Zynq UltraScale+ MPSoCs.....	12
Chapter 2: PCB Guidelines for Memory Interfaces in Kintex	
UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices....	23
Overview.....	23
Reference Stackup.....	23
General Memory Routing Guidelines.....	26
Adjusting for Different Stack-Ups.....	39
PCB Guidelines for DDR4 SDRAM (PL and PS).....	42
PCB Guidelines for DDR3/3L SDRAM (PL and PS).....	54
DDR3/DDR4 UDIMM/RDIMM/SODIMM/LRDIMM Routing Guidelines (PL and PS).....	60
PCB Guidelines for LPDDR4 Memories without ECC (PS).....	69
PCB Guidelines for LPDDR4 Memories with ECC (PS).....	78
PCB Routing Guidelines for LPDDR4 Memories in High-Density Interconnect Boards...	92
PCB Guidelines for LPDDR3 SDRAM (PL and PS).....	96
PCB Guidelines for RLDRAM 3 Memory.....	106
PCB Guidelines for QDR II+ SRAM.....	117
PCB Guidelines for QDR-IV SRAM.....	126
Chapter 3: PCB Guidelines for Memory Interfaces in Spartan	
UltraScale+ Devices.....	131
Overview.....	131
Required Memory Routing Guidelines for All Interfaces.....	131
PCB Guidelines for DDR4 SDRAM.....	144

PCB Routing Guidelines for LPDDR5 Interfaces.....	155
PCB Routing Guidelines for LPDDR4x Interfaces.....	160
Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS.....	168
PDN Guidelines for RFSoc Digital Power Rails.....	168
ADC and DAC PCB Guidelines.....	170
Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC.....	205
Boot Mode.....	205
CAN.....	205
DisplayPort.....	205
eMMC.....	206
Ethernet MAC RGMII.....	206
Ethernet MAC SGMII.....	206
I2C.....	206
JTAG.....	207
NAND Memory Controller.....	207
PCIe.....	207
PS_INIT_B, PS_PROG_B, and PS_DONE.....	208
PS Reference Clock.....	208
PS Reset (External System Reset and POR Reset).....	208
QSPI.....	209
Real-Time Clock.....	209
SATA.....	209
SD/SDIO.....	209
SPI.....	210
Trace Port Interface Unit.....	210
Triple Time Counter.....	210
UART.....	211
USB 2.0.....	211
USB 3.0.....	211
Watchdog Timer.....	211
PS-GTR Transceiver Interfaces.....	211
AC Coupled Reference Clock.....	217
Unused Reference Clocks.....	217
Reference Clock Power.....	218
Power Supply Noise.....	218

PCB Design Checklist.....	218
Chapter 6: PCB Routing Guidelines for MIPI D-PHY.....	220
Chapter 7: Migration between UltraScale Devices and Packages..	222
UltraScale Device Migration Checklist.....	222
Chapter 8: Migration between UltraScale and UltraScale+ FPGAs.	243
UltraScale+ FPGA Migration Checklist.....	243
Chapter 9: Migration between Zynq UltraScale+ MPSoCs and Packages.....	258
Zynq UltraScale+ MPSoC Migration Checklist.....	258
Chapter 10: Migration between Virtex UltraScale+ VU13P Devices and VU27P/VU29P Devices.....	270
Power Supply Differences.....	270
Package Height and Thermal Design Differences.....	274
Package Flight Time Differences.....	275
Bonded/Unbonded Bank/Quad Differences.....	276
CMAC/PCIe Location Differences.....	277
GTY/GTM Channel and Clock Mapping Differences.....	277
100G Ethernet Layout Guidelines.....	279
Chapter 11: SelectIO Signaling.....	281
Interface Types.....	281
Single-Ended Signaling.....	282
Chapter 12: PCB Technology Basics.....	294
Introduction.....	294
PCB Structures.....	294
Basic PDS Principles.....	297
PCB Materials and Traces.....	313
Simulation Methods.....	319
PDS Measurements.....	320
Troubleshooting.....	326
Chapter 13: Design of Transitions for High-Speed Signals.....	328
Excess Capacitance and Inductance.....	328
Time Domain Reflectometry.....	328

BGA Package.....	330
SMT Pads.....	331
Differential Vias.....	336
P/N Crossover Vias.....	339
SMA Connectors.....	339
Backplane Connectors.....	339
Microstrip/Stripline Bends.....	340
Appendix A: Maximum Current Draw for V_{CCINT} in UltraScale+ Devices.....	344
Appendix B: Memory Derating Tables.....	353
Appendix C: Material Properties and Insertion Losses.....	361
Appendix D: Additional Resources and Legal Notices.....	364
Finding Additional Documentation.....	364
Support Resources.....	365
References.....	365
Revision History.....	366
Please Read: Important Legal Notices.....	380

Power Distribution System in UltraScale Devices

Introduction to the UltraScale Architecture

The AMD UltraScale™ architecture is the first ASIC-class architecture to enable multi-hundred gigabit-per-second levels of system performance with smart processing, while efficiently routing and processing data on-chip. UltraScale architecture-based devices address a vast spectrum of high-bandwidth, high-utilization system requirements by using industry-leading technical innovations, including next-generation routing, ASIC-like clocking, 3D ICs, multiprocessor SoC (MPSoC) technologies, and new power reduction features. The devices share many building blocks, providing scalability across process nodes and product families to leverage system-level investment across platforms.

AMD Spartan™ UltraScale+™ devices provide high I/O to logic ratio, integrated memory controllers, and advanced I/O that help Industrial, Vision and Healthcare (IVH), Audio, Video and Broadcast (AVB), Automotive, and other FPGA application developers looking to build cost-optimized solutions. Available in a wide array of packaging options, this family delivers a balance of cost, power, performance, and size.

AMD Artix™ UltraScale+™ devices provide high serial bandwidth and signal compute density in a cost-optimized device for critical networking applications, vision and video processing, and secured connectivity. Coupled with the innovative InFO packaging, which provides excellent thermal and power distribution, Artix UltraScale+ FPGAs are perfectly suited to applications requiring high compute density in a small footprint.

AMD Kintex™ UltraScale+™ devices provide the best price/performance/watt balance in a 16 nm FinFET node, delivering the most cost-effective solution for high-end capabilities, including transceiver and memory interface line rates as well as 100G connectivity cores. Our newest mid-range family is ideal for both packet processing and DSP-intensive functions and is well suited for applications including wireless MIMO technology, Nx100G networking, and data center.

AMD Kintex™ UltraScale™ devices provide the best price/performance/watt at 20 nm and include the highest signal processing bandwidth in a mid-range device, next-generation transceivers, and low-cost packaging for an optimum blend of capability and cost-effectiveness. The family is ideal for packet processing in 100G networking and data centers applications as well as DSP-intensive processing needed in next-generation medical imaging, 8k4k video, and heterogeneous wireless infrastructure.

AMD Virtex™ UltraScale+™ devices provide the highest performance and integration capabilities in a 16 nm FinFET node, including both the highest serial I/O and signal processing bandwidth, as well as the highest on-chip memory density. As the industry's most capable FPGA family, the Virtex UltraScale+ devices are ideal for applications including 1+Tb/s networking and data center and fully integrated radar/early-warning systems.

Virtex UltraScale devices provide the greatest performance and integration at 20 nm, including serial I/O bandwidth and logic capacity. As the industry's only high-end FPGA at the 20 nm process node, this family is ideal for applications including 400G networking, large scale ASIC prototyping, and emulation.

AMD Zynq™ UltraScale+™ devices provide 64-bit processor scalability while combining real-time control with soft and hard engines for graphics, video, waveform, and packet processing. Integrating an Arm®-based system for advanced analytics and on-chip programmable logic for task acceleration creates unlimited possibilities for applications including 5G Wireless, next generation ADAS, and industrial Internet-of-Things.

The UltraScale architecture documentation suite is available at docs.amd.com.

Introduction

This chapter documents the power distribution system (PDS) for UltraScale devices, including decoupling capacitor selection, placement, and PCB geometries. A simple decoupling method is provided for each device. Basic PDS design principles are covered, as well as simulation and analysis methods. This chapter contains the following sections:

- [PCB Decoupling Capacitors](#)
- [V_{CC_PSDDR_PLL} Supply](#)
- [Video Codec Unit \(MPSoC EV Devices Only\)](#)
- [Transceiver PCB Routing Guidelines](#)
- [Power Supply Consolidation Solutions for Zynq UltraScale+ MPSoCs](#)

PCB Decoupling Capacitors


The recommended decoupling capacitor quantities for UltraScale+ devices can be found in the Power Design Manager (PDM) tool (download at www.amd.com/power). The capacitor quantities provided in the PDM tool are derived specifically for each design based on actual device utilization. It is for this reason that the decoupling tables in this user guide have been removed. However, the decoupling tables have been moved to [Answer Record 39033](#) to provide continuity as customers move to the PDM tool.

$V_{CC_PSDDR_PLL}$ Supply

$V_{CC_PSDDR_PLL}$ is a 1.8V nominal supply that provides power to the PLL used for the PS DDR controller. It can be powered separately or derived from the V_{CC_PSAUX} supply. If powered by V_{CC_PSAUX} , $V_{CC_PSDDR_PLL}$ must be filtered through a 120Ω @ 100 MHz, size 0603 ferrite bead and a $10\ \mu\text{F}$ or larger, size 0603 decoupling capacitor. In both cases, a $1.0\ \mu\text{F}$ 0201 or $10\ \mu\text{F}$ 0402 capacitor must be placed near the $V_{CC_PSDDR_PLL}$ BGA via.

The PCB construction of the $V_{CC_PSDDR_PLL}$ power supply must be carefully managed. The recommended connection between the 0603 capacitor and the $V_{CC_PSDDR_PLL}$ BGA ball is a planelet with a minimum width of 80 mil (2 mm) and a length of less than 3,000 mil (76 mm). If a planelet cannot be used, a trace with a maximum impedance of 40Ω and a length of less than 2,000 mil (50.8 mm) must be used. The 0201 or 0402 capacitor should be placed as close as possible to the FPGA, along with the shortest possible trace length. The following figure shows an example of the filtering and local capacitor circuit used when $V_{CC_PSDDR_PLL}$ is derived from V_{CC_PSAUX} .


Figure 1: Connecting VCC_PSDDR_PLL



UG583_c5_25_060619

The following figure shows an example of the layout of the same filtering circuit.

Figure 2: Filtering Circuit Layout



UG583_c8_26_031616

The recommended components are:

- Ferrite bead: Murata BLM18SG121TN1
- 10 µF (or larger) capacitor: Murata GRM188C80E476ME05
- 1.0 µF 0201 or 10 µF 0402 capacitor: Murata GRM155C80J106ME11D

Video Codec Unit (MPSoC EV Devices Only)

V_{CCINT_VCU} is a 0.90V rail that provides power to the video codec unit (VCU) within Zynq UltraScale+ MPSoC EV devices. V_{CCINT_VCU} must be powered as a standalone power rail and cannot be combined with any other power rails.

If the VCU is not going to be used, the V_{CCINT_VCU} pins can be grounded to reduce leakage current.

V_{CCINT_VCU} Decoupling Capacitors

Refer to Power Design Manager (PDM) tool (download at www.amd.com/power) for V_{CCINT_VCU} decoupling capacitor recommendations.


V_{CCINT_VCU} Plane Design and Power Delivery

The V_{CCINT_VCU} plane(s) or planelet(s) should be wide enough to accommodate up to the required amount of current (per PDM) at their narrowest points, factoring in all width reductions due to vias and other keepouts. However, planes and planelets need to be drawn even wider to reduce resistive losses between the regulator and the BGA pins. While properly routed sense lines offer assistance in offsetting resistive losses, regulator efficiency is maximized when resistive loss is minimized. For these reasons, AMD recommends drawing V_{CCINT_VCU} planes at 500–800 mil width while outside of the BGA region, and 400 mils under the BGA before the plane gradually narrows to the V_{CCINT_VCU} pins. The following figure shows an example V_{CCINT_VCU} plane layout.

Note: If there are an unusually large number of vias or other keepouts within the plane, a wider plane might need to be drawn to reduce resistive losses and to maintain the ability to provide at least 3A of current.

In addition, AMD recommends placing a ground plane directly above or below the plane that contains V_{CCINT_VCU} to reduce inductance.

Figure 3: Example Plane Layout for V_{CCINT_VCU}



X20168-020618

Transceiver PCB Routing Guidelines

Refer to the *UltraScale Architecture GTH Transceivers User Guide* ([UG576](#)) or the *UltraScale Architecture GTY Transceivers User Guide* ([UG578](#)).

Power Supply Consolidation Solutions for Zynq UltraScale+ MPSoCs

Zynq UltraScale+ MPSoCs contain a number of power rails that correspond to the variety of features contained within the devices. A complete list of the power rails for Zynq UltraScale+ MPSoCs is listed in *Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics* ([DS925](#)).

For most purposes, assuming the decoupling requirements and filtering requirements are met, many of the power rails can be consolidated to reduce the total number of power regulators required to power a Zynq UltraScale+ MPSoC. In many applications, only five power regulators are required to power a Zynq UltraScale+ MPSoC. (Some of the five power regulators might be shared with other devices in the system.) However, the amount of consolidation that is possible depends on user requirements.

For example, many applications periodically need to enter ultra-low power, reduced functionality states to extend battery life or to reduce their overall power consumption. To facilitate this, Zynq UltraScale+ MPSoCs have four independent power domains that can be individually isolated: low-power domain (LPD), full-power domain (FPD), PL power domain (PLPD), and battery power domain (BPD). Applications requiring access to these ultra-low power states that cannot be achieved via power gating options in software might need independent control of the power rails associated with each power domain to enable the power associated with unused domains to reduce to zero. This results in the need for more power regulators or for the use of load switches. In this document, this use case is referred to as full power management flexibility. More detail on Zynq UltraScale+ MPSoC power domains can be found in *Zynq UltraScale+ Device Technical Reference Manual* ([UG1085](#)) and *Managing Power and Performance with the Zynq UltraScale+ MPSoC* ([WP482](#)). The benefits of leveraging the power domains and having full power management flexibility can be quantified using the Power Design Manager (PDM) tool (download at www.amd.com/power).

Other applications nearly always operate at full performance or can meet their power targets by disabling functionality via IP power-gating options (for example, disabling processor cores or clock gating portions of the programmable logic). These applications do not need individual control of the power rails associated with each power domain. Therefore, these applications can achieve the maximum amount of power supply consolidation. For simplicity in this user guide, these applications are referred to as *always on*, and are further classified into three unique use cases. Together with the first use case (full power management flexibility), a total of four key use cases are discussed in this section:

- Always on: Cost optimized (-1 and -2 devices)
- Always on: Power/efficiency optimized (-1L and -2L devices)
- Always on: PL performance optimized (-3 devices)

- Full power management flexibility (all devices/speed grades)

The remainder of this section illustrates the power supply consolidation that can be achieved for the above use cases. In addition, [Power Management Partners](#) provides links to power delivery solutions from a range of AMD power delivery partners for the use cases mentioned above.

Power Supply Consolidation Solutions

Always On: Optimized for Cost (-1 and -2 Devices)

In always on, cost-optimized applications, where standard -1 or -2 speed grade devices are used, a significant amount of power rail consolidation is possible, assuming the decoupling and filtering requirements on the individual supplies as outlined in this document are met. The following table broadly defines the power rail consolidation that is possible for this use case. For further clarity, the following figure shows the possible power rail consolidation graphically.


As shown in the following table, the minimum number of regulators required to power a Zynq UltraScale+ MPSoC in this use case is five. In many cases, five or more regulators can be provided by a single integrated power management integrated circuit (PMIC). For example power delivery solutions for AMD products, refer to [Power Management Partners](#).

Table 1: Always On: Cost-Optimized Power Rail Consolidation (-1 and -2 Devices)

	Power Regulator	Sequence	Possible Power Rail Consolidation
Required	1	See <i>Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)</i> .	VCCINT, VCC_PSINTFP, VCC_PSINTLP, VCC_PSINTFP_DDR, VCCINT_IO, VCCBRAM
	2		VCC_PSAUX, VCC_PSADC ¹ , VCC_PSDDR_PLL ² , VCCAUX, VCCAUX_IO, and VCCADC ¹
	3		VCC_PSPLL, VMGTAVTT (GTH), and VMGTAVTT (GTY)
	4		VCCO_PSDDR
	5		VCCO_PSIO[0:3] assuming all PS I/Os run from same voltage
Required: EV -1 and -2	9		VCCINT_VCU
User-defined	6		VPS_MGTRAVCC
	7		VPS_MGTRAVTT, VMGTVCCAUX (GTH), and VMGTVCCAUX (GTY)
	8		VMGTAVCC (GTH) and VMGTAVCC (GTY)
	10		Optional PL and PS I/O voltages

Notes:

1. Assuming rail is filtered as per *UltraScale Architecture System Monitor User Guide (UG580)*.
2. Assuming rail is filtered as per [VCC_PSDDR_PLL Supply](#).
3. When consolidating, ensure that all tolerances are met.

Figure 4: Always On: Cost-Optimized Power Rail Consolidation


Note: In the previous figure, the dashed lines are dependent on the user configuration. Identical voltages between I/O supplies and digital supplies can be combined if operating at the same voltage. This might alter the power sequence, moving the I/O supply up to the sequence slot of the rail it is combined with.

Depending on user configuration, more than five regulators might be required. The following table details the number of regulators required to power Zynq UltraScale+ MPSoCs for a variety of device configurations. The following figure shows the power rail consolidation for the configuration requiring only five regulators using CG or EG devices.

Additional power regulators are typically required for termination voltages and reference voltages of the DDR memory associated to the PS or PL.


Table 2: Number of Rails for “Always On: Optimized for Cost (-1 and -2 Devices)” Scenario

Configuration	Devices (Speed Grades)	Number of Power Regulators
PS and PL DDR, 1.8V I/O, and 2.5V I/O or 3.3V I/O with no MGTs	CG (-1, -2)	Five (1, 2, 3, 4, 5)
	EG (-1, -2)	
	EV (-1, -2)	Six (1, 2, 3, 4, 5, 9)

Table 2: Number of Rails for “Always On: Optimized for Cost (-1 and -2 Devices)” Scenario (cont'd)

Configuration	Devices (Speed Grades)	Number of Power Regulators
PS and PL DDR, 1.8V I/O, 2.5V or 3.3V I/O and PS MGTs	CG (-1, -2) EG (-1, -2)	Seven (1, 2, 3, 4, 5, 6, 7)
	EV (-1, -2)	Eight (1, 2, 3, 4, 5, 6, 7, 9)
PS DDR, 1.8V I/O, 3.3V I/O, PS MGTs and PL MGTs	CG (-1, -2) EG (-1, -2)	Eight (1, 2, 3, 4, 5, 6, 7, 8)
	EV (-1, -2)	Nine (1, 2, 3, 4, 5, 6, 7, 8, 9)

Figure 5: Always On: Cost-Optimized with Only Five Power Regulators (CG or EG Devices)



X18634-120417

Always On: Optimized for Power and/or Efficiency (-1L and -2L Devices)

For applications requiring the lowest power dissipation and/or the highest efficiency, AMD offers the -1L and -2L speed grades. To achieve the highest efficiency and lowest power, the -1L and -2L devices can operate with a VCCINT voltage of 0.72V. To enable VCCINT to run at 0.72V an additional power regulator is required.

Other than the need for an additional power regulator for VCCINT, the possible power rail consolidation is similar to the previous use case. The following table broadly defines what power rail consolidation is possible for this use case, assuming the decoupling and filtering requirements on the individual supplies as outlined in this document are met. For further clarity, the following figure shows the possible power rail consolidation graphically.

As shown in the following table, the minimum number of regulators required for this use case is six.


Table 3: Always On - Power/Efficiency Rail Consolidation for Low-Power Devices

	Power Regulator	Sequence	Possible Power Rail Consolidation
Required	1	See <i>Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)</i> .	VCCINT
	2		VCC_PSINTFP, VCC_PSINTLP, VCC_PSINTFP_DDR, VCCINT_IO, and VCCBRAM
	3		VCC_PSAUX, VCC_PSADC ¹ , VCC_PSDDR_PLL ² , VCCAUX, VCCAUX_IO and VCCADC ¹
	4		VCC_PSPLL, VMGTAVTT (GTH), and VMGTAVTT (GTY)
	5		VCCO_PSDDR
	6		VCCO_PSIO[0:3] assuming all PS I/Os run from same voltage
Required - EV devices	10		VCCINT_VCU
User-defined	7		VPS_MGTRAVCC
	8		VPS_MGTRAVTT, VMGTVCCAUX (GTH), and VMGTVCCAUX (GTY)
	9		VMGTAVCC (GTH) and VMGTAVCC (GTY)
	11		Optional PL and PS I/O voltages

Notes:

1. Assuming rail is filtered as per *UltraScale Architecture System Monitor User Guide (UG580)*.
2. Assuming rail is filtered as per [VCC_PSDDR_PLL Supply](#).

Figure 6: Always On: Power and/or Efficiency-Optimized Power Rail Consolidation for Low-Power Devices



X18637-120417

Note: In the previous figure, the dashed lines are dependent on the user configuration. Identical voltages between I/O supplies and digital supplies can be combined if operating at the same voltage. This might alter the power sequence, moving the I/O supply up to the sequence slot of the rail it is combined with.

The following table shows the number of power regulators for a variety of common configurations of Zynq UltraScale+ MPSoC for this use case.

Table 4: Number of Power Regulators Needed for “Always On: Optimized for Power and/or Efficiency (-1L and -2L Devices)” Scenario

Configuration	Number of Power Regulators
PS and PL DDR, 1.8V I/O and 2.5V or 3.3V I/O with no MGTs	Seven (1, 2, 3, 4, 5, 6, 10)
PS and PL DDR, 1.8V I/O, 2.5V I/O and 3.3V I/O with no MGTs	Eight (1, 2, 3, 4, 5, 6, 10, 11)
PS and PL DDR, 1.8V I/O, 2.5V or 3.3V I/O and PS MGTs	Nine (1, 2, 3, 4, 5, 6, 7, 8, 10)
PS DDR, 1.8V I/O, 3.3V I/O, PS MGTs and PL MGTs	Ten (1, 2, 3, 4, 5, 6, 7, 8, 9, 10)

Always On: Optimized for PL Performance (-3 Devices)

For the highest level of PL performance, AMD offers Zynq UltraScale+ MPSoCs in a -3 speed grade. For the -3 speed grade device, all the core rails (VCCINT, VCCINT_VCU, VCCBRAM, VCCINT_IO, VCC_PSINTLP, VCC_PSINTFP, and VCC_PSINTFP_DDR) are run at a nominal voltage of 0.9V.

The following table broadly defines the power rail consolidation that is possible for this use case, assuming the decoupling and filtering requirements on the individual supplies as outlined in this document are met. For further clarity, the following figure shows the possible power rail consolidation graphically. As shown in the following table, the minimum number of power regulators required to power a Zynq UltraScale+ MPSoC in this use case is five for non-EV devices, and six for EV devices.


Table 5: Always On: PL Performance-Optimized Rail Consolidation

	Power Regulator	Sequence	Possible Power Rail Consolidation
Required	1	See <i>Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)</i> .	VCCINT, VCC_PSINTFP, VCC_PSINTLP, VCC_PSINTFP_DDR, VCCINT_IO, VCCBRAM
	2		VCC_PSAUX, VCC_PSADC ¹ , VCC_PSDDR_PLL ² , VCCAUX, VCCAUX_IO, and VCCADC ¹
	3		VCC_PSPLL, VMGTAVTT (GTH), and VMGTAVTT (GTY)
	4		VCCO_PSDDR
	5		VCCO_PSIO[0:3] assuming all PS I/Os run from same voltage
	10		VCCINT_VCU (if used)
User-defined	6	VPS_MGTRAVTT, VMGTVCCAUX (GTH), and VMGTVCCAUX (GTY)	
	7	VMGTAVCC (GTH), and VMGTAVCC (GTY)	
	8	Additional PL and PS I/O voltages	
	9	VPS_MGTRAVCC	

Notes:

1. Assuming rail is filtered as per *UltraScale Architecture System Monitor User Guide (UG580)*.
2. Assuming rail is filtered as per [VCC_PSDDR_PLL Supply](#).

Figure 7: Always On: PL Performance-Optimized Rail Consolidation



X18636-020618

Note: In the previous figure, the dashed lines are dependent on the user configuration. Identical voltages between I/O supplies and digital supplies can be combined if operating at the same voltage. This might alter the power sequence, moving the I/O supply up to the sequence slot of the rail it is combined with.

The following table shows the number of regulators for common configurations of Zynq UltraScale+ MPSoC used in this use case.

Table 6: Number of Power Regulators Required for “Always On: Optimized for PL Performance (-3 Devices)” Scenario

Configuration	Number of Power Regulators ¹
PS and PL DDR, 1.8V I/O and 2.5V or 3.3V I/O with no MGTs	Five (1, 2, 3, 4, 5)
PS and PL DDR, 1.8V I/O, 2.5V I/O and 3.3V I/O with no MGTs	Six (1, 2, 3, 4, 5, 8)
PS and PL DDR, 1.8V I/O, 2.5V or 3.3V I/O and PS MGTs	Seven (1, 2, 3, 4, 5, 6, 9)
PS DDR, 1.8V I/O, 3.3V I/O, PS MGTs and PL MGTs	Eight (1, 2, 3, 4, 5, 6, 7, 9)

Notes:

1. Add regulator 10 to each scenario if using VCCINT_VCU.

Full Power Management Flexibility (All Speed Grades/Devices)

As mentioned earlier in this section, many applications need to enter ultra-low power, reduced functional states to maximize battery life or reduce overall power consumption. To facilitate this, Zynq UltraScale+ MPSoCs have four independent power domains: LPD, FPD, PLPD, and BPD. To enable the power associated with each of these domains to go to zero when unused, the power rails associated with each domain need to be separated, either via independent power regulators or via load switches.

The following table broadly defines the power rail consolidation that is possible for this use case, assuming the decoupling and filtering requirements on the individual supplies as outlined in this document are met. For further clarity, the following figure shows the possible power rail consolidation graphically. As shown in the following table, the minimum number of power regulators required to power a Zynq UltraScale+ MPSoC in this use case is nine. Alternatively, to achieve the same isolation, a combination of power regulator and load switch could be used, reducing the overall number of power regulators required.


Table 7: Full Power Management Flexibility Rail Consolidation

	Power Regulator	Sequence	Possible Power Rail Consolidation
Required	1	See <i>Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)</i> .	VCC_PSINTLP
	2		VCC_PSAUX and VCC_PSADC ¹
	3		VCC_PSPLL
	4		VCCO_PSIO[0:3] assuming run off same voltage
	5		VCC_PSINTFP and VCC_PSINTFP_DDR
	6		VCC_PSDDR_PLL
	7		VCCO_PSDDR
	8		VCCINT ² , VCCBRAM, and VCCINT_IO
	9		VCCAUX, VCCAUX_IO, VCCADC ¹
Required EV devices	15		VCCINT_VCU
User-defined	10		VPS_MGTRAVCC
	11		VPS_MGTRAVTT
	12		VMGTAVTT (GTH) and VMGTAVTT (GTY)
	13		VMGTAVCC (GTH) and VMGTAVCC (GTY)
	14		VMGTVCCAUX (GTH) and VMGTVCCAUX (GTY)
	16		Optional PL and PS I/O voltages

Notes:

1. Assuming rail is filtered as per *UltraScale Architecture System Monitor User Guide (UG580)*.
2. Additional voltage regulator required for -2LI or -2LE devices where user wants to run VCCINT = 0.72V.

Figure 8: Full Power Domain Flexibility Consolidation



X18638-120417

Note: In the previous figure, the dashed lines are dependent on the user configuration. Identical voltages between I/O supplies and digital supplies can be combined if operating at the same voltage. This might alter the power sequence, moving the I/O supply up to the sequence slot of the rail it is combined with.

The following figure shows the number of regulators for common configurations of Zynq UltraScale+ MPSoC used in this use case. For applications that use -1LI or -2L devices, where the user wants to run VCCINT at 0.72V, an additional voltage regulator is required relative to the following table. An additional regulator is required for VCCINT_VCU when using an EV device.

Table 8: Number of Power Rails for Typical Configurations for Application Requiring Full Power Domain Flexibility

Configuration	Devices (Speed Grades)	Number of Power Regulators
Full power domain control, VCCINT = 0.85, PS DDR, 1.8V PL I/O and 1.8/2.5/3.3V PS I/O with no MGTs	CG (-1, -2) EG (-1, -2)	Nine (1, 2, 3, 4, 5, 6, 7, 8, 9)

Table 8: Number of Power Rails for Typical Configurations for Application Requiring Full Power Domain Flexibility (cont'd)

Configuration	Devices (Speed Grades)	Number of Power Regulators
Full power domain control, VCCINT = 0.9V, PS DDR, 1.8V PL I/O and 1.8/2.5/3.3V PS I/O with no MGTs	EG (-3)	Nine (1, 2, 3, 4, 5, 6, 7, 8, 9)
Full power domain control, VCCINT = 0.9V, PS DDR, 1.8V PL I/O and 1.8/2.5/3.3V PS I/O with no MGTs	EV (-3)	Ten (1, 2, 3, 4, 5, 6, 7, 8, 9, 15)
Full power domain control, VCCINT = 0.85V, PS DDR, 1.8V PL I/O and 1.8/2.5/3.3V PS I/O with no MGTs	CG (-1, -2) EG (-1, -2)	Ten (1, 2, 3, 4, 5, 6, 7, 8, 9, 16)
Full power domain control, VCCINT = 0.9V, PS DDR, 1.8V PL I/O and 1.8/2.5/3.3V PS I/O with no MGTs	EG (-3)	Ten (1, 2, 3, 4, 5, 6, 7, 8, 9, 16)
Full power domain control, VCCINT = 0.9V, PS DDR, 1.8V PL I/O, 1.8/2.5/3.3V PS I/O with no MGTs	EV (-3)	Eleven (1, 2, 3, 4, 5, 6, 7, 8, 9, 15, 16)
Full power domain control, VCCINT = 0.85V, PS DDR, 1.8V PL I/O and 1.8/2.5/3.3V PS I/O with no MGTs	EV (-1, -2)	Eleven (1, 2, 3, 4, 5, 6, 7, 8, 9, 15, 16)
Full power domain control, VCCINT = 0.72V/ programmable, PS DDR, 1.8V PL I/O and 1.8/2.5/3.3V PS I/O with no MGTs	CG (-1L, -2L) EG (-1L, -2L)	Eleven (1, 2, 3, 4, 5, 6, 7, 8, 9, 16, 17)
Full power domain control, VCCINT = 0.72V/ programmable, PS DDR, 1.8V PL I/O, DDR/2.5/3.3V PL I/O and 1.8/2.5/3.3V PS I/O with no MGTs	EV (-1L, -2L)	Twelve (1, 2, 3, 4, 5, 6, 7, 8, 9, 15, 16, 17)
Full power domain control, VCCINT = 0.85V, PS DDR, 1.8V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	CG (-1, -2) EG (-1, -2)	Fourteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14)
Full power domain control, VCCINT = 0.9V, PS DDR, 1.8V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	EG (-3)	Fourteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14)
Full power domain control, VCCINT = 0.9V, PS DDR, 1.8V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	EV (-3)	Fifteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15)
Full power domain control, VCCINT = 0.85V, PS DDR, 1.8V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	EV (-1, -2)	Fifteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15)
Full power domain control, VCCINT = 0.85V, PS DDR, 1.8V PL I/O, DDR/2.5/3.3V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	CG (-1, -2) EG (-1, -2)	Fifteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 16)
Full power domain control, VCCINT = 0.9V, PS DDR, 1.8V PL I/O, DDR/2.5/3.3V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	EG (-3)	Fifteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 16)
Full power domain control, VCCINT = 0.9V, PS DDR, 1.8V PL I/O, DDR/2.5/3.3V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	EV (-3)	Sixteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16)
Full power domain control, VCCINT = 0.85V, PS DDR, 1.8V PL I/O, DDR/2.5/3.3V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	EV (-1, -2)	Sixteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16)
Full power domain control, VCCINT = 0.72/programmable, PS DDR, 1.8V PL I/O, DDR/2.5/3.3V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	CG (-1L, -2L) EG (-1L, -2L)	Sixteen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16)
Full power domain control, VCCINT = 0.72/programmable, PS DDR, 1.8V PL I/O, DDR/2.5/3.3V PL I/O, 1.8/2.5/3.3V PS I/O, PL MGTs and PS MGTs	EV (-1L, -2L)	Seventeen (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17)

PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices

Overview

The AMD UltraScale™ architecture provides solutions for interfacing with the following memory interfaces:

- DDR4, for both programmable logic (PL) and the processing system (PS)
 - DDR3/3L (PL and PS)
 - LPDDR4 (PS)
 - LPDDR3 (PL and PS)
 - RDRAM 3 (PL)
 - QDR-IV (PL)
 - QDR II+ (PL)
-

Reference Stackup

All electrical routing constraints are defined upon the reference stackup (the following table). The actual stackup might be different from this reference stackup. The related constraints such as width and spacing should be adjusted accordingly to meet target impedance and crosstalk in the design guide at reference stackup. For reference, this particular stackup results in an inner signal layer propagation time of 169.5 ps/in.

Table 9: Reference Stackup

Layer	Thickness (mil): Copper/ Core	Description	
		Copper Weight (oz)	Layer
L1	0.6	0.5	Top
	2.9		
L2	0.6	0.5	Ground
	4.5		
L3	0.6	0.5	Signal
	4.5		
L4	0.6	0.5	Ground
	4.5		
L5	0.6	0.5	Signal
	4.5		
L6	1.2	1.0	Power/Ground
	8.0		
L7	1.2	1.0	Power/Ground
	8.0		
L8	1.2	1.0	Power/Ground
	8.0		
L9	1.2	1.0	Power/Ground
	8.0		
L10	1.2	1.0	Power/Ground
	8.0		
L11	1.2	1.0	Power/Ground
	4.5		
L12	0.6	0.5	Signal
	4.5		
L13	0.6	0.5	Ground
	4.5		
L14	0.6	0.5	Signal
	4.5		
L15	0.6	0.5	Ground
	2.9		
L16	0.6	0.5	Bottom


Notes:

1. The material for this reference stackup is Isola High-Tg FR-4, 370HR with Er = 4.0.



IMPORTANT! To achieve the highest memory interface performance, all the high-speed signals are recommended to be routed on the upper signal layers such as L3 and L5, as shown in the following figure, to minimize device pin field via crosstalk impact. Deeper signal layers could be used but signal routing spacing needs to take trade-offs into account with system-level signal integrity simulations.


Figure 9: Layer 3 Routing Example



UG583_c2_11_073014

★ **IMPORTANT!** Routing high-speed signals on lower signal layers comes with more board via coupling jitter depending on board thickness. Signal spacing in the same layer routing needs to be compromised to mitigate deep boards via crosstalk impact.

Figure 10: Layer 14 Routing Example



UG583_c2_12_073014

To determine system timing margins in this design following the AMD memory simulation guidelines, system designers should run system-level memory channel simulations to confirm actual timing margin in customer-specific layout practices.

General Memory Routing Guidelines

1. Include package delay in routing constraints when determining signal trace lengths. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values.
2. UltraScale device breakout specifications assume two signal routes between pads. If routing one signal between pads, traces can be 39Ω instead of 50Ω . To minimize coupling, minimize trace lengths as much as possible when routing two signals between pads.

Figure 11: UltraScale Device Breakout Single-Ended Routing Options



UG583_c2_50_120114


3. Signal-to-signal skew constraints in this chapter are presented in the form *X to Y*, where the *Y* signal is the reference point. Within the specified constraint, signal *X* can be shorter or longer than signal *Y*. If signal *X* is part of a bus, then the shortest and longest signals in the bus must be within the listed specification. If the *Y* signal is a differential clock or signal group, then *Y* is defined as the mid-point between shortest and longest signals in the pair/group.

The following table shows an example *DQ to DQS* signal-to-signal constraint specification. The following figure illustrates an interpretation of the constraint. The shortest *DQ* in the data bus can be no more than 5 ps ahead of the *DQS* signal (mid-point of *DQS_P* and *DQS_N*), while the longest *DQ* in the bus can be no more than 5 ps after the *DQS* signal.

Table 10: Example Signal to Signal Skew Constraint

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
DQ to DQS	±5	±29

Figure 12: Skew Example for DQ to DQS



UG583_c2_104_010816


4. Differential signal group skew constraints in this chapter are presented in the form X. The difference between the longest signal and the shortest signal in the X group must be no greater than the specified constraint.

The following table shows an example *clock* group constraint specification. The following figure illustrates an interpretation of the constraint. The difference in delay between the longest clock signal in the bus and the shortest clock in the bus can be no greater than 2 ps.

Table 11: Example Signal Group Skew Constraint

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
ck_p and ck_n	2	12

Figure 13: Skew Example for Clock Signals



ug583_c2_105_060719

- The CK to DQS specification for DDR3 and DDR4 component interfaces encompasses a wide range, as shown in the following table. The wide range is to ensure that proper write leveling can take place at all memory devices, from the first in the chain to the last.


Table 12: Example CK to DQS Skew Constraint

Signal Group	Skew Constraint (ps)	Skew Constraint (mil)
CK to DQS	-149 to 1796	-879 to 10,600

The UltraScale device memory controller can internally delay the DQS line to account for negative skew, which helps because the CK line encounters more capacitive load than each individual DQS pair because the CK lines touch each memory device in the chain. This slows the CK line relative to each DQS, which only touches one memory device. The specification routes the CK and DQS lines where the skew between CK to DQS is no less than -149 ps from the UltraScale device to the first memory device in the chain and no more than 1796 ps to the furthest memory device in the chain. This is illustrated in the following figure. As long as the first memory device and the last memory device in the chain are bounded by this constraint, all memory devices can be write leveled.

For example, if the DQS delay between the UltraScale device and the first memory device is 200 ps, the clock delay from the UltraScale device to the first memory device should be at least 51 ps (200 ps - 149 ps). If the DQS delay from the UltraScale device to the furthest memory device is 700 ps, the clock delay from the UltraScale device to the last memory device should be less than 2496 ps (700 ps + 1796 ps).

Figure 14: Interpretation of CK to DQS Skew Specification




TIP: DIMM guidelines have a much narrower CK to DQS specification because the PCB is only routed from the UltraScale device to the first memory device in the DIMM. The routing from the first memory device to the last memory device is managed by the DIMM.

6. Skew constraints assume the fastest interface speeds. For slower speeds, certain skews can be relaxed. Refer to [Appendix B: Memory Derating Tables](#) for derating tables for the various memory interfaces.
7. Skew and length constraints assume a propagation time of 169.5 ps/in based on the stackup shown in [Table 9](#). Maximum trace lengths can be relaxed or tightened depending on the actual propagation time of the board. To convert, multiply the specified trace length by 169.5 ps, then divide by the actual propagation time of the board. See [Adjusting for Different Stack-Ups](#).
8. Trace widths and spacing are based on the stackup shown in [Table 9](#). If another stackup scheme is used, the widths and spacing can be modified to meet the impedance targets. See [Adjusting for Different Stack-Ups](#).
9. DQ and DQS signals in the same byte group should be routed in the same layer from UltraScale device to DRAM/DIMM, except in the breakout areas. Include the data mask (DM) in the byte group as applicable.
10. Do not change layers when routing from one DIMM to the next, if applicable. In addition, for DIMM routing, it is recommended to route data byte groups on the highest signal layers (closest to the DIMM connector) as much as possible, especially for byte groups located near the center of the DIMM.
11. For fly-by routing, address, command, and control signals can be routed on different layers, though it is recommended to use as few as possible. Do not route any individual signal on more than two layers to minimize inductive loops that can lead to crosstalk issues. Any signal layer switching via needs to have one ground via within a 50 mil perimeter range.

12. UltraScale device and memory drive strengths depend on the particular memory standard and memory component(s) being used. Memory guidelines in this document are based on the default standard used by the Memory Interface Generator (MIG).
13. When using the internal V_{REF} , the dedicated V_{REF} pin can be tied to GND with a 500Ω or 1 kΩ resistor. Take care to minimize via coupling so as to reduce noise coupling through the V_{REF} pin.
Note: When internal V_{REF} is used, this pin cannot be used as an I/O.
14. If the system clock is connected to a bank that is also used for memory, terminate as shown for LVDS in the following figure with the appropriate pull-up voltage. This termination circuit is necessary because of the different I/O standard of the memory bank (HSTL, SSTL, or POD).


Figure 15: System Clock



UG583_c2_57_100623

15. Signal lines must be routed over a solid reference plane. Avoid routing over voids (see the following figure).


Figure 16: Signal Routing Over Solid Reference Plane



UG583_c2_13_050614


16. Avoid routing over reference plane splits (see the following figure).

Figure 17: Signal Routing Over Reference Plane Split



17. Keep the routing at least 30 mils away from the reference plane and void edges with the exception of breakout regions (see [Figure 16: Signal Routing Over Solid Reference Plane](#)).
18. In the breakout region, route signal lines in the middle of the via void aperture. Avoid routing at the edge of via voids (see the following figure).

Figure 18: Breakout Region Routing




UG583_c2_15_051915

- Use chevron-style routing to allow for ground stitch vias. The following figure shows recommended routing for fly-by configurations, while [Figure 20: Example of Ground Stitching \(Clamshell\) Red: Power, Green: Ground](#) shows recommended routing to accommodate ground stitch vias in a more congested clamshell configuration.


IMPORTANT! Crosstalk issues leading to data errors can occur with lack of proper ground stitching, especially in areas where there are fewer ground pins, such as near address pins of memory devices.

Figure 19: Example of Ground Stitching (Fly-by)




UG583_c2_16_000614

Figure 20: Example of Ground Stitching (Clamshell) Red: Power, Green: Ground



The following figure shows simulated eye diagrams for a DDR4 command/address/control bit with and without ground stitching vias. The simulation on the left shows an eye height of 180 mV with ground stitch vias, while the simulation on the right shows an eye height of only 99 mV when not using ground stitch vias.

Figure 21: Simulations With and Without Ground Stitching Vias



20. Add ground vias as much as possible around the edges and inside the device (FPGA, MPSoC, memory component, DIMM) to make a better ground return path for signals and power, especially corners. Corner or edge balls are generally less populated as grounds.
21. For address/command/control V_{TT} termination, every four termination resistors should be accompanied by one 0.1 μF capacitor, physically interleaving among resistors, as shown in the following figure. Refer to the memory vendor's data sheet for specifications regarding noise limits on the address/command/control V_{TT} lines.

Figure 22: Example of V_{TT} Termination Placement






Figure 23: Schematic Example of VTT Resistor and Capacitor Connections



UG583_c2_60_090816

For five components, the recommendation is shown in the following figure.

Figure 25: Component Placement Recommendations for Five Components



UG583_c2_19_073014

Adjusting for Different Stack-Ups

The trace width, spacing, length, and skew constraints presented in this chapter are based on the reference stackup shown in [Table 9: Reference Stackup](#). When not using this stackup, trace width, lengths, layer heights, spacings, and dielectric material could need adjustments to meet the impedance, length, and skew specifications.

The PCB fabrication house can adjust these factors to achieve the desired impedance and propagation delay targets. The effects of each of these items are listed in this section. In addition, a two-dimensional field solver utility allows for various combinations to be tested.

Dielectric Material

Each dielectric material has its own relative dielectric constant (D_K) and dissipation factor/loss tangent (D_F) that are contributing factors to line impedance (Z_0), signal propagation delay (T_{PD}), and signal loss (α). As D_K increases, impedance decreases, while the signal propagation delay and signal loss increase, and vice versa. A typical D_K found on PCBs range from 3.4 to 4.6. The propagation delay in a given dielectric material is constant and is not affected by any other board parameters such as layer height, conductor width, or conductor spacing. Propagation delay is affected by frequency, but the effect is minimal with regards to typical memory speeds. Signal loss is also affected by frequency, with the loss increasing with increasing frequency.

The following equation shows the calculation for propagation delay (T_{PD}), with D_K as the dielectric constant, and c as the speed of light in free space (2.998×10^8 m/s or 1.180×10^{10} in/s).

Equation 1: T_{PD}

$$T_{PD} = \frac{\sqrt{D_K}}{c}$$

The associations in the following table show the effect of the dielectric constant (D_K) on the impedance (Z_0), propagation delay (T_{PD}), and signal loss (α).

$D_K \uparrow$	$Z_0 \downarrow$	$T_{PD} \uparrow$	$\alpha \uparrow$
$D_K \downarrow$	$Z_0 \uparrow$	$T_{PD} \downarrow$	$\alpha \downarrow$

Trace Width

As the trace width (W) increases, the impedance decreases, while the signal propagation delay remains unchanged, and vice-versa. Any adjustments in width should include adjustments in spacing (S) to maintain immunity to crosstalk effects. Spacing factors are roughly 1.0X to 3X depending on the particular type of memory and signal. The associations in the following table show the effect of width on impedance, spacing, and propagation delay.

$W \uparrow$	$Z_0 \downarrow$	$S \uparrow$	T_{PD} (no change)
$W \downarrow$	$Z_0 \uparrow$	$S \downarrow$	T_{PD} (no change)

Layer Height

As the layer height (H) is increased, the impedance increases, while the signal propagation delay remains unchanged, and vice-versa. When reducing or increasing layer height, consider that layer heights that are too low can be more expensive to reliably manufacture, because the PCB fabricator must avoid plane shorts. Layer heights that are too high can lead to aspect ratio violations. The associations in the following table show the effect of layer height on impedance and propagation delay.

H ↑	Z_0 ↑	T_{PD} (no change)
H ↓	Z_0 ↓	T_{PD} (no change)

Base Copper Weight

As the base copper weight (Cu oz) of the conducting lines is increased, the impedance will decrease, while the signal propagation delay remains unchanged, and vice-versa. The most common base copper weight used in PCBs is 0.5, known as half-ounce copper. The effect of changing the base copper weight is minimal, however, and is not recommended due to the potential for increased costs in relation to the small benefit typically gained.

Example of Tuning Design Parameters to Meet Impedance Targets

The following table shows a target of 36Ω for stripline impedance in the main L1 PCB area. The trace width (W) is specified as 7.0 mils along with 8.0 mil spacing (S). The reference PCB stackup in [Table 9: Reference Stackup](#) defines the layer composition and dielectric material that will result in reaching 36Ω . This consists of a 4.0 dielectric constant (D_K), and 9.6 mil layer height (H) (4.5 + 0.6 + 4.5). The following table shows tuning of the trace width (W) and height (H) required to meet the impedance target. The spacing (S) shows how it should change to continue to meet performance requirements, though the spacing is not a factor in the line impedance. When increasing D_K , the signal loss also increases per [Dielectric Material](#).

Table 16: Example of Tuning PCB Parameters to Meet Impedance Target

Target (Z_0) Ω	Material	D_K	W (mil)	S (mil)	H (mil)	Description
36	Isola FR-4 370H	4.00	7.0	8.0	9.6	See Table 9: Reference Stackup .
36	Megtron 6	3.71	7.0	8.0	9.2 ↓	Same width and spacing as reference.
36	Megtron 6	3.71	7.5 ↑	8.5 ↑	9.6	Same height as reference.

Table 16: Example of Tuning PCB Parameters to Meet Impedance Target (cont'd)

Target (Z_0) Ω	Material	D_K	W (mil)	S (mil)	H (mil)	Description
36	FR4 (standard)	4.60	7.0	8.0	10.6 ↑	Same width and spacing as reference.
36	FR4 (standard)	4.60	6.0 ↓	7.0 ↓	9.6	Same height as reference.

PCB Guidelines for DDR4 SDRAM (PL and PS)

Overview

This section provides electrical design guidelines for typical DDR4 SDRAM interfaces from system-level signal integrity simulations for the UltraScale architecture.

Note: For applications that use the Zynq UltraScale+ device Video Codec Unit (VCU) and DDR4 memory in the PL, refer to [Answer Record 71209](#) for important design considerations.

IMPORTANT! All routing guidelines in this section must be followed to achieve the maximum data rates specified for the DDR4 SDRAM interface for typical system designs. Customers could have unique or specific designs with particular violations of some rules. In these scenarios, design or routing trade-offs have to be taken in other routing parameters to mitigate the risk. System-level channel signal integrity simulations are required to evaluate such trade-offs. It is important to read [General Memory Routing Guidelines](#) before continuing with this section.

DDR4 SDRAM Interface Signal Description

The DDR4 SDRAM interface consists of clock, control, address, and data signals as shown in the following table.

Table 17: DDR4 SDRAM I/O Signal Description

Signal Name	Description
Clock Signals	
ck_t, ck_c	Differential clock
Address and Command Signals	
a[17,13:0]	Address inputs
ras_n/a[16]	Row address strobe, address bit 16
cas_n/a[15]	Column address strobe, address bit 15
we_n/a[14]	Write enable, address bit 14
bg[1:0]	Bank group inputs

Table 17: DDR4 SDRAM I/O Signal Description (cont'd)

Signal Name	Description
ba[1:0]	Bank address inputs
act_n	Activation command input
par	Command and address parity input
Control Signals	
cke	Clock enable
cs_n[3:0]	Chip select
odt	On-die termination enable
reset_n	See reset_n
Data Signals	
dq[79:0]	Data input/output
dqs_t/dqs_c[9:0]	Data strobe (differential)
dm_n/dbi_n	Data mask and data bus inversion
Other Signals	
TEN	Test connectivity mode. CMOS level. Connect 500Ω resistor to ground at memory devices.
alert_n	See alert_n
VRP (PL)	240Ω to GND
ZQ (PS)	240Ω to GND


Notes:

- Actual signal list might vary based on configuration.

Fly-by and Clamshell Topologies

Two topology types are supported for DDR4 SDRAM: fly-by, and clamshell. The fly-by topology (see the following figure) consists of all memory devices on one layer, usually in-line. This type of topology is generally easier to route and can offer the best signal integrity, but can take up precious board real estate.

Figure 26: Fly-by Topology




UG583_c2_61_090816

Clamshell topology (see the following figure) requires more intricate routing, but is optimal for designs where board space is at a premium.

★ IMPORTANT! Clamshell is a supported DDR4 SDRAM topology in the MIG tool and is selectable for Programmable Logic banks only. The PS in the Zynq UltraScale+ MPSoC does not have a selectable clamshell configuration option. However, the PS can be configured as clamshell if set up as a dual-rank configuration with the first rank on the top layer, and the second rank mirrored on the bottom layer. When using this topology within the AMD Vivado™ tools, refer to the Clamshell Topology section in *UltraScale Architecture-Based FPGAs Memory IP LogiCORE IP Product Guide (PG150)* for additional information.

Figure 27: Clamshell Topology



Using Address Mirroring to Ease Clamshell Routing

While the clamshell topology has advantages such as reduced board space, the asymmetrical nature of the memory device package pin locations can lead to higher routing congestion in the areas under the memory devices between the top and bottom layers. This routing congestion could potentially lead to higher crosstalk due to the inability to include as many ground-return vias as in other topologies (see item 21 in [General Memory Routing Guidelines](#)), as well as longer stub traces on the top and bottom layers. Address mirroring can be used to change the function of certain pins on a memory device to correspond to the pin directly above or below it. One via can be used for the signal, along with a short stub to the landing pad for each device. Address mirroring is defined in JEDEC specification JESD21-C. Twelve command/address/control pins can be mirrored for DDR4 SDRAM, according to the following table.

Table 18: Allowable Mirror Pins for DDR4 SDRAM

Memory Controller Pin	DRAM Pin (Non-Mirrored)	DRAM Pin (Mirrored)
A3	A3	A4
A4	A4	A3
A5	A5	A6
A6	A6	A5
A7	A7	A8
A8	A8	A7
A11	A11	A13
A13	A13	A11
BA0	BA0	BA1
BA1	BA1	BA0
BG0 ¹	BG0	BG1

Table 18: Allowable Mirror Pins for DDR4 SDRAM (cont'd)

Memory Controller Pin	DRAM Pin (Non-Mirrored)	DRAM Pin (Mirrored)
BG1 ¹	BG1	BG0

Notes:

1. BG0 and BG1 can only be mirrored when pin BG1 is present on the memory device.

ECC Connection Rules for DDR4 SDRAM

For DDR4 interfaces when using ECC with x16 components, the lower data byte (DQ[7:0]) of the last component must be connected to the ECC byte of the controller. If the lower data byte is not connected and instead the upper data byte is used, the per-DRAM addressability feature does not function correctly, which results in V_{REF} calibration errors. This results in FSBL hangs for Zynq UltraScale+ MPSoC based designs. Functionality is not impacted for soft controller designs because V_{REF} calibration is not used. However, if V_{REF} calibration is enabled, it results in functional errors.

For DDR4-based interfaces, the unused upper data byte must be connected as follows (unless explicitly stated otherwise in the memory vendor's datasheet):

- DQ[15:8] can be left floating.
- UDM or UDM_n/UDBI_n terminate to V_{DD} with a pull-up resistor $\leq 4.7\text{ K}\Omega$.
- UDQS or UDQS_t terminate to V_{DD} with a pull-up resistor $\leq 4.7\text{ K}\Omega$.
- UDQS# or UDQS_c terminate to V_{SS} with a pull-down resistor $\leq 4.7\text{ K}\Omega$.

Routing Rule Changes for Thicker Printed Circuit Boards

The DDR4 routing rules provided in this chapter allow for the interface to run at the maximum supported data rates as specified in the *Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)* while also assuming a maximum board thickness of 100 mil. For boards that are thicker than 100 mil, the maximum interface speed might need to be adjusted downward. See the following table for maximum memory data rates per board thickness.


Table 19: DDR4 Maximum Data Rates per Board Thickness

DDR4	Board Thickness (mil)			
	≤ 100	101 - 145	101 - 145 + Stub-free ¹	145 - 275
Component	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	Two memory speed grades lower from data rates in <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	Three memory speed grades lower from data rates in <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>
1-Rank DIMM	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	One memory speed grade lower from data rates in <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>
2-Rank DIMM	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	One memory speed grade lower from data rates in <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>
4-Rank DIMM	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	See <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>	One memory speed grade lower from data rates in <i>Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)</i>

Notes:


1. Stub-free refers to alternating the routing of the signals between upper and lower layers such that long via stubs are avoided. See [Figure 29: Thicker PCB: Command/Address/Control Bus Stub-Free Routing](#) for an example of alternate layer routing as compared to conventional routing shown in the following figure.

Figure 28: Thicker PCB: Command/Address/Control Bus Conventional Routing



X25066-012821

Figure 29: Thicker PCB: Command/Address/Control Bus Stub-Free Routing




X25068-012821

Topology and Routing Guidelines for DDR4 SDRAM

DDR4 SDRAM Address, Command, and Control Fly-by and Clamshell Topologies

The following figure illustrates fly-by topology for address, command, and control signals. Each clock, address, command, and control pin on each SDRAM is connected to a single trace and terminated at the far end.

Figure 30: Address, Command, and Control Fly-by Termination for DDR4 SDRAM




UG583_c2_23_062122

reset_n

The following figure shows the termination for *reset_n*. The 4.7 kΩ to ground is to keep *reset_n* Low during FPGA power-up. If self-refresh is required during FPGA power-down, circuitry needs to be added to ensure that *reset_n* stays High during that time.

Figure 31: Termination for reset_n in DDR4 DRAM



alert_n

For Zynq UltraScale+ MPSoC designs with component interfaces, connect the PS_DDR_ALERT_N pin to the ALERT_N pins of the DDR4 devices in fly-by routing and terminate to V_{DD} with a 50Ω pull-up resistor. For Zynq UltraScale+ MPSoC designs with DIMMs, connect the PS_DDR_ALERT_N pin to the ALERT_N pin of the connector. For PL-based designs, the controller does not generate an ALERT_N pin but the DDR4 component ALERT_N signals must be connected together via fly-by routing and terminated to V_{DD} with a 50Ω pull-up resistor. For PL-based designs with DIMMs, leave the ALERT_N pin floating at the connector. The ALERT_N signal does not have any skew or length matching requirements.

The following table shows the DDR4 SDRAM fly-by impedance, length, and spacing guidelines for address, command, and control signals.

Table 20: DDR4 SDRAM Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R_{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	36±10%	50±10%	50±10%	39±10%	Ω
Trace width	4.0	7.0	4.0	4.0	6.0	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	0.0~0.1	0.35~0.75	0~1	inches
Spacing in address, command, and control signals (minimum)	4.0	8.0 ²	4.0	8.0	8.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	30	30	30	mil

Table 20: DDR4 SDRAM Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals (cont'd)


Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Maximum PCB via count			7			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 12.0.

For DDR4 SDRAM clamshell topology, an alternating fly-by topology is recommended for control/address/command signals. The alternating layer routing properly balances the signal loads at each memory device. As depicted in the following figure with the FPGA located at the top layer, the inner layer routing to top layer devices 1, 3, 5, 7, and 9 is closer to the top layer, while the inner layer routing to bottom layer devices 2, 4, 6, and 8 is closer to the bottom layer.

Figure 32: Address/Command/Control Topology for DDR4 SDRAM Clamshell



Note: The end-termination resistor can be located on either the top or bottom layer.

DDR4 SDRAM clamshell topology uses two individual chip select (CS) signals, one for the top layer DRAMs, and one for the bottom layer DRAMs. Those signals should be routed as in the standard fly-by topology as defined in [Figure 30: Address, Command, and Control Fly-by Termination for DDR4 SDRAM](#) and the previous table.

The following table shows the DDR4 SDRAM clamshell impedance, length, and spacing guidelines for address, command, and control signals. Note the extra length of the L2 segment.

Table 21: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Address/Command/Control Signals

Parameter	L0 FPGA Breakout	L1	L2	L3	L4	L5	L6 V _{TT} Stub	Units
Layer (recommended)	Upper inner	Upper inner	Lower inner	Upper inner	Top	Bottom	Bottom	
Impedance Z ₀	50	50	50	50	50	50	39	Ω
Length	0.0~1.5 ¹	0.0~4.0	L3+0.2	0.45~0.85	For mirrored case, L4 = L5; for non-mirrored case, make as short as possible		≤ 0.95	inch
Width	4.0	4.0	4.0	4.0	4.0	4.0	6.0	mil

Table 21: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Address/Command/Control Signals (cont'd)

Parameter	L0 FPGA Breakout	L1	L2	L3	L4	L5	L6 V _{TT} Stub	Units
Spacing within group	4.0	8.0 ¹	8.0	8.0	8.0	8.0	8.0	mil
Spacing to clocks	8.0	20	20	20	20	20	20	mil
Spacing to other groups	8.0	30	30	30	30	30	30	mil

Notes:

1. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 12.0.




IMPORTANT! Add as many ground vias as possible to help avoid crosstalk issues. See Item 20 in [General Memory Routing Guidelines](#).

DDR4 SDRAM Clock Fly-By and Clamshell Termination

Inherent to fly-by topology, the timing skew between the clock and dq_s signals is deskewed by the write-leveling feature on DDR4 SDRAM (see the following figure).

Figure 33: Clock Fly-By Termination for DDR4 SDRAM



The following table shows the DDR4 SDRAM impedance, length, and spacing guidelines for clock signals.

Table 22: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Clock differential impedance Z _{DIFF}	86±10%	76±10%	86±10%	90±10%	76±10%	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	4.0/5.0/4.0	6.0/6.0/6.0	mil

Table 22: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals (cont'd)


Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	0.0~0.1	0.35~0.75	0~1.0	inches
Spacing in address, command, and control signals (minimum)	8.0	20 ²	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	30	mil
Maximum PCB via count per signal	7					-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 30.

For DDR4 SDRAM clamshell topology, an alternating fly-by topology is recommended for clock signals. The alternating layer routing properly balances the signal loads at each memory device. As depicted in [Table 36: Impedance, Length, and Spacing Guidelines for One-Slot DIMM Address, Command, and Control Signals](#) with the FPGA located at the top layer, the inner layer routing to top layer devices 1, 3, 5, 7, and 9 is closer to the top layer, while the inner layer routing to bottom layer devices 2, 4, 6, and 8 is closer to the bottom layer.

Figure 34: Clock Topology for DDR4 SDRAM Clamshell



Note: The end-termination components can be located on either the top or bottom layer.

The following table shows the DDR4 SDRAM clamshell impedance, length, and spacing guidelines for clock signals. Note the extra length of the L2 segment.

Table 23: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Clock Signals

Parameter	L0 FPGA Breakout	L1	L2	L3	L4	L5	L6 V _{TT} Stub	Units
Layer (recommended)	Upper inner	Upper inner	Lower inner	Upper inner	Top	Bottom	Bottom	
Differential impedance Z ₀	86	76	90	90	86	86	76	Ω
Length	0.0~1.5 ¹	0.0~4.0	L3+0.2	0.45~0.85	For mirrored case, L4 = L5; for non-mirrored case, make as short as possible		≤ 1.0	inch
Width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/5.0/4.0	4.0/5.0/4.0	4.0/4.0/4.0	4.0/4.0/4.0	6.0/6.0/6.0	mil
Spacing to addr/cmd/ctrl	8.0	20 ¹	20	20	20	20	20	mil
Spacing to other groups	8.0	30	30	30	30	30	30	mil

Notes:


1. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 30.

★ IMPORTANT! Add as many ground vias as possible to help avoid crosstalk issues. See item 20 in [General Memory Routing Guidelines](#).

DDR4 SDRAM Data Signals Point-to-Point for Fly-by and Clamshell Configurations

The following figure shows the data signals (DQ, DM, and DQS) point-to-point for DDR4 SDRAM. The guidelines are the same for fly-by and clamshell configurations.

Figure 35: Data Signals Point-to-Point for DDR4 SDRAM



The following table shows the DDR4 SDRAM impedance, length, and spacing guidelines for data signals.

Table 24: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
dq single-ended impedance Z_0	50±10%	39±10%	50±10%	Ω
dqs differential impedance Z_{DIFF}	86±10%	76±10%	86±10%	Ω
Trace width (nominal)	4.0	6.0	4.0	mil
Differential trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	mil
Trace length	0.0~1.5 ¹²	0.0~4.0	0.0~0.1	inches
Spacing in byte (minimum)	4.0	8.0 ²	4.0	mil
Spacing byte to byte (minimum)	4.0	20	4.0	mil
dq to dqs spacing (minimum)	4.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	2			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 12.0.

DDR4 SDRAM Routing Constraints

There are two constraints requirements for each signal group in the DDR4 memory interface:

- Total length/delay constraints
- Skew constraints

The total length/delay constraints are shown in the following table.

Table 25: DDR4 SDRAM Total Length/Delay Constraints

Signal Group	Reference Figure	Maximum Length/Delay Constraints
address/command/control: x1/x2/x3/x4/x5/x6/x7/x8/x9 (includes package routing P0)	Figure 30: Address, Command, and Control Fly-by Termination for DDR4 SDRAM	7.15/7.9/8.65/9.4/10.15/10.9/11.65/12.4/13.15 inches 1211/1339/1466/1593/1720/1847/1974/2101/2228 ps
Data Signals P0+L0+L1+L2	Figure 35: Data Signals Point-to-Point for DDR4 SDRAM	6.0 inches, 1017 ps

The skew constraints are listed in the following tables.

Table 26: DDR4 SDRAM Data Group Skew Constraints

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
Data ¹ to DQS	±10	±58

Table 26: DDR4 SDRAM Data Group Skew Constraints (cont'd)

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
dqs_p and dqs_n	2	12
CK to DQS	-149 to 1,796	-879 mil to 10.59 inches

Notes:

1. The data group includes dq and dm_n/dbi_n.
2. Delays are based on 169.5 ps/in. See item 8 in [General Memory Routing Guidelines](#).
3. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.

The data group skew constraints are listed in the following table.

Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints

Signal	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
address/command/control ¹ to CK	UltraScale device to each memory device	±8 ps	±47
ck_p and ck_n	UltraScale device to each memory device	2	12

Notes:

1. The signal reset_n is not required to meet the skew constraints in this table.
2. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

PCB Guidelines for DDR3/3L SDRAM (PL and PS)

Overview

This section provides electrical design guidelines for typical DDR3 SDRAM interfaces (up to nine devices) from system-level signal integrity simulations for the UltraScale architecture. A five-component DDR3 80-bit wide interface using x16 devices is used as a design example in this document.

★ IMPORTANT! All routing guidelines in this section must be followed to achieve the maximum data rates specified for the DDR3/3L SDRAM interface for typical system designs. Customers could have unique or specific designs with particular violations of some rules. In these scenarios, design or routing trade-offs have to be taken in other routing parameters to mitigate the risk. System-level channel signal integrity simulations are required to evaluate such trade-offs. It is important to read the [General Memory Routing Guidelines](#) section before continuing with this section.

DDR3 SDRAM Interface Signal Description

The DDR3 SDRAM interface consists of clock, control, address, command, and data signals as shown in the following table.

Table 28: DDR3 SDRAM Interface Signal Description

Signal Name	Description
Clock Signals	
ck_p/n[1:0]	Differential clock
Control Signals	
cke[1:0]	Clock enable
cs_n[1:0]	Chip select
odt[1:0]	On-die termination enable
reset_n	See reset_n
Address Signals	
a[15:0]	Memory address bus
ba[2:0]	Bank address
Command Signals	
ras_n	Row address select
cas_n	Column address select
we_n	Write enable
Data Signals	
dq[79:0]	Data bus
dqs_p/n[9:0]	Differential data strobe
dm[9:0]	Data mask
Other Signals	
VRP (PL)	240Ω to GND
ZQ (PS)	240Ω to GND

Notes:

- Actual signal list might vary based on configuration.

Table 29: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z ₀	50±10%	36±10%	50±10%	50±10%	39±10%	Ω
Trace width	4.0	7.0	4.0	4.0	6.0	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	0.0~0.1	0.35~0.75	0~1.0	inches
Spacing in address, command, and control signals (minimum)	4.0	8.0	4.0	8.0	8.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	30	30	30	mil
Maximum PCB via count	7					-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 12.0.

DDR3 SDRAM Clock Fly-by Termination

Inherent to fly-by topology, the timing skew between the clock and dq_s signals is deskewed by the write leveling feature on DDR3 SDRAM (see the following figure).

Figure 38: Clock Fly-by Termination for DDR3 SDRAM



The following table shows the DDR3 SDRAM impedance, length, and spacing guidelines for clock signals.

Table 30: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Clock differential impedance Z _{DIFF}	86±10%	76±10%	86±10%	90±10%	76±10%	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	4.0/5.0/4.0	6.0/6.0/6.0	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	0.0~0.1	0.35~0.75	0~1.0	inches
Spacing in address, command, and control signals (minimum)	8.0	20 ²	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	8	30	30	mil
Maximum PCB via count per signal	7					-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 30.

DDR3 SDRAM Data Signals Point-to-Point

The following figure shows the data signals (DQ, DM, and DQS) point-to-point for DDR3 SDRAM.

Figure 39: Data Signals Point-to-Point for DDR3 SDRAM



The following table shows the DDR3 SDRAM impedance, length, and spacing guidelines for data signals.

Table 31: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
dq single-ended impedance Z_0	50±10%	39±10%	50±10%	Ω
dqs differential impedance Z_{DIFF}	86±10%	76±10%	86±10%	Ω
Trace width (nominal)	4.0	6.0	4.0	mil
Differential trace width/ space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	mil
Trace length (nominal)	0.0~1.5 ^{1,2}	0.0~5.0	0.0~0.1	inches
Spacing in byte (minimum)	4.0	8.0 ²	4.0	mil
Spacing byte to byte (minimum)	4.0	20	4.0	mil
dq to dqs spacing (minimum)	4.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	2			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 12.0.

DDR3 SDRAM Routing Constraints

There are two constraints requirements for each signal group in the DDR3 memory interface:

- Total length/delay constraints
- Skew constraints

The total length/delay constraints are shown in the following table.

Table 32: DDR3 SDRAM Total Length/Delay Constraints

Signal Group	Reference Figure	Maximum Length/Delay Constraints
Address/command/control: x1/x2/x3/x4/x5/x6/x7/x8/x9 (includes package routing P0)	Figure 36: Address Fly-by Termination for DDR3 SDRAM	6.15/6.9/7.65/8.4/9.15/9.9/10.65/11.4/12.15 inches 1042/1169/1296/1423/1550/1678/1805/1932/2110 ps
Data Signals P0+L0+L1+L2	Figure 39: Data Signals Point-to-Point for DDR3 SDRAM	7.0 inches, 1186 ps

The skew constraints are shown in the following tables.

Table 33: DDR3 SDRAM Data Group Skew Constraints

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
Data ¹ to DQS	±10	±58
dqs_p and dqs_n	2	12
CK to DQS	-149 to 1,796	-879 mil to 10.6 inches

Notes:

1. The data group includes dq and dm.
2. Delays are based on 169.5 ps/in. See item 8 in [General Memory Routing Guidelines](#).
3. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3-8.

The address, command, and control clock group skew constraints are listed in the following table.

Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints

Signal	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
Address/command/control ¹ to CK	UltraScale device to each memory device	±8 ps	±47
ck_p and ck_n	UltraScale device to each memory device	2	12

Notes:

1. The signal reset_n is not required to meet the skew constraints in this table.
2. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3-8.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

DDR3/DDR4 UDIMM/RDIMM/SODIMM/LRDIMM Routing Guidelines (PL and PS)

For the case of two dual in-line memory modules (DIMMs), each DIMM has its own set of point-to-point clock and control signals, while data, command, and address route to both DIMMs in a fly-by topology.



IMPORTANT! The Processing System (PS) does not support multiple DIMMs.




IMPORTANT! [Figure 42: CMD/ADDR Fly-by Routing for DDR3/DDR4 DIMM Topology](#), [Figure 44: DQS Fly-by Routing for DDR3/DDR4 DIMM Topology](#), and [Figure 45: DQ Fly-by Routing for DDR3/DDR4 DIMM Topology](#) show configurations with multiple DIMMs. To reduce the effect of SI reflections, always place DIMM #0 on the furthest connector from the FPGA. DIMM #1, if present, should be placed on the connector nearest to the FPGA.

DDR3/DDR4 DIMM Clock Point-to-Point Routing

The following figure and table define the routing topology and guidelines for DDR3/DDR4 DIMM clock signals.

Figure 40: CK Point-to-Point Routing for DDR3/DDR4 DIMM Topology



UG583_c2_43_060719

Table 35: Impedance, Length, and Spacing Guidelines for DIMM Clock Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	Units
Trace type	Stripline	Stripline	-
Differential impedance Z_0	86±10%	66±10%	Ω
Trace width/space/width	4.0/4.0/4.0	8.0/8.0/8.0	mil
Trace length	0.0~1.5 ¹²	0.0~4.0	inches
Spacing to address, command, and control signals (minimum)	8.0	20 ²	mil
Spacing to other group signals (minimum)	8.0	30	mil
Maximum PCB via count	2		-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 30.

DDR3/DDR4 DIMM Control, Command, and Address Routing


The following figures and tables define the routing topologies and routing guidelines for DDR3/DDR4 DIMM address, command, and control signals. Place four 0402 1.0 μ F capacitors in the CMD/CTRL pin area of the DIMM connector. All capacitors should connect from V_{CC0} to GND.

Figure 41: CTRL Point-to-Point Routing for DDR3/DDR4 DIMM Topology



UG583_c2_44_060719

Figure 42: CMD/ADDR Fly-by Routing for DDR3/DDR4 DIMM Topology



UG583_c2_45_080919

Note: The Zynq UltraScale+ MPSoC DDR interface does not support dual-DIMM topologies.

Table 36: Impedance, Length, and Spacing Guidelines for One-Slot DIMM Address, Command, and Control Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	Units
Trace type	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	Ω
Trace width	4.0	6.0	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	inches
Spacing in address, command, and control signals (minimum)	4.0	12 ²	mil
Spacing to clock signals (minimum)	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	mil
Maximum PCB via count	2		-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 16.0.

Table 37: Impedance, Length, and Spacing Guidelines for Two-Slot DIMM Address, Command, and Control Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DIMM to DIMM)	Units
Trace type	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	34±10% addr/cmd 39±10% ctrl	34±10% addr/cmd 39±10% ctrl	Ω

Table 37: Impedance, Length, and Spacing Guidelines for Two-Slot DIMM Address, Command, and Control Signals (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DIMM to DIMM)	Units
Trace width	4.0	6.9 addr/cmd 6.0 ctrl	6.9 addr/cmd 6.0 ctrl	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	<0.5	inches
Spacing in address, command, and control signals (minimum)	4.0	12.0 ctrl, 13.8 addr/cmd ²	12.0 ctrl, 13.8 addr/cmd	mil
Spacing to clock signals (minimum)	8.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	Addr/cmd: 3 Ctrl: 2			-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 18.0 ctrl, and 20.7 addr/cmd.

reset_n

The following figure shows the termination for reset_n. The 4.7 kΩ to ground is to keep reset_n Low during FPGA power-up. If self-refresh is required during FPGA power-down, circuitry needs to be added to ensure that reset_n stays High during that time.

Figure 43: Termination for reset_n in DDR4 DRAM




UG583_c2_58_111115

DDR3/DDR4 DIMM Data (DQ and DQS) Routing

The following figures and table define routing topologies and guidelines for DDR3/DDR4 DIMM data routing.


Figure 44: DQS Fly-by Routing for DDR3/DDR4 DIMM Topology



UG583_c2_47_080919

Note: The Zynq UltraScale+ MPSoC DDR interface does not support dual-DIMM topologies.

Figure 45: DQ Fly-by Routing for DDR3/DDR4 DIMM Topology



UG583_c2_46_080919

Note: The Zynq UltraScale+ MPSoC DDR interface does not support dual-DIMM topologies.

Table 38: Impedance, Length, and Spacing Guidelines for DIMM Data Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DIMM to DIMM)	Units
Trace type	Stripline	Stripline	Stripline	-
dq single-ended impedance Z_0	50±10%	39±10%	39±10%	Ω
dqs differential impedance Z_{DIFF}	86	66	66	Ω
dq trace width (nominal)	4.0	6.0	6.0	mil
Dqs differential trace width/ space/width	4.0/4.0/4.0	8.0/8.0/8.0	8.0/8.0/8.0	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	≤ 0.5	inches
Spacing in byte (minimum)	4.0	15 ²	15	mil
Spacing byte to byte (minimum)	4.0	30	30	mil
dq to dqs spacing (minimum)	4.0	21	21	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	2 for 1 DIMM, 3 for two DIMMs			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 20.0.

DDR3 UDIMM/RDIMM/SODIMM/LRDIMM Routing Constraints

There are two constraints requirements for each signal group in the DDR3 DIMM interface:

- Total length/delay constraints
- Skew constraints

The total length/delay constraints are shown in the following table.

Table 39: DDR3 DIMM Total Length/Delay Constraints

Signal Group	Reference Figure	Maximum Length/Delay Constraints
Address/command/control: x1/x2 (includes package routing P0)	Figure 42: CMD/ADDR Fly-by Routing for DDR3/DDR4 DIMM Topology	6.15/6.65 inches, 1042/1127 ps
Data Signals P0+L0+L1+L2	Figure 44: DQS Fly-by Routing for DDR3/DDR4 DIMM Topology	7.0 inches, 1186 ps

The skew constraints are shown in the following tables.

Table 40: DDR3 DIMM Data Group Skew Constraints

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
Data ¹ to DQS	±10	±58
dqs_p and dqs_n	2	12
CK to DQS ¹	±150	±885

Notes:

1. Clock to DQS constraints should be from the UltraScale device to the DIMM.
2. The data group includes dq and dm.
3. Delays are based on 169.5 ps/in. See item 8 in [General Memory Routing Guidelines](#).
4. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.

Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints

Signal	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
Address/command/control ¹ to CK	UltraScale device to each memory device	±8 ps	±47
ck_p and ck_n	UltraScale device to each DIMM	2	12

Notes:

1. The signal `reset_n` is not required to meet the skew constraints in this table.
2. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

DDR4 UDIMM/RDIMM/SODIMM/LRDIMM Routing Constraints

There are two constraints requirements for each signal group in the DDR4 DIMM interface:

- Total length/delay constraints
- Skew constraints

The total length/delay constraints are shown in the following table.

Table 42: DDR4 DIMM Total Length/Delay Constraints

Signal Group	Reference Figure	Maximum Length/Delay Constraints
Address/command/control: x1/x2 (includes package routing P0)	Figure 42: CMD/ADDR Fly-by Routing for DDR3/DDR4 DIMM Topology	7.15/7.65 inches 1212/1296 ps

Table 42: DDR4 DIMM Total Length/Delay Constraints (cont'd)

Signal Group	Reference Figure	Maximum Length/Delay Constraints
Data Signals P0+L0+L1+L2	Figure 44: DQS Fly-by Routing for DDR3/DDR4 DIMM Topology	6.0 inches 1017 ps

The skew constraints are listed in the following tables.

Table 43: DDR4 DIMM Skew Constraints

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
Data ¹ to DQS	±10	±58
dqs_p and dqs_n	2	12
CK to DQS ²	±150	±885

Notes:

1. The data group includes dq and dm_n/dbi_n.
2. Clock to DQS constraints should be from the UltraScale device to the DIMM.
3. Delays are based on 169.5 ps/in. See item 8 in [General Memory Routing Guidelines](#).
4. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.

Table 44: DDR4 DIMM Address, Command, and Control Skew Constraints

Signal	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
Address/command/ control ¹ to CK	UltraScale device to each memory device	±8 ps	±47
ck_p and ck_n	UltraScale device to each DIMM	2	12

Notes:

1. The signal reset_n is not required to meet the skew constraints in this table.
2. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

PCB Guidelines for LPDDR4 Memories without ECC (PS)

Overview

This section provides electrical design guidelines for 32-bit LPDDR4 single-die package (SDP) and dual-die package (DDP) interfaces without error correcting code (ECC).

IMPORTANT! All guidelines in this section must be followed to achieve the maximum data rates specified for the LPDDR4 memory interface.

LPDDR4 x32 without ECC Memory Interface Signals and Connections

Signals and Connections for x32 Dual Channel LPDDR4 SDP without ECC

The required signals for a x32 dual channel LPDDR4 SDP without ECC interface are shown in the following table. The table shows which FPGA memory interface signals connect to which pins on the LPDDR4 device, along with any required termination.

Table 45: Signal Connection Matrix for x32 Dual Channel LPDDR4 SDP without ECC

FPGA Pins	LPDDR4 Pins	PCB Termination at Far End
Clock Signals		
CK0_P CK0_N	CK_t_A CK_c_A	None
CK1_P CK1_N	CK_t_B CK_c_B	None
Address and Command Signals		
A[5:0]	CA[5:0]_A	None
A[15:10]	CA[5:0]_B	None
Control Signals (CKE, CS, and ODT¹)		
CKE0	CKE_A CKE_B	160Ω to V _{DDQ} / 160Ω to GND
CS0	CS_A CS_B	None
	ODT_CA_A	Direct connect to V _{DD2}
	ODT_CA_B	Direct connect to V _{DD2}
Data Signals		
DQ[15:0]	DQ[15:0]_A	None

Table 45: Signal Connection Matrix for x32 Dual Channel LPDDR4 SDP without ECC (cont'd)

FPGA Pins	LPDDR4 Pins	PCB Termination at Far End
DQ[31:16]	DQ[15:0]_B	None
DM0	DMI0_A	None
DM1	DMI1_A	None
DM2	DMI0_B	None
DM3	DMI1_B	None
DQS0_P/ DQS0_N	DQS0_t_A/ DQS0_c_A	None
DQS1_P/ DQS1_N	DQS1_t_A/ DQS1_c_A	None
DQS2_P/ DQS2_N	DQS0_t_B/ DQS0_c_B	None
DQS3_P/ DQS3_N	DQS1_t_B/ DQS1_c_B	None
Reset		
RESET_N	RESET_n	4.7 kΩ to GND
-	ZQ0	240Ω to V _{DDQ}
ZQ	-	240Ω to GND

Notes:

1. The FPGA ODT pins should be left unconnected. The ODT pin(s) of the memory device should be directly connected as specified in this table.
2. The ZQ pin on the LPDDR4 devices should be tied to V_{DDQ} through a 240Ω resistor.
3. Actual signal list might vary based on configuration.

Signals and Connections for x32 Dual Channel LPDDR4 DDP without ECC

The required signals for a x32 dual channel LPDDR4 DDP non-ECC interface are shown in the following table. The table shows which FPGA memory interface signals connect to which pins on the LPDDR4 device, along with any required termination.

Table 46: Signal Connection Matrix for x32 Dual Channel LPDDR4 DDP without ECC

FPGA Pins	LPDDR4 Pins	PCB Termination at Far End
Clock Signals		
CK0_P CK0_N	CK_t_A CK_c_A	None
CK1_P CK1_N	CK_t_B CK_c_B	None
Address and Command Signals		
A[5:0]	CA[5:0]_A	None
A[15:10]	CA[5:0]_B	None

Table 46: Signal Connection Matrix for x32 Dual Channel LPDDR4 DDP without ECC (cont'd)

FPGA Pins	LPDDR4 Pins	PCB Termination at Far End
Control Signals (CKE, CS, and ODT¹)		
CKE0	CKE0_A CKE0_B	160Ω to V _{DDQ} / 160Ω to GND
CKE1	CKE1_A CKE1_B	160Ω to V _{DDQ} / 160Ω to GND
CS0	CS0_A CS0_B	None
CS1	CS1_A CS1_B	None
	ODT_CA_A	Direct connect to V _{DD2}
	ODT_CA_B	Direct connect to V _{DD2}
Data Signals		
DQ[15:0]	DQ[15:0]_A	None
DQ[31:16]	DQ[15:0]_B	None
DM0	DMI0_A	None
DM1	DMI1_A	None
DM2	DMI0_B	None
DM3	DMI1_B	None
DQS0_P DQS0_N	DQS0_t_A DQS0_c_A	None
DQS1_P DQS1_N	DQS1_t_A DQS1_c_A	None
DQS2_P DQS2_N	DQS0_t_B DQS0_c_B	None
DQS3_P DQS3_N	DQS1_t_B DQS1_c_B	None
Reset		
RESET_N	RESET_n	4.7 kΩ to GND

Notes:

1. The FPGA ODT pins should be left unconnected. The ODT pin(s) of the memory device should be directly connected as specified in this table.
2. Each ZQ pin on the LPDDR4 devices should be individually tied to V_{DDQ} through a 240Ω resistor.
3. The ZQ pin on the FPGA should be tied to GND through a 240Ω resistor.

LPDDR4 Address Copy

The PS DDR interface includes a feature called address copy to reduce loading on the address/command bus when operating in LPDDR4 mode. A[5:0] are replicated on A[15:10] to allow fewer loads on the CA bus in the following table.

Table 47: LPDDR4 Address Copy

MPSoC Pin	DRAM Pin	MPSoC Pin Copy	DRAM Pin
PS_DDR_A0	CA0_A	PS_DDR_A10	CA0_B
PS_DDR_A1	CA1_A	PS_DDR_A11	CA1_B
PS_DDR_A2	CA2_A	PS_DDR_A12	CA2_B
PS_DDR_A3	CA3_A	PS_DDR_A13	CA3_B
PS_DDR_A4	CA4_A	PS_DDR_A14	CA4_B
PS_DDR_A5	CA5_A	PS_DDR_A15	CA5_B


LPDDR4 Interface without ECC Topology and Routing Guidelines

★ IMPORTANT! To achieve the specified performance, all rules in the tables must be followed. This includes but is not limited to the trace type, width, spacing, impedance, and via count.

LPDDR4 Interface without ECC Address/Command and Control Routing

The following figure shows the address/command routing for the LPDDR4 interface without ECC.

Figure 46: Address/Command Routing for LPDDR4 without ECC



The following table shows the impedance, length, and spacing guidelines for address, command, and control signals in the LPDDR4 interface without ECC.

Table 48: Impedance, Length, and Spacing Guidelines for Address and Command Signals in LPDDR4 Interface without ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	Ω
Trace width	4.0	6.0	3.5	mil
Trace length	0.0-0.55/0.95 ¹	≤ 4.0	≤ 0.3	inches
Spacing in addr/cmd (minimum)	4.0	8.0	4.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	8.0	mil
Maximum PCB via count	2			-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 Interface without ECC Chip Select Routing

The following figure shows the chip select (CS) routing for the LPDDR4 interface without ECC.

Figure 47: Chip Select Routing for LPDDR4 without ECC



The following table shows the impedance, length, and spacing guidelines for chip select in the LPDDR4 interface without ECC.

Table 49: Impedance, Length, and Spacing Guidelines for Chip Select in LPDDR4 Interface without ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-

Table 49: Impedance, Length, and Spacing Guidelines for Chip Select in LPDDR4 Interface without ECC (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	Units
Single-ended impedance Z_0	50±10%	39±10%	52±10%	50±10%	Ω
Trace width	4.0	6.0	3.5	4.0	mil
Trace length	0.0~0.55/0.95 ¹	≤ 4.0	≤ 0.3	≤ 0.3	inches
Spacing to clock signals (minimum)	8.0	20	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	mil
Maximum PCB via count	3				-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 without ECC Clock Enable Routing

The following figure shows the clock enable (CKE) routing for an LPDDR4 interface without ECC.

Figure 48: Clock Enable Routing for LPDDR4 without ECC



The following table shows the impedance, length, and spacing guidelines for clock enable in the LPDDR4 interface without ECC.

Table 50: Impedance, Length, and Spacing Guidelines for Clock Enable in LPDDR4 Interface without ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (Stub)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	50±10%	50±10%	Ω

Table 50: Impedance, Length, and Spacing Guidelines for Clock Enable in LPDDR4 Interface without ECC (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (Stub)	Units
Trace width	4.0	6.0	3.5	4.0	4.0	mil
Trace length	0.0~0.55/0.95 ¹	≤ 4.0	≤ 0.3	≤ 0.3	≤ 0.5	inches
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	30	mil
Maximum PCB via count	3					-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 without ECC CK and DQS Differential Point-to-Point Routing

The following figure shows the CK and DQS differential point-to-point routing for the LPDDR4 interface without ECC.

Figure 49: CK and DQS Routing for LPDDR4 without ECC



The following table shows the impedance, length, and spacing guidelines for CK and DQS in the LPDDR4 interface without ECC.

Table 51: Impedance, Length, and Spacing Guidelines for CK and DQS in LPDDR4 Interface without ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
Clock differential impedance Z_{diff}	86±10%	76±10%	88±10%	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	3.5/3.5/3.5	mil
Trace length	≤ 0.55/0.95 ¹	≤ 4.0	≤ 0.3	inches
Spacing in addr/cmd/ctrl (minimum)	8.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	8.0	mil
Maximum PCB via count per signal	2			-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 without ECC Data (DQ and DM) Point-to-Point Routing

The following figure shows the DQ and DM point-to-point routing for the LPDDR4 interface without ECC.

Figure 50: DQ and DM Routing for LPDDR4 without ECC



UG583_c8_07_011516

The following table shows the impedance, length, and spacing guidelines for DQ and DM in the LPDDR4 interface without ECC.

Table 52: Impedance, Length, and Spacing Guidelines for DQ and DM in LPDDR4 Interface without ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	Ω

Table 52: Impedance, Length, and Spacing Guidelines for DQ and DM in LPDDR4 Interface without ECC (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace width	4.0	6.0	3.5	mil
Trace length	$\leq 0.55/0.95^1$	≤ 4.0	≤ 0.3	inches
Spacing in byte (including DQS) (minimum)	4.0	8.0	4.0	mil
Spacing byte to byte (minimum)	8.0	30	8.0	mil
Spacing to other group signals (minimum)	8.0	20	8.0	mil
Maximum PCB via count	2			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 without ECC Length and Skew Constraints

The constraints requirement for each signal group in the LPDDR4 without ECC memory interface consists of two parts:

- Maximum length constraints
- Skew constraints

The maximum length constraints are shown in the following table.

Table 53: LPDDR4 without ECC Maximum Length Constraints

Signal Group	Reference Figure	Maximum Length Constraints (inches)
Address/command P0+L0+L1+L2	Figure 46: Address/Command Routing for LPDDR4 without ECC	6.2
CS P0+L0+L1+L2+L3	Figure 47: Chip Select Routing for LPDDR4 without ECC	6.8
CKE P0+L0+L1+L2+L3	Figure 48: Clock Enable Routing for LPDDR4 without ECC	7.9
Data signals P0+L0+L1+L2	Figure 50: DQ and DM Routing for LPDDR4 without ECC	6.2

The skew constraints are listed in the following table.

Table 54: LPDDR4 without ECC Skew Constraints

Signal Group	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
Address/command/control to CK (A) Address/command/control to CK (B)	MPSoC to memory device	± 8	± 47

Table 54: LPDDR4 without ECC Skew Constraints (cont'd)

Signal Group	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
Data (DQ/DM) to DQS (A) Data (DQ/DM) to DQS (B)	MPSoC to memory device	±5	±29
CK (A) to DQS0/1 (A) CK (B) to DQS0/1 (B)	MPSoC to memory device	-500 to 2500	-2950 to 14750
DQ/DM (slowest to fastest) (A) DQ/DM (slowest to fastest) (B)	MPSoC to memory device	5	29
CK_T and CK_C (A) CK_T and CK_C (B)	MPSoC to memory device	2	11
DQS_T and DQS_C (A) DQS_T and DQS_C (B)	MPSoC to memory device	2	11

Notes:

- For skew specifications, refer to items 3–8 in [General Memory Routing Guidelines](#).
- The signal reset_n is not required to meet the skew constraints in this table.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

PCB Guidelines for LPDDR4 Memories with ECC (PS)

Overview

This section provides electrical design guidelines for 32-bit LPDDR4 interfaces with ECC.



IMPORTANT! All guidelines in this section must be followed to achieve the maximum data rates specified for the LPDDR4 interface.

LPDDR4 x32 with ECC Memory Interface Signals and Connections

The required signals for a x32 dual channel LPDDR4 single-rank with ECC interface are shown in the following table. The table shows how to connect the FPGA memory interface signals to the correct pins on the LPDDR4 devices, along with any required termination.

Table 55: Signal Connection Matrix for x32 Dual Channel LPDDR4 Single-Rank with ECC

FPGA Pins	LPDDR4 Pins MAIN DEVICE	LPDDR4 Pins ECC DEVICE	PCB Termination at Far End
Clock Signals			
CK0_P CK0_N	CK_t_A CK_c_A		40Ω to GND 40Ω to GND
CK1_P CK1_N	CK_t_B CK_c_B	CK_t_A CK_c_A	40Ω to GND 40Ω to GND
Address and Command Signals			
A[5:0]	CA[5:0]_A		40Ω to GND
A[15:10]	CA[5:0]_B	CA[5:0]_A	40Ω to GND
Control Signals (CKE, CS, and ODT¹)			
CKE0	CKE_A CKE_B	CKE_A	160Ω to V _{DDQ} / 160Ω to GND
CS0	CS_A CS_B	CS_A	40Ω to GND
	ODT_CA_A	ODT_CA_A	Direct connect to GND
	ODT_CA_B		Direct connect to GND
Data Signals			
DQ[15:0]	DQ[15:0]_A		None
DQ[31:16]	DQ[15:0]_B		None
DQ[71:64]		DQ[7:0]_A	None
DM0	DMI0_A		None
DM1	DMI1_A		None
DM2	DMI0_B		None
DM3	DMI1_B		None
DM8		DMI0_A	None
DQS0_P DQS0_N	DQS0_t_A DQS0_c_A		None
DQS1_P DQS1_N	DQS1_t_A DQS1_c_A		None
DQS2_P DQS2_N	DQS0_t_B DQS0_c_B		None
DQS3_P DQS3_N	DQS1_t_B DQS1_c_B		None
DQS8_P DQS8_N		DQS0_t_A DQS0_c_A	None

Table 55: Signal Connection Matrix for x32 Dual Channel LPDDR4 Single-Rank with ECC (cont'd)

FPGA Pins	LPDDR4 Pins MAIN DEVICE	LPDDR4 Pins ECC DEVICE	PCB Termination at Far End
Reset			
RESET_N	RESET_n	RESET_n	4.7 kΩ to GND

Notes:

1. The FPGA ODT pins should be left unconnected. The ODT pin(s) of the memory device should be directly connected as specified in this table.
2. Unused inputs to the LPDDR4 ECC device can be grounded through a common 100Ω resistor.
3. Each ZQ pin on the LPDDR4 devices should be individually tied to V_{DDQ} through a 240Ω resistor.
4. The ZQ pin on the FPGA should be tied to GND through a 240Ω resistor.
5. Actual signal list might vary based on configuration.

The required signals for a x32 dual channel LPDDR4 dual-rank with ECC interface are shown in the following table. The table shows how to connect the FPGA memory interface signals to the correct pins on the LPDDR4 devices, along with any required termination.

Table 56: Signal Connection Matrix for x32 Dual Channel LPDDR4 Dual-Rank with ECC

FPGA Pins	LPDDR4 Pins Main Device	LPDDR4 Pins ECC Device	PCB Termination at Far End
Clock Signals			
CK0_P	CK_t_A		40Ω to GND
CK0_N	CK_c_A		40Ω to GND
CK1_P	CK_t_B	CK_t_A	40Ω to GND
CK1_N	CK_c_B	CK_c_A	40Ω to GND
Address and Command Signals			
A[5:0]	CA[5:0]_A		40Ω to GND
A[15:10]	CA[5:0]_B	CA[5:0]_A	40Ω to GND
Control Signals (CKE, CS, and ODT¹)			
CKE0	CKE0_A CKE0_B	CKE0_A	160Ω to V _{DDQ} / 160Ω to GND
CKE1	CKE1_A CKE1_B	CKE1_A	160Ω to V _{DDQ} / 160Ω to GND
CS0	CS0_A CS0_B	CS0_A CS0_B	40Ω to GND
CS1	CS1_A CS1_B	CS1_A CS1_B	40Ω to GND
	ODT_CA_A	ODT_CA_A	Direct connect to GND
	ODT_CA_B		Direct connect to GND
Data Signals			
DQ[15:0]	DQ[15:0]_A		None
DQ[31:16]	DQ[15:0]_B		None

Table 56: Signal Connection Matrix for x32 Dual Channel LPDDR4 Dual-Rank with ECC (cont'd)

FPGA Pins	LPDDR4 Pins Main Device	LPDDR4 Pins ECC Device	PCB Termination at Far End
DQ[71:64]		DQ[7:0]_A	None
DM0	DMI0_A		None
DM1	DMI1_A		None
DM2	DMI0_B		None
DM3	DMI1_B		None
DM8		DMI0_A	None
DQS0_P DQS0_N	DQS0_t_A DQS0_c_A		None
DQS1_P DQS1_N	DQS1_t_A DQS1_c_A		None
DQS2_P DQS2_N	DQS0_t_B DQS0_c_B		None
DQS3_P DQS3_N	DQS1_t_B DQS1_c_B		None
DQS8_P DQS8_N		DQS0_t_A DQS0_c_A	None
Reset			
RESET_B	RESET_n	RESET_n	4.7 kΩ to GND

Notes:

1. The FPGA ODT pins should be left unconnected. The ODT pin(s) of the memory device should be directly connected as specified in this table.
2. Unused inputs to the LPDDR4 ECC device can be grounded through a common 100Ω resistor.
3. Each ZQ pin on the LPDDR4 devices should be individually tied to V_{DDQ} through a 240Ω resistor.
4. The ZQ pin on the FPGA should be tied to GND through a 240Ω resistor.
5. Actual signal list might vary based on configuration.

LPDDR4 Address Copy

The PS DDR interface includes a feature called address copy to reduce loading on the address/command bus when operating in LPDDR4 mode. A[5:0] are replicated on A[15:10] to allow fewer loads on the CA bus in the following table.

Table 57: LPDDR4 Address Copy

MPSoC Pin	DRAM Pin	MPSoC Pin Copy	DRAM Pin
PS_DDR_A0	CA0_A	PS_DDR_A10	CA0_B
PS_DDR_A1	CA1_A	PS_DDR_A11	CA1_B
PS_DDR_A2	CA2_A	PS_DDR_A12	CA2_B
PS_DDR_A3	CA3_A	PS_DDR_A13	CA3_B
PS_DDR_A4	CA4_A	PS_DDR_A14	CA4_B

Table 57: LPDDR4 Address Copy (cont'd)

MPSoC Pin	DRAM Pin	MPSoC Pin Copy	DRAM Pin
PS_DDR_A5	CA5_A	PS_DDR_A15	CA5_B


LPDDR4 with ECC Topology and Routing Guidelines

★ IMPORTANT! To achieve the specified performance, all rules in the tables must be followed. This includes but is not limited to the trace type, width, spacing, impedance, and via count.

LPDDR4 with ECC Address/Command and Control Routing (CA_A)

The following figure shows the CA_A address/command routing for the LPDDR4 interface with ECC. The R_{TT} termination exists to terminate the fly-by route.

Figure 51: CA_A Address/Command Routing for LPDDR4 with ECC



The following table shows the impedance, length, and spacing guidelines for CA_A address/command signals in the LPDDR4 interface with ECC.

Table 58: Impedance, Length, and Spacing Guidelines for CA_A Address/Command Signals in LPDDR4 Interface with ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (Stub)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	50±10%	Ω
Trace width	4.0	6.0	3.5	4.0	mil
Trace length	≤ 0.55/0.95 ¹	≤ 4.0	≤ 0.3	≤ 0.5	inches
Spacing to addr/cmd (minimum)	4.0	6.0	4.0	6.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	mil
Maximum PCB via count	3				-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 with ECC Address/Command and Control Routing (CA_B)

The following figure shows the CA_B address/command routing for the LPDDR4 interface with ECC. The R_{TT} termination exists to terminate the fly-by route.

Figure 52: CA_B Address/Command Routing for LPDDR4 with ECC



The following table shows the impedance, length, and spacing guidelines for CA_B address/command signals in the LPDDR4 interface with ECC.

Table 59: Impedance, Length, and Spacing Guidelines for CA_B Address/Command Signals in LPDDR4 Interface with ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (Chip to Chip)	L4 (Stub)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	52±10%	50±10%	Ω
Trace width	4.0	6.0	3.5	3.5	4.0	mil
Trace length	$\leq 0.55/0.95^1$	≤ 4.0	≤ 0.3	≤ 1.5	≤ 0.5	inches
Spacing in addr/cmd (minimum)	4.0	6.0	4.0	6.0	6.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	30	mil
Maximum PCB via count	4					-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 with ECC Chip Select Routing

The following figure shows the chip select (CS0/CS1) routing for the LPDDR4 interface with ECC. The R_{TT} termination exists to terminate the fly-by route.

Figure 53: CS0/CS1 Routing for LPDDR4 with ECC



The following table shows the impedance, length, and spacing guidelines for CS0/CS1 routing in the LPDDR4 interface with ECC.

Table 60: Impedance, Length, and Spacing Guidelines for CS0/CS1 in LPDDR4 Interface with ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (Channel to Channel)	L4 (Chip to Chip)	L5 (Stub)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	52±10%	52±10%	50±10%	Ω
Trace width	4.0	6.0	3.5	3.5	3.5	4.0	mil
Trace length	≤ 0.55/0.95 ¹	≤ 4.0	≤ 0.3	≤ 0.5	≤ 1.5	≤ 0.5	inches
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	30	30	mil
Maximum PCB via count	4						-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 with ECC Clock Enable Routing

The following figure shows the clock enable (CKE0/CKE1) routing for the LPDDR4 interface with ECC.

Figure 54: CKE0/CKE1 Routing for LPDDR4 with ECC



UG583_c8_12_102521

The following table shows the impedance, length, and spacing guidelines for CKE0/CKE1 in the LPDDR4 interface with ECC.

Table 61: Impedance, Length, and Spacing Guidelines for CKE0/CKE1 in LPDDR4 Interface with ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (Channel to Channel)	L4 (Chip to Chip)	L5 (Stub)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	52±10%	52±10%	50±10%	Ω
Trace width	4.0	6.0	3.5	3.5	3.5	4.0	mil
Trace length	0.0-0.55/0.95 ¹	≤ 4.0	≤ 0.3	≤ 0.5	≤ 1.5	≤ 0.5	inches
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	30	30	mil
Maximum PCB via count	5						-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 with ECC CK0 Differential Point-to-Point Routing

The following figure shows the CK0 differential point-to-point routing for the LPDDR4 interface with ECC.

Figure 55: CK0 Routing for LPDDR4 with ECC



The following table shows the impedance, length, and spacing guidelines for CK0 in the LPDDR4 interface with ECC.

Table 62: Impedance, Length, and Spacing Guidelines for CK0 in LPDDR4 Interface with ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (Stub)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-
Clock differential impedance Z_{diff}	$86 \pm 10\%$	$76 \pm 10\%$	$88 \pm 10\%$	$93 \pm 10\%$	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	3.5/3.5/3.5	4.0/6.0/4.0	mil
Trace length	$\leq 0.55/0.95^1$	≤ 4.0	≤ 0.3	≤ 0.5	inches
Spacing in addr/cmd/ctrl (minimum)	8.0	20	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	mil
Maximum PCB via count per signal	3				-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 with ECC CK1 Differential Fly-by Routing

The following figure shows the CK1 differential fly-by routing for the LPDDR4 interface with ECC.

Figure 56: CK1 Routing for LPDDR4 with ECC



The following table shows the impedance, length, and spacing guidelines for CK1 in the LPDDR4 interface with ECC.

Table 63: Impedance, Length, and Spacing Guidelines for CK1 in LPDDR4 Interface with ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (Chip to Chip)	L4 (Stub)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Clock differential impedance Z_{diff}	86±10%	76±10%	88±10%	93±10%	93±10%	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	3.5/3.25/3.5	3.5/4.0/3.5	4.0/6.0/4.0	mil
Trace length	≤ 0.55/0.95 ¹	≤ 4.0	≤ 0.3	≤ 1.5	≤ 0.5	inches
Spacing in addr/cmd/ctrl (minimum)	8.0	20	8.0	20	20	mil

Table 63: Impedance, Length, and Spacing Guidelines for CK1 in LPDDR4 Interface with ECC (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (Chip to Chip)	L4 (Stub)	Units
Spacing to other group signals (minimum)	8.0	30	8.0	30	30	mil
Maximum PCB via count per signal	4					-


Notes:

- See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 with ECC Data (DQ) Point-to-Point Routing

The following figure shows the DQ point-to-point routing for the LPDDR4 interface with ECC.

Figure 57: DQ Routing for LPDDR4 with ECC



The following table shows the impedance, length, and spacing guidelines for DQ in the LPDDR4 interface with ECC.

Table 64: Impedance, Length, and Spacing Guidelines for DQ in LPDDR4 Interface with ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	$50 \pm 10\%$	$39 \pm 10\%$	$52 \pm 10\%$	Ω
Trace width	4.0	6.0	3.5	mil
Trace length	$\leq 0.55/0.95^1$	≤ 4.0	≤ 0.3	inches
Spacing in byte (including DQS) (minimum)	4.0	8.0	4.0	mil

Table 64: Impedance, Length, and Spacing Guidelines for DQ in LPDDR4 Interface with ECC (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Spacing byte to byte (minimum)	8.0	20	4.0	mil
Spacing to other group signals (minimum)	8.0	30	8.0	mil
Maximum PCB via count	2			-


Notes:

- See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 with ECC DQS Differential Point-to-Point Routing

The following figure shows the DQS point-to-point routing for the LPDDR4 interface with ECC.

Figure 58: DQS Routing for LPDDR4 with ECC



The following table shows the impedance, length, and spacing guidelines for DQS in the LPDDR4 interface with ECC.

Table 65: Impedance, Length, and Spacing Guidelines for DQS in LPDDR4 Interface with ECC

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
Clock differential impedance Z_{diff}	86±10%	76±10%	88±10%	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	3.5/3.5/3.5	mil
Trace length	≤ 0.55/0.95 ¹	≤ 4.0	≤ 0.3	inches
Spacing in addr/cmd/ctrl (minimum)	4.0	20	4.0	mil
Spacing to other group signals (minimum)	8.0	30	8.0	mil
Maximum PCB via count per signal	2			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR4 with ECC Length and Skew Constraints

The constraints requirement for each signal group in the LPDDR4 with ECC memory interface consists of two parts:

- Maximum length constraints
- Skew constraints

The maximum length constraints are shown in following table.

Table 66: LPDDR4 with ECC Maximum Length Constraints

Signal Group	Reference Figure	Maximum Length Constraints (inches)
CA_A Address/command P0+L0+L1+L2+L3	Figure 51: CA_A Address/Command Routing for LPDDR4 with ECC	7.0
CA_B Address/command P0+L0+L1+L2+L3+L4	Figure 52: CA_B Address/Command Routing for LPDDR4 with ECC	8.5
CS0/CS1 P0+L0+L1+L2+L3+L4+L5	Figure 53: CS0/CS1 Routing for LPDDR4 with ECC	9.0
CKE0/CKE1 P0+L0+L1+L2+L3+L4+L5	Figure 54: CKE0/CKE1 Routing for LPDDR4 with ECC	9.0
Data signals P0+L0+L1+L2	Figure 57: DQ Routing for LPDDR4 with ECC	6.2

The skew constraints are listed in following table.

Table 67: LPDDR4 with ECC Skew Constraints

Signal Group	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
Address/command/control to CK (A) ¹ Address/command/control to CK (B) ¹	MPSoC to memory device	±8	±47
Data (DQ/DM) to DQS (A) Data (DQ/DM) to DQS (B)	MPSoC to memory device	±5	±29
CK (A) to DQS0/1 (A) CK (B) to DQS0/1 (B)	MPSoC to memory device	-500 to 2500	-2950 to 14750
DQ/DM (slowest to fastest) (A) DQ/DM (slowest to fastest) (B)	MPSoC to memory device	5	29
CK_T and CK_C (A) CK_T and CK_C (B)	MPSoC to memory device	2	11
DQS_T and DQS_C (A) DQS_T and DQS_C (B)	MPSoC to memory device	2	11

Notes:

1. Consider CS and CKE to be part of the A group.
2. For skew specifications, refer to items 3–8 in [General Memory Routing Guidelines](#).
3. The signal reset_n is not required to meet the skew constraints in this table.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

PCB Routing Guidelines for LPDDR4 Memories in High-Density Interconnect Boards

The LPDDR4 routing guidelines for AMD devices that are placed on high-density interconnect (HDI) boards are different from those for standard printed circuit board technologies. The recommended routing of LPDDR4 signals between the FPGA and DRAM device is shown in the following figure, with a reference PCB shown in [Figure 60: Reference PCB with Routing Areas Denoted](#). Routing begins from the BGA ball with a short stub to a blind via (top layer to inner layer) which travels to layer three. The routing from the BGA ball to where it is no longer under the BGA is referred to as the *breakout area* and can typically be the most difficult area to route due to tight spacing of pins and vias. After traveling outside of the breakout area into the *main routing area*, the signal then travels under the DRAM breakout area, through a blind via, to the DRAM pin. The following figure also shows how blind and buried vias (vias between inner layers) can be used in cases where spacing is too tight for just one blind via.

Figure 59: Examples of Stripline Routing for HDI LPDDR4 Signals





Figure 60: Reference PCB with Routing Areas Denoted



X25152-030221

Refer to the following table for the physical and timing rules.

Table 68: Routing Guidelines for HDI LPDDR4 Signals

Parameter	Value	Notes
Bitrate	1066 Mb/s	
PCB thickness	37 mil ± 10%	Typical board thickness for HDI. Assumes FPGA and memory component on same side of PCB.
FPGA/DRAM component placement	Top layer	
Number of routing layers to route LPDDR4 signals	3 (for FGAA530 package)	Top, Layer 3, Layer 5
Maximum distance between FPGA ball and blind via	70 mil	
Maximum distance between DRAM ball and blind via	300 mil	
Maximum distance between FPGA ball and DRAM ball	2000 mil	
Via type	Blind	Maximum 4 (see next row)
Blind/buried via routing (see Figure 59: Examples of Stripline Routing for HDI LPDDR4 Signals)	Top-L3-Top	2 Vias
	Top-L5-Top	2 Vias
	Top-L3-L5-L3-Top	4 Vias
Maximum skew: P/N	±2 ps	DQS, CK signals
Maximum skew: DQ to DQS	±5 ps	
Maximum skew: CAC to CK	±8 ps	
Trace Impedance (single-ended)	50Ω ± 10%	
Trace Impedance (differential)	100Ω ± 10%	
Minimum spacing: CAC Signals	2.5 × H ¹	Within same channel
Minimum spacing: CAC Signals to CK	4.5 × H	Within same channel
Minimum spacing: CAC/CK to DQ/DQS	7.5 × H ¹	Within same channel
Minimum spacing: DQ Signals	2.5 × H ¹	Within same byte (main)
Minimum spacing: DQ to DQS	2.5 × H ¹	Within same byte (main)
Minimum spacing: DQ/DQS Signals	4.5 × H ¹	To different bytes
Minimum spacing: DQ/DQS Signals	4.5 × H ¹	To other signals in the same channel
Minimum spacing	7.5 × H ¹	To signals in different channels or interfaces

Notes:

1. H is the distance between the routing layer and closest reference plane.

PCB Guidelines for LPDDR3 SDRAM (PL and PS)

Overview

This section provides electrical design guidelines for typical LPDDR3 SDRAM interfaces for the PS architecture in Zynq UltraScale+ MPSoCs based on system-level signal integrity simulation. A 32-bit wide one 8 GB SDP SDRAM system and a 64-bit wide two 8 GB SDP SDRAM system are used as typical design examples. For both cases, clock, data, and address/command signals are point-to-point connections. For a x32 single component system, control signals are also point-to-point connections. In a x64 two-component system, they are connected in fly-by topology from the FPGA to two SDRAMs.

 **IMPORTANT!** All routing guidelines in this section must be followed to achieve the maximum data rates specified for the LPDDR3 SDRAM interface for typical system designs. Customers could have unique or specific designs with particular violations of some rules. In these scenarios, design or routing trade-offs have to be taken in other routing parameters to mitigate the risk. System-level channel signal integrity simulations are required to evaluate such trade-offs.

LPDDR3 SDRAM Interface Signal Description

The LPDDR3 SDRAM interface consists of clock, control, address, command, and data signals as shown in the following table.

Table 69: LPDDR3 SDRAM Interface Signal Description

Signal Name	Description
CK_t, CK_c	Differential clock
Control Signals	
CKE[1:0]	Clock enable
CS[1:0]_n	Chip select
ODT	DQ ODT control
Command/Address Signals	
CA[9:0]	Command/address bus
Data Signals	
DQ[32:0]	Data bus
DQS[3:0]_t/c	Differential data strobe
DM[3:0]	Input data mask

Notes:

- Actual signal list might vary based on configuration.

LPDDR3 Address Copy

The PS DDR interface includes a feature called address copy to reduce loading on the address/command bus when operating in LPDDR3 mode. A[9:0] are replicated on ACT_N, BG[0], BA[1:0], and A[15:10] to allow fewer loads on the CA bus (see the following table).

Table 70: LPDDR3 Address Copy

MPSoC Pin	DRAM Pin	MPSoC Pin Copy	DRAM Pin
PS_DDR_A0	CA0_A	PS_DDR_A10	CA0_B
PS_DDR_A1	CA1_A	PS_DDR_A11	CA1_B
PS_DDR_A2	CA2_A	PS_DDR_A12	CA2_B
PS_DDR_A3	CA3_A	PS_DDR_A13	CA3_B
PS_DDR_A4	CA4_A	PS_DDR_A14	CA4_B
PS_DDR_A5	CA5_A	PS_DDR_A15	CA5_B
PS_DDR_A6	CA6_A	PS_DDR_BA0	CA6_B
PS_DDR_A7	CA7_A	PS_DDR_BA1	CA7_B
PS_DDR_A8	CA8_A	PS_DDR_BG0	CA8_B
PS_DDR_A9	CA9_A	PS_DDR_ACT_N	CA9_B


VREFCA and VREFDQ

Both PL and PS logic use ODT for data signals, so the voltage levels for VREFCA and VREFDQ will be different. Refer to the memory vendor data sheet for specific requirements.

Design Example for x32 and x64 LPDDR3 SDRAM

To optimize the signal routing, the recommendation for a one-component 32-bit system placement is shown in the following figure.


Figure 61: Design Example for One-Component x32 LPDDR3 SDRAM



UG583_c8_27_061016

The recommendation for a two-component 64-bit system is shown in the following figure.

Figure 62: Design Example for Two-Component x64 LPDDR3 SDRAM




UG583_c8_28_040717

Topology and Routing Guidelines for LPDDR3 SDRAM

LPDDR3 SDRAM Clk Point-to-Point Termination

The following figure shows the clock point-to-point termination for LPDDR3 SDRAM.

Figure 63: Clk Point-to-Point Termination for LPDDR3 SDRAM



UG583_c8_29_061316

The following tables show the LPDDR3 SDRAM impedance, length, and spacing guidelines for clock signals for x32 and x64 systems, respectively. For LPDDR3 SDRAM, clock signals have on-board termination. There is no on-die termination available for clock signals. For both x32 and x64 systems, clock signals are point-to-point connections to achieve the highest operation frequency.

Table 71: LPDDR3 x32 SDRAM System Impedance, Length, and Spacing Guidelines for Clock Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (OBT)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	39±10%	Ω
Differential impedance Z_{0d}	86±10%	76±10%	88±10%	76±10%	Ω
Trace width/spacing/width (nominal)	4.0/4.0/4.0	6.0/6.0/6.0	3.5/3.25/3.5	6.0/6.0/6.0	mil
Trace length	0.0~0.55	1.0~4.0	0.0~0.3	0.0~0.5	inch
Spacing to clock signals (minimum)	8.0	20	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	mil
Maximum PCB via count	3				mil

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

Table 72: LPDDR3 x64 SDRAM System Impedance, Length, and Spacing Guidelines for Clock Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (OBT)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	39±10%	Ω
Differential impedance Z_{0d}	86±10%	76±10%	88±10%	76±10%	Ω
Trace width/spacing/width (nominal)	4.0/4.0/4.0	6.0/6.0/6.0	3.5/3.25/3.5	6.0/6.0/6.0	mil
Trace length	0.0~0.55	1.0~4.0	0.0~0.3	0.0~0.5	inch
Spacing to clock signals (minimum)	8.0	20	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	mil
Maximum PCB via count	3				mil


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR3 SDRAM Address, Command, and Control Point-to-Point Termination

With high-speed signaling in LPDDR3 SDRAM, point-to-point topology is used for address, command, and control signals to achieve the best signal integrity in X32 system. Each address, command, and control signal by itself is routed continuously in the same layer from device pin to far-end termination, except in breakout areas. In other words, each individual address, command, or control signal routing is not broken into routings on multiple layers. The following figure shows the address, command, and control point-to-point termination for LPDDR3 SDRAM.

Figure 64: Address, Command, and Control Point-to-Point Termination for LPDDR3 SDRAM



UG583_c8_30_060816

The following tables show the LPDDR3 SDRAM impedance, length, and spacing guidelines for address, command, and control signals for x32 and x64 systems, respectively. In the LPDDR3 DRAM package, the vertical pitch is only 0.65 mm. Due to limited space, the trace impedance in L2 could be up to 60 Ω . Address, command, and control signals have on-board termination. There is no on-die termination available. All address, command, and control signals are point-to-point connections for a x32 single-SDRAM system. For a x64 two-SDRAM system, address and command signals are point-to-point connections while control signals including CS, CKE, and ODT are connected from the FPGA to the two SDRAMs in fly-by topology.

Table 73: LPDDR3 x32 SDRAM Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signal Point-to-Point Connection

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (OBT)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	39±10%	Ω
Trace width (nominal)	4.0	6	3.5	6	mil
Trace length	0.0~0.55	1.0~4.0	0.0~0.2	0.0~0.5	inch

Table 73: LPDDR3 x32 SDRAM Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signal Point-to-Point Connection (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (OBT)	Units
Spacing in cmd/addr/ctrl signals (minimum)	4.0	8.0	4.0	4.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	mil
Maximum PCB via count	3				mil

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

Table 74: LPDDR3 x64 SDRAM Impedance, Length, and Spacing Guidelines for Address and Command Signal Point-to-Point Connection

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (OBT)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	52±10%	39±10%	Ω
Trace width (nominal)	4.0	6	3.5	6	mil
Trace length	0.0~0.55	1.0~4.0	0.0~0.2	0.0~0.5	inch
Spacing in cmd/addr/ctrl signals (minimum)	4.0	8.0	4.0	4.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	mil
Maximum PCB via count	3				mil


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. Control signals (CS/CKE/ODT) are routed under different constraints.

LPDDR3 SDRAM Control Fly-by Termination for x64 Two-SDRAM System

With high-speed signaling in LPDDR3 SDRAMs, point-to-point topology is used for address and command signals to achieve the best signal integrity. However, due to pin limits on the FPGA side, control signals are routed from the FPGA to two SDRAMs in fly-by topology in a x64 two-SDRAM system. The following figure shows the control fly-by termination for LPDDR3 SDRAM.

Figure 65: Control Fly-by Termination for LPDDR3 SDRAM



UG583_c8_31_060816

The following table shows the LPDDR3 SDRAM impedance, length, and spacing guidelines for control signals for a x64 two-SDRAM system.

Table 75: LPDDR3 SDRAM Impedance, Length, and Spacing Guidelines for Control Signals (CKE/CS/ODT)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (DRAM to DRAM)	L4 (OBT)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	$52 \pm 10\%$	$39 \pm 10\%$	$52 \pm 10\%$	$52 \pm 10\%$	$50 \pm 10\%$	Ω
Trace width (nominal)	3.5	6.0	3.5	3.5	4.0	mil
Trace length	0.0~0.55	1.2~4.0	0.0~0.3	2.0~2.4	0.0~0.5	inch
Spacing in cmd/addr/ctrl signals (minimum)	4.0	8.0	4.0	4.0	8	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	30	mil
Maximum PCB via count	4					mil


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. The entire trace length between two SDRAM control signal balls including breakout region should be between 2.0 to 2.4 inches.
3. Clock, address, and command point-to-point traces at the second SDRAM need to have an extra 0.5 inch added when compared to the control trace to compensate for control double loading.

LPDDR3 SDRAM Data Signals Point-to-Point

The following figure shows the data signals point-to-point for LPDDR3 SDRAM. For DQ/DM/DQS signals, ODT is available on both the FPGA and SDRAM side.

Figure 66: Data Signals Point-to-Point for LPDDR3 SDRAM



The following table shows the LPDDR3 SDRAM impedance, length, and spacing guidelines for data signals for both x32 and x64 systems.

Table 76: LPDDR3 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
DQ single-ended impedance Z_0	50±10%	39±10%	52±10%	Ω
DQS differential impedance Z_{diff}	86±10%	76±10%	88±10%	Ω
Trace width (nominal)	4.0	6.0	3.5	mil
Differential trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	3.5/3.25/3.5	mil
Trace length (nominal)	0.0~0.55	1.0~4.0	0.0~0.3	inch
Spacing in byte (minimum)	4.0	8.0	4.0	mil
Spacing byte to byte (minimum)	4.0	20	4.0	mil
DQ to DQS spacing (minimum)	4.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	8.0	mil
Maximum PCB via count	2			mil

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

LPDDR3 SDRAM Routing Constraints

There are two constraints requirements for each signal group in the LPDDR3 memory interface:

- Maximum length constraints
- Skew constraints

The maximum length constraints are shown in the following tables. The maximum length is counted from the FPGA, through the PCB board, to the SDRAM package.

Table 77: LPDDR3 x32 SDRAM Maximum Length Constraints

Signal Group	Reference Figure	Maximum Length Constraints (inches)
Address/command/control	Figure 64: Address, Command, and Control Point-to-Point Termination for LPDDR3 SDRAM	5.83
Data signals	Figure 66: Data Signals Point-to-Point for LPDDR3 SDRAM	5.83
Clock	Figure 63: Clk Point-to-Point Termination for LPDDR3 SDRAM	5.83

Table 78: LPDDR3 x64 SDRAM Maximum Length Constraints

Signal Group	Reference Figure	Maximum Length Constraints (inches)
Address/command/control	Figure 65: Control Fly-by Termination for LPDDR3 SDRAM	9.00
Data signals	Figure 66: Data Signals Point-to-Point for LPDDR3 SDRAM	5.83
Clock	Figure 63: Clk Point-to-Point Termination for LPDDR3 SDRAM	9.00

The skew constraints are shown in the following table.

Table 79: LPDDR3 SDRAM Memory Skew Constraints

Signal Group	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
Address/command/control to clock ¹	FPGA to memory device	±4	±23
Data to DQS	FPGA to memory device	±5	±29
DQS_p and DQS_n	FPGA to memory device	2	12
Clock to strobe	FPGA to memory device	-149 to 1,796	-879 to 10,600

Notes:


1. This does not take into account the extra 0.5 in trace on the clock, address, and command signals in a x64 system.
2. Constraints are referred to signals within the same SDRAM.
3. For skew specifications, refer to items 3-8 in [General Memory Routing Guidelines](#).



IMPORTANT! Package routing length must be included in both maximum length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values.

Due to the extra loading at the second DRAM in the control bus, additional trace length is needed in address, command, and clock traces at the second SDRAM in a x64 two-SDRAM system shown in the following figure.

Figure 67: LPDDR3 x64 Two-SDRAM System



UG583_c8_33_040717

The matching constraints for control, address, command, and clock signals are shown here:

- L1 = L3
- L2 = L3 + L4 + 0.5 in

PCB Guidelines for RLDRAM 3 Memory

Overview

This section provides electrical design guidelines for typical RLDRAM 3 memory interfaces from system-level signal integrity simulations for the UltraScale architecture. A two-component RLDRAM 3 72-bit wide interface using x36 devices is used as a typical design example in this document. Clamshell and fly-by topologies are discussed.

★ IMPORTANT! All routing guidelines in this section must be followed to achieve the maximum data rates specified for the RLDRAM 3 memory interface for typical system designs. Customers could have unique or specific designs with particular violations of some rules. In these scenarios, design or routing trade-offs have to be taken in other routing parameters to mitigate the risk. System-level channel signal integrity simulations are required to evaluate such trade-offs. It is important to read the [General Memory Routing Guidelines](#) section before continuing with this section.

RLDRAM 3 Memory Interface Signal Description

The RLDRAM 3 memory interface consists of clock, control, address, command, and data signals as shown in the following table.

Table 80: RLDRAM 3 Memory Interface Signal Description

Signal Name	Description
Clock Signals	
ck_p/n[1:0]	Command and address clock
dk_p/n[3:0]	Write data clock
qk_p/n[7:0]	Read data clock (to device)
Control Signals	
reset_n	Tie low to 4.7 kΩ resistor at far-end near RLD3 device
Address Signals	
a[20:0]	Memory address bus
ba[3:0]	Bank address
Command Signals	
cs_n	Chip select - works with ref_n and we_n to define command
ref_n	Works with cs_n and we_n to define command
we_n	Works with cs_n and ref_n to define command
Data Signals	
dq[71:0]	Data bus
dm[3:0]	Write data mask
qvld[3:0]	Read data valid
Other Signals	
VRP	240Ω to GND


Notes:

- Actual signal list might vary based on configuration.

RLDRAM 3 Memory Topology and Routing Guidelines for Clamshell and Fly-by Configurations

RLDRAM 3 Memory Clamshell Configuration with Width Expansion


Figure 68: Clamshell RLDRAM 3 Memory with Width Expansion



UG583_c2_26_081415

RLDRAM 3 Memory Fly-by Configuration with Width Expansion


Figure 69: Fly-by RLDRAM 3 Memory with Width Expansion



RLDRAM 3 Memory Address and Command Clamshell Routing and Termination

Using clamshell topology, each address and command signal is routed to opposite memory devices with matched topologies due to the mirroring ability of RLDRAM 3 devices (see the following figure).

Figure 70: Address and Command Clamshell Routing for Clamshell-mounted RLDRAM 3 Memories



The following table shows the clamshell routing impedance, length, and spacing guidelines for address and command signals.

Table 81: Clamshell Routing Impedance, Length, and Spacing Guidelines for Address and Command Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (To R _{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z ₀	50±10%	39±10%	50±10%	39±10%	Ω
Trace width	4.0	6.0	4.0	6.0	mil
Trace length	0.0~0.8/1.2 ¹	0.0~4.0	0.0~0.25	0~1.0	inches
Spacing in address and command signals (minimum)	4.0	8.0	4.0	8.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	mil
Spacing to other group signals (minimum)	8.0	30	30	30	mil
Maximum PCB via count	6				-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

RLDRAM 3 Memory Address and Command Fly-by Routing and Termination

Using fly-by topology, each address and command signal is routed to drop by each memory device and connect to its corresponding pin. (see the following figure).

Figure 71: Address, Command, and Control Fly-by Termination for RLDRAM 3 Memories



The following table shows the RLDRAM 3 memory fly-by impedance, length, and spacing guidelines for address, command, and control signals.

Table 82: RLDRAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z ₀	50±10%	36±10%	50±10%	50±10%	39±10%	Ω
Trace width	4.0	7.0	4.0	4.0	6.0	mil
Trace length	0.0~0.8/1.2 ¹	0.0~4.0	0.0~0.25	0.35~1.05	0~1.0	inches
Spacing in address, command, and control (minimum)	4.0	8.0	4.0	8.0	8.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	30	30	30	mil
Maximum PCB via count	6					-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

RLDRAM 3 Memory CK_P/N Clamshell Routing

The following figure shows the ck_p/n clamshell routing for clamshell-mounted RLDRAM 3 memories.

Figure 72: **ck_p/n Clamshell Routing for Clamshell-mounted RLDRAM 3 Memories**



The following table shows the RLDRAM 3 memory clamshell impedance, length, and spacing guidelines for ck_p/n signals.

Table 83: RLDRAM 3 Memory Clamshell Impedance, Length, and Spacing Guidelines for ck_p/n Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (To Term)	Units
Trace type	Stripline	Stripline	Stripline	Stripline	-
Clock differential impedance Z_{DIFF}	86±10%	76±10%	86±10%	76±10%	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	6.0/6.0/6.0	mil
Trace length	0.0~0.8/1.2 ¹	0.0~3.0	0.0~0.25	0~1.0	inches

Table 83: RLDRAM 3 Memory Clamshell Impedance, Length, and Spacing Guidelines for ck_p/n Signals (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (To Term)	Units
Spacing to other group signals (minimum)	8.0	30	30	20	mil
Maximum PCB via count per signal	3				-


Notes:

- See item 2 in [General Memory Routing Guidelines](#).

RLDRAM 3 Memory CK_P/N Fly-by Routing

The following figure shows the clock differential fly-by routing for RLDRAM 3 memory.

Figure 73: Clock Differential Fly-by Routing for RLDRAM 3 Memory



The following table shows the RLDRAM 3 memory fly-by impedance, length, and spacing guidelines for clock signals.

Table 84: RLDRAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Clock Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Clock differential impedance Z _{DIFF}	86±10%	76±10%	86±10%	90±10%	76±10%	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	4.0/5.0/4.0	6.0/6.0/6.0	mil
Trace length	0.0~0.8/1.2 ¹	0.0~3.0	0.0~0.25	0.35~1.05	0~1.0	inches

Table 84: RLDRAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Clock Signals (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Spacing in address, command, and control signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	30	30	30	mil
Maximum PCB via count per signal	4					-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

RLDRAM 3 Memory DK_P/N and QK_P/N Differential Point-to-Point Routing

The following figures show the RLDRAM 3 memory dk_p/n and qk_p/n differential point-to-point routing.

Figure 74: dk_p/n Point-to-Point Routing





Figure 75: qk_p/n Point-to-Point Routing



The following table shows the RLD RAM 3 memory impedance, length, and spacing guidelines for dk and qk signals.

Table 85: RLD RAM 3 Memory Impedance, Length, and Spacing Guidelines for dk and qk Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
ck/qk differential impedance Z_{DIFF}	$86 \pm 10\%$	$76 \pm 10\%$	$86 \pm 10\%$	Ω
Trace width (nominal)	4.0	6.0	4.0	mil
Differential trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	mil
Trace length (nominal)	0.0~0.8/1.2 ¹	0.0~4.0	0.0~0.25	inches
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	2			-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

RLDRAM 3 Memory Data Signals Point-to-Point

The following figure shows the RLD RAM 3 memory data signals point-to-point routing.

Figure 76: Data Signals Point-to-Point Routing



UG583_c2_33_073014

The following figure shows the RLD RAM 3 memory impedance, length, and spacing guidelines for data signals.

Table 86: RLD RAM 3 Memory Impedance, Length, and Spacing Guidelines for Data Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
dq impedance Z_0	50±10%	39±10%	50±10%	Ω
Trace width (nominal)	4.0	6.0	4.0	mil
Trace length (nominal)	0.0~0.8/1.2 ¹	0.0~4.0	0.0~0.25	inches
Spacing in byte (minimum)	4.0	8.0	4.0	mil
Spacing byte to byte (minimum)	4.0	20	4.0	mil
dq to dk/qk spacing (minimum)	4.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	2			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

RLDRAM 3 Memory Routing Constraints

There are two constraints requirements for each signal group in the RLD RAM 3 memory interface:

- Maximum delay constraints
- Skew constraints

The maximum delay constraints are shown in following table.

Table 87: RLD RAM 3 Memory Maximum Delay Constraints

Signal Group	Reference Figure	Maximum Delay Constraints (ps)
address/command	Figure 70: Address and Command Clamshell Routing for Clamshell-mounted RLD RAM 3 Memories, Figure 71: Address, Command, and Control Fly-by Termination for RLD RAM 3 Memories	1017
Data signals P0+L0+L1+L2	Figure 76: Data Signals Point-to-Point Routing	1017

The skew constraints are listed in the following table.

Table 88: RLD RAM 3 Memory Skew Constraints

Signals	Skew Constraints (ps)
dq/dm to dk_p/n	±5
dq to qk_p/n	±5
qvld to qk_p/n	±5
dk_p/n to ck_p/n	±5
addr/cmd to ck_p/n	±5
Differential signals p to n	2

Notes:

- For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.

★ IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

PCB Guidelines for QDR II+ SRAM

Overview

This section provides electrical design guidelines for typical QDR II+ SRAM interfaces from system-level signal integrity simulations for the UltraScale architecture. A two-component QDR II+ 36-bit wide interface using x18 devices is used as a typical design example in this document.

★ IMPORTANT! All guidelines in this section must be followed to achieve the maximum data rates specified for the QDR II+ SRAM interface. It is important to read the [General Memory Routing Guidelines](#) section before continuing with this section.

QDR II+ SRAM Interface Signal Description

The QDR II+ SRAM interface consists of clock, control, address, command, and data signals as shown in the following table.

Table 89: QDR II+ SRAM Interface Signal Description

Signal Name	Description
Clock Signals	
k/k_b[1:0]	Write data/address/command clock
cq/cq_b[1:0]	Read data clock
Control Signals	
zq	Output impedance match. (Do not make a NC or connect to GND.)
odt	On-die termination select
Address Signals	
a[19:0]	Memory address bus
Command Signals	
wps_n	Write port select
rps_n	Read port select
Data Signals	
d[35:0]	Data input bus
bws[3:0]	Byte write select
q[35:0]	Data output bus
Other Signals	
VRP	240Ω to GND


Notes:

- Actual signal list might vary based on configuration.

Design Example for Dual QDR II+ SRAM Devices


The following figure shows a design example for QDR II+ SRAM devices with width expansion.

Figure 77: Design Example for QDR II+ SRAM Devices with Width Expansion



The following figure shows a design example for a single QDR II+ SRAM device.

Figure 78: Design Example for Single QDR II+ SRAM Device




QDR II+ SRAM Topology and Routing Guidelines for Fly-by Configuration

QDR II+ SRAM Address and Command Fly-by Routing and Termination


Fly-by topology is used for address and command signals to achieve the best signal integrity (see the following figure).

Figure 79: Address and Command Fly-by Routing and Termination for QDR II+ SRAM



The following figure shows the QDR II+ SRAM address and command point-to-point routing and termination.

Figure 80: Address and Command Point-to-Point Routing and Termination for QDR II+ SRAM



The following table shows the QDR II+ SRAM impedance, length, and spacing guidelines for address and command signals.

Table 90: QDR II+ SRAM Impedance, Length, and Spacing Guidelines for Address and Command Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (SRAM Breakout)	L3 (Main)	L4 (To R _{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z ₀	50±10%	39±10%	50±10%	50±10%	36±10%	Ω
Trace width	4.0	6.0	4.0	4.0	7.0	mil
Trace length	0.0~0.8/1.2 ¹	0.0~3.0	0.0~0.3	0.5~0.8	0.0~0.4	inches
Spacing in address, command, and control signals (minimum)	4.0	8.0	4.0	8.0	8.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	30	30	30	mil
Maximum PCB via count	4					-

Notes:


1. See item 2 in [General Memory Routing Guidelines](#).

Additional QDR II+ SRAM Routing Guidelines

QDR II+ SRAM Topology and Routing Guidelines for T-Branch Configuration

The following figure and table show the QDR II+ SRAM address and command T-branch routing, termination topology, and routing guidelines.

Figure 81: Address and Command T-Branch Routing and Termination for QDR II+ SRAM



UG583_c2_48_102414

Table 91: Address and Command T-Branch Routing and Termination for QDR II+ SRAM

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3 (DRAM to DRAM)	L4 (to RTT)	Units
Trace type	Stripline	Stripline	Microstrip	Microstrip	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	60±10%	60±10%	39±10%	Ω
Routing layer	Upper/Lower	Upper/Lower	Surface	Surface	Upper/Lower	
Trace width	4.0	6.0	4.0	4.0	6.0	mil
Trace length	0.0~0.8/1.2 ¹	<3.0	<0.4	<0.7	<0.5	inches
Spacing in address, command, control signals (minimum)	4.0	12.0	4.0	8.0	8.0	mil
Spacing to clock signals (minimum)	8.0	12.0	8.0	12.0	12.0	mil
Spacing to other group signals (minimum)	8.0	30	30	30	30	mil
Maximum PCB via count	3					-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

QDR II+ SRAM Clock (K, K_B, and BWS) T-Branch Routing and Termination

It is preferred to have dedicated K and K_B for each SRAM, but if that is not possible, the T-branch topology is an option. The following figure and the following table define the topology and routing guidelines for QDR II+ SRAM clock signals.

Figure 82: Clock (K, K_B) and BWS T-Branch Routing and Termination for QDR II+ SRAM



UG583_c2_49_102414

Table 92: Impedance, Length, and Spacing Guidelines for Clock Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Microstrip	-
Single-ended impedance Z_0	50±10%	39±10%	60±10%	Ω
Routing layers	Upper/Lower	Upper/Lower	Upper/Lower	
Trace width	4.0	6.0	4.0	mil
Trace length	0.0~0.8/1.2 ¹	<4.0	<0.7	inches
Spacing in byte (minimum)	4.0	12.0	4.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	2			-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

QDR II+ SRAM Clock and Data Signals (d/k/k_b) Point-to-Point Routing

The following figure shows the QDR II+ SRAM clock and data signals point-to-point routing.


Figure 83: Clock and Data Signals (d/k/k_b) Point-to-Point Routing



QDR II+ SRAM Clock and Data Signals (q/cq/cq_b) Point-to-Point Routing

The following figure shows the QDR II+ SRAM clock and data signals point-to-point routing.

Figure 84: Clock and Data Signals (q/cq/cq_b) Point-to-Point Routing



UG583_c2_38_101714

The following table shows the QDR II+ SRAM impedance, length, and spacing guidelines for clock and data signals.

Table 93: QDR II+ SRAM Impedance, Length, and Spacing Guidelines for Clock and Data Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (SRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
d and q single-ended impedance Z_0	50±10%	39±10%	50±10%	Ω
Trace width (nominal)	4.0	6.0	4.0	mil
Trace length (nominal)	0.0~0.8/1.2 ¹	0.0~4.0	0.0~0.5	inches
Spacing in byte (minimum)	4.0	8.0	4.0	mil
Spacing byte to byte (minimum)	4.0	20	4.0	mil
CQ_P/N to other spacing	4.0	12.0	4.0	mil
K_P/N to other spacing	4.0	12.0	4.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	2			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

QDR II+ SRAM Routing Constraints

There are two constraints requirements for each signal group in the QDR II+ SRAM interface:

- Maximum length/delay constraints
- Skew constraints

The maximum length/delay constraints are shown in the following table.

Table 94: QDR II+ SRAM Maximum Length/Delay Constraints

Signal Group	Reference Figure	Maximum Length/Delay Constraints
address/command	Figure 79: Address and Command Fly-by Routing and Termination for QDR II+ SRAM	6.0 inches 1017 ps
Data signals P0+L0+L1+L2	Figure 84: Clock and Data Signals (q/cq/cq_b) Point-to-Point Routing	6.0 inches 1017 ps



IMPORTANT! Address and data signals must match lengths to each respective QDR II+ SRAM device.

The skew constraints are listed in the following table.

Table 95: QDR II+ SRAM Skew Constraints

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
Data to clock (d to k, q to cq)	±6	±35
Data	±5	±29
T legs	±2	±12
CMD/ADDR (CA)	±6	±35
CMD/ADDR to CLK (CLK point-to-point and CA fly-by)	±6	±35
CMD/ADDR to CLK (CLK t-branch and CA fly-by)	±34	±200
CMD/ADDR to CLK (CLK t-branch and CA t-branch)	±6	±35
k to k_b	2	12
q to q_b	2	12

Notes:

- For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

PCB Guidelines for QDR-IV SRAM

Overview

This section provides electrical design guidelines for typical QDR-IV SRAM interfaces from system-level signal integrity simulations for the UltraScale architecture. A one-component QDR-IV 36-bit wide interface is used as a typical design example in this document.

 **IMPORTANT!** All guidelines in this section must be followed to achieve the maximum data rates specified for the QDR-IV SRAM interface. It is important to read the [General Memory Routing Guidelines](#) section before continuing with this section.

QDR-IV SRAM Interface Signal Description

The QDR-IV SRAM interface consists of clock, control, address, command, and data signals as shown in the following table.

Table 96: QDR-IV SRAM Interface Signal Description

Signal Name	Description
Clock Signals	
CK/CK#	Address/command clock
DK/DK#	Write data clock
QK/QK#	Read data clock
Control Signals	
ZQ	Output impedance match
Address Signals	
a[19:0]	Memory address bus (20:0 if using x18 data widths)
Command Signals	
LD#	Synchronous load input
Data Signals	
DQ[35:0]	Data bus
Other Signals	
VRP	240Ω to GND


Notes:

- Actual signal list might vary based on configuration.

Design Example for Dual QDR-IV SRAM Devices

The following figure shows a design example for dual QDR-IV SRAM devices.

Figure 85: Design Example for QDR-IV SRAM Devices



QDR-IV Topology and Routing Guidelines

★ IMPORTANT! To achieve the specified performance, all rules in the tables must be followed. This includes but is not limited to the trace type, width, spacing, impedance, and via count.

QDR-IV SRAM DQ/Address/Command Routing

The following figure and table define topology and routing guidelines for QDR-IV SRAM DQ, address, and command signals.

Figure 86: DQ/Address/Command Routing for QDR-IV SRAM




Table 97: QDR-IV SRAM Impedance, Length, and Spacing Guidelines for DQ/Address/Command Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	50±10%	39±10%	50±10%	Ω
Trace width	4.0	6.0	4.0	mil
Trace length	0.0~0.8/1.2 ¹	0.0~4.0	0.0~0.1	inches
Spacing in address, command, and control signals (minimum)	4.0	8.0	4.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	4 (addr/cmd) 2 DQ			-


Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

QDR-IV SRAM CK/DK/QK Routing and Termination


The following figures and table define topology and routing guidelines for QDR-IV SRAM clock signals.

Figure 87: DK and CK Routing for QDR-IV SRAM



UG583_c2_55_041715

Figure 88: QK Routing for QDR-IV SRAM



UG583_c2_56_041715

Table 98: QDR-IV SRAM Impedance, Length, and Spacing Guidelines for CK/DK/QK Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
Clock differential impedance Z_{DIFF}	$86 \pm 10\%$	$76 \pm 10\%$	$86 \pm 10\%$	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	mil
Trace length	0.0~0.8/1.2 ¹	0.0~3.0	0.0~0.1	inches
Spacing in address, command, and control signals (minimum)	8.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count per signal		2		-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).

QDR-IV SRAM Routing Constraints

There are two constraints requirements for each signal group in the QDR-IV SRAM interface:

- Maximum length/delay constraints
- Skew constraints

The maximum length/delay constraints are shown in the following table.

Table 99: QDR-IV SRAM Maximum Length/Delay Constraints

Signal Group	Reference Figure	Maximum Length/Delay Constraints
Address/command P0+L0+L1+L2	Figure 86: DQ/Address/Command Routing for QDR-IV SRAM	6.0 inches 1017 ps
Data Signals P0+L0+L1+L2	Figure 86: DQ/Address/Command Routing for QDR-IV SRAM	6.0 inches 1017 ps

The skew constraints are listed in the following table.

Table 100: QDR-IV SRAM Skew Constraints

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
DQ to clock (DQ to DK, DQ to QK)	±5	±29
Address/command to CK	±6	±35
DQ (slowest to fastest)	2	12
Address/command (slowest to fastest)	2	12
Clock to clock# (CK/DK/QK)	2	12
CK to DK	50	295
CK to QK	50	295

Notes:

- For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

PCB Guidelines for Memory Interfaces in Spartan UltraScale+ Devices

Overview

The Spartan UltraScale+ FPGA provides solutions for interfacing with the following memory architectures:

- DDR4
- LPDDR4x
- LPDDR5/5x

Before the individual architecture guidelines is a section on required guidelines that apply to all of the above memory interfaces.


Required Memory Routing Guidelines for All Interfaces

The following list contains guidelines that apply to all memory interfaces.

1. Include package delay in routing constraints when determining signal trace lengths unless otherwise specified. When minimum and maximum values are available for the package delay, use the midpoint/average between the minimum and maximum values.
2. DQ and DQS signals in the same byte group should be routed on the same layer from Spartan UltraScale+ device to DRAM. Include the data mask (DM) in the byte group as applicable.

3. For fly-by routing, address, command, and control signals can be routed on different layers, but it is recommended to use as few as possible. Do not route any individual signal on more than two layers to minimize inductive loops that can lead to crosstalk issues. Any signal layer switching via needs to have one ground via within a 50 mil radius.
4. If the system clock is connected to a bank that is also used for memory interfaces, the incoming clock signals must be biased so that they adhere to the signal level requirements of the I/O standard in the bank. Refer to *UltraScale Architecture SelectIO Resources User Guide (UG571)* for specific requirements, as well as [Answer Record 76062](#). The following figure shows the biasing structure from those reference documents for a DDR4 use case. AMD also recommends using DQS_BIAS with the unpopulated bias circuitry in place as a fallback option.


Figure 89: AC-Coupled with DC-Biased Differential Clock Input



X22485-103020


5. Signal lines must be routed over a solid reference plane. Avoid routing over voids, as shown in the following figure.

Figure 90: Signal Routing Over Solid Reference Plane




6. Avoid routing over reference plane splits, as shown in the following figure.

Figure 91: Signal Routing Over Reference Plane Split



7. Keep the routing at least 30 mils away from the reference plane and void edges with the exception of breakout regions, as shown in the following figure.

Figure 92: Breakout Region Routing



- Use chevron-style routing to allow for ground stitch vias. [Figure 93: Example of Ground Stitching \(Fly-by\)](#) shows recommended routing for fly-by configurations, while [Figure 94: Example of Ground Stitching \(Clamshell\)](#) Red: Power, Green: Ground shows recommended routing to accommodate ground stitch vias in a more congested clamshell configuration.

Figure 93: Example of Ground Stitching (Fly-by)






Figure 94: Example of Ground Stitching (Clamshell) Red: Power, Green: Ground



The following figure shows simulated eye diagrams for a DDR4 command/address/control bit with and without ground stitching vias. The simulation on the left shows an eye height of 180 mV with ground stitch vias, while the simulation on the right shows an eye height of only 99 mV when not utilizing ground stitch vias.

Figure 95: Simulations With and Without Ground Stitching Vias



- Add ground vias as much as possible around the edges and inside the device (FPGA, memory component) to make a better ground return path for signals and power, especially corners. Corner or edge balls are generally less populated as grounds.

- For address/command/control VTT termination, every four termination resistors should be accompanied by one 0.1 μ F capacitor, physically interleaving among resistors, as shown in the following figure. Refer to the memory vendor's data sheet for specifications regarding noise limits on the address/command/control VTT lines.

Figure 96: Schematic Example of VTT Resistor and Capacitor Connections






Figure 97: Example of V_{TT} Termination Placement

11. For clamshell configurations that use address mirroring, ensure that both chip select lines have adequate decoupling at their terminations as well as sufficient plane/trace thicknesses to/from V_{TT} .
12. Ensure that all PCB traces are within a 50 mil radius from a ground via to ensure impedance continuity through the PCB.

Figure 98: Ground Via Radius



Reference Material Specifications

The guidelines in this chapter were determined based on the following PCB material specifications. Signal integrity simulations are encouraged for any material or specifications that deviate from those listed here.

- Board material: Megtron 6 class
- Dielectric constant: ~3.7
- Dissipation factor: 0.002
- Board thickness:
 - 65 mil for LPDDR4x and LPDDR5
 - 100 mil for DDR4
- Layer Order: Ground-signal-ground

Adjusting for Different Stack-Ups

The trace width, spacing, length, and skew constraints presented in this chapter are based on the reference PCB materials listed in [Reference Material Specifications](#). When not using these specific parameters, trace width, lengths, layer heights, spacings, and dielectric material could need adjustments to meet the impedance, length, and skew specifications.

The PCB fabrication house can adjust these factors to achieve the desired impedance and propagation delay targets. The effects of each of these items are listed in this section. In addition, a two-dimensional field solver utility allows for various combinations to be tested.

Dielectric Material

Each dielectric material has its own relative dielectric constant (D_K) and dissipation factor/loss tangent (D_F) that are contributing factors to line impedance (Z_0), signal propagation delay (T_{PD}), and signal loss (α). As D_K increases, impedance decreases while the signal propagation delay and signal loss increase, and vice versa. A typical D_K found on PCBs ranges from 3.4 to 4.6. The propagation delay in a given dielectric material is constant and is not affected by any other board parameters such as layer height, conductor width, or conductor spacing. Propagation delay is affected by frequency, but the effect is minimal with regards to typical memory speeds. Signal loss is also affected by frequency, with the loss increasing with increasing frequency.

The following equation shows the calculation for propagation delay (T_{PD}), with D_K as the dielectric constant, and c as the speed of light in free space (2.998×10^8 m/s or 1.180×10^{10} in/s).

Equation 2: Propagation Delay Calculation

$$T_{PD} = \frac{\sqrt{D_K}}{c}$$

The associations in the following table show the effect of the dielectric constant (D_K) on the impedance (Z_0), propagation delay (T_{PD}), and signal loss (α).

$D_K \uparrow$	$Z_0 \downarrow$	$T_{PD} \uparrow$	$\alpha \uparrow$
$D_K \downarrow$	$Z_0 \uparrow$	$T_{PD} \downarrow$	$\alpha \downarrow$

Trace Width

As the trace width (W) increases, the impedance decreases while the signal propagation delay remains unchanged, and vice versa. Any adjustments in width should include adjustments in spacing (S) to maintain immunity to crosstalk effects. Spacing factors are roughly 1.0X to 3X depending on the particular type of memory and signal. The associations in the following table show the effect of width on impedance, spacing, and propagation delay.

$W \uparrow$	$Z_0 \downarrow$	$S \uparrow$	T_{PD} (no change)
$W \downarrow$	$Z_0 \uparrow$	$S \downarrow$	T_{PD} (no change)

Layer Height

As the layer height (H) is increased, the impedance increases while the signal propagation delay remains unchanged, and vice versa. When reducing or increasing layer height, consider that layer heights that are too low can be more expensive to reliably manufacture because the PCB fabricator must avoid plane shorts. Layer heights that are too high can lead to aspect ratio violations. The associations in the following table show the effect of layer height on impedance and propagation delay.

H ↑	Z_0 ↑	T_{PD} (no change)
H ↓	Z_0 ↓	T_{PD} (no change)

Base Copper Weight

As the base copper weight (Cu oz) of the conducting lines is increased, the impedance decreases while the signal propagation delay remains unchanged, and vice versa. The most common base copper weight used in PCBs is 0.5, known as half-ounce copper. The effect of changing the base copper weight is minimal, however, and is not recommended due to the potential for increased costs in relation to the small benefit typically gained.

Example of Tuning Design Parameters to Meet Impedance Targets

The following table shows a target of 45Ω for stripline impedance in the main L1 PCB area. The trace width (W) is specified as 7.0 mils, DK of 3.71, and height of 6.7 mils. The table also shows tuning of the dielectric constant, trace width (W), and height (H) required to meet the impedance target. The spacing (S) shows how it should change in order to continue to meet performance requirements, though the spacing is not a factor in the line impedance. When increasing DK, the signal loss also increases as shown in the table.

Table 104: Example of Tuning PCB Parameters to Meet Impedance Target

Target ($[Z_0]$ Ω)	Material	DK	W (mil)	S (mil)	H (mil)	Change from Reference
45	Megtron 6	3.71	7.0	-	6.7	Reference
45	Megtron 6	3.71	7.5 ↑	↑	7.15 ↑	Same DK, bigger width, taller height
45	Megtron 6	3.4 ↓	7.0	-	6.27 ↓	Lower DK, same width, lower height
45	Megtron 6	3.4 ↓	7.5 ↑	↑	6.7	Lower DK, bigger width, same height

PCB Guidelines for DDR4 SDRAM

Overview

This section provides electrical design guidelines for typical DDR4 SDRAM interfaces from system-level signal integrity simulations for Spartan UltraScale+ FPGAs.

 **IMPORTANT!** All routing guidelines in this section must be followed to achieve the maximum data rates specified in Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930). Customers could have unique or specific designs with particular violations of some rules. In these scenarios, design or routing trade-offs have to be taken in other routing parameters to mitigate the risk. System-level channel signal integrity simulations are required to evaluate such trade-offs. It is important to read [Required Memory Routing Guidelines for All Interfaces](#) before continuing with this section.

DDR4 SDRAM Interface Signal Description

The DDR4 SDRAM interface consists of clock, control, address, and data signals as shown in the following table.

Table 105: DDR4 SDRAM I/O Signal Description

Signal Name	Description
Clock Signals	
ck_t, ck_c	Differential clock
Address and Command Signals	
a[17,13:0]	Address inputs
ras_n/a[16]	Row address strobe, address bit 16
cas_n/a[15]	Column address strobe, address bit 15
we_n/a[14]	Write enable, address bit 14
bg[1:0]	Bank group inputs
ba[1:0]	Bank address inputs
act_n	Activation command input
par	Command and address parity input
Control Signals	
cke	Clock enable
cs_n	Chip select
odt	On-die termination enable
reset_n	See reset_n
Data Signals	
dq	Data input/output
dqs_t/dqs_c	Data strobe (differential)
dm_n/dbi_n	Data mask and data bus inversion

Table 105: DDR4 SDRAM I/O Signal Description (cont'd)

Signal Name	Description
Other Signals	
TEN	Test connectivity mode. CMOS level. Connect 500Ω resistor to ground at memory devices.
alert_n	See alert_n
VRP (PL)	240Ω to GND
ZQ (PS)	240Ω to GND


Notes:

- Actual signal list might vary based on configuration.

Fly-by and Clamshell Topologies


Two topology types are supported for DDR4 SDRAM: fly-by, and clamshell. The fly-by topology (see the following figure) consists of all memory devices on one layer, usually in-line. This type of topology is generally easier to route and can offer the best signal integrity, but can take up precious board real estate.

Figure 99: Fly-by Topology



Clamshell topology (see the following figure) requires more intricate routing, but is optimal for designs where board space is at a premium.

Figure 100: Clamshell Topology



Using Address Mirroring to Ease Clamshell Routing

While the clamshell topology has advantages such as reduced board space, the asymmetrical nature of the memory device package pin locations can lead to higher routing congestion in the areas under the memory devices between the top and bottom layers. This routing congestion could potentially lead to higher crosstalk due to the inability to include as many ground-return vias as in other topologies (see item 21 in [General Memory Routing Guidelines](#)), as well as longer stub traces on the top and bottom layers. Address mirroring can be used to change the function of certain pins on a memory device to correspond to the pin directly above or below it. One via can be used for the signal, along with a short stub to the landing pad for each device. Address mirroring is defined in JEDEC specification JESD21-C. Twelve command/address/control pins can be mirrored for DDR4 SDRAM, according to the following table.

Table 106: Allowable Mirror Pins for DDR4 SDRAM

Memory Controller Pin	DRAM Pin (Non-Mirrored)	DRAM Pin (Mirrored)
A3	A3	A4
A4	A4	A3
A5	A5	A6
A6	A6	A5
A7	A7	A8
A8	A8	A7
A11	A11	A13
A13	A13	A11
BA0	BA0	BA1
BA1	BA1	BA0
BG0 ¹	BG0	BG1
BG1 ¹	BG1	BG0

Notes:

1. BG0 and BG1 can only be mirrored when pin BG1 is present on the memory device.

ECC Connection Rules for DDR4 SDRAM

For DDR4 interfaces when using ECC with x16 components, the lower data byte (DQ[7:0]) of the last component must be connected to the ECC byte of the controller. If the lower data byte is not connected and instead the upper data byte is used, the per-DRAM addressability feature does not function correctly, which results in V_{REF} calibration errors. This results in FSBL hangs for Zynq UltraScale+ MPSoC based designs. Functionality is not impacted for soft controller designs because V_{REF} calibration is not used. However, if V_{REF} calibration is enabled, it results in functional errors.

For DDR4-based interfaces, the unused upper data byte must be connected as follows (unless explicitly stated otherwise in the memory vendor's datasheet):

- DQ[15:8] can be left floating.
- UDM or UDM_n/UDBI_n terminate to V_{DD} with a pull-up resistor ≤ 4.7 KΩ.
- UDQS or UDQS_t terminate to V_{DD} with a pull-up resistor ≤ 4.7 KΩ.
- UDQS# or UDQS_c terminate to V_{SS} with a pull-down resistor ≤ 4.7 KΩ.

Routing Rule Changes for Thicker Printed Circuit Boards

The DDR4 routing rules provided in this chapter allow for the interface to run at the maximum supported data rates as specified in the *Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930)* while also assuming a maximum board thickness of 100 mil. For boards that are thicker than 100 mil, the maximum interface speed might need to be adjusted downward. See the following table for maximum memory data rates per board thickness.


Table 107: DDR4 Maximum Data Rates per Board Thickness

DDR4	Board Thickness (mil)		
	≤ 100	101 – 145	101 – 145 + Stub-free ¹
Component	See <i>Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930)</i>	Two memory speed grades lower from data rates in <i>Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930)</i>	See <i>Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930)</i>

Notes:


1. Stub-free refers to alternating the routing of the signals between upper and lower layers such that long via stubs are avoided. See [Figure 29: Thicker PCB: Command/Address/Control Bus Stub-Free Routing](#) for an example of alternate layer routing as compared to conventional routing shown in the following figure.

Figure 101: Thicker PCB: Command/Address/Control Bus Conventional Routing



X25066-012821

Figure 102: Thicker PCB: Command/Address/Control Bus Stub-Free Routing




X25068-012821

Topology and Routing Guidelines for DDR4 SDRAM

DDR4 SDRAM Address, Command, and Control Fly-by and Clamshell Topologies

The following figure illustrates fly-by topology for address, command, and control signals. Each clock, address, command, and control pin on each SDRAM is connected to a single trace and terminated at the far end.

Figure 103: Address, Command, and Control Fly-by Termination for DDR4 SDRAM




UG583_c2_23_062122

reset_n

The following figure shows the termination for reset_n. The 4.7 kΩ to ground is to keep reset_n Low during FPGA power-up. If self-refresh is required during FPGA power-down, circuitry needs to be added to ensure that reset_n stays High during that time.

Figure 104: Termination for reset_n in DDR4 DRAM



alert_n

Connect the PS_DDR_ALERT_N pin to the ALERT_N pins of the DDR4 devices in fly-by routing and terminate to V_{DD} with a 50Ω pull-up resistor. For PL-based designs, the controller does not generate an ALERT_N pin but the DDR4 ALERT_N signals must be connected together via fly-by routing and terminated to V_{DD} with a 50Ω pull-up resistor. The ALERT_N signal does not have any skew or length matching requirements.

The following table shows the DDR4 SDRAM fly-by impedance, length, and spacing guidelines for address, command, and control signals.

Table 108: DDR4 SDRAM Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals


Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R_{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Single-ended impedance Z_0	$50\pm 10\%$	$36\pm 10\%$	$50\pm 10\%$	$50\pm 10\%$	$39\pm 10\%$	Ω
Trace width	4.0	7.0	4.0	4.0	6.0	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	0.0~0.1	0.35~0.75	0~1	inches
Spacing in address, command, and control signals (minimum)	4.0	8.0 ²	4.0	8.0	8.0	mil
Spacing to clock signals (minimum)	8.0	20	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	30	30	30	mil
Maximum PCB via count	7					-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 12.0.

For DDR4 SDRAM clamshell topology, an alternating fly-by topology is recommended for control/address/command signals. The alternating layer routing properly balances the signal loads at each memory device. As depicted in the following figure with the FPGA located at the top layer, the inner layer routing to top layer devices 1, 3, 5, 7, and 9 is closer to the top layer, while the inner layer routing to bottom layer devices 2, 4, 6, and 8 is closer to the bottom layer. The example below is for nine devices, but the same idea holds for interfaces with fewer devices.

Figure 105: Address/Command/Control Topology for DDR4 SDRAM Clamshell



X18996-040517

Note: The end-termination resistor can be located on either the top or bottom layer.

DDR4 SDRAM clamshell topology uses two individual chip select (CS) signals, one for the top layer DRAMs, and one for the bottom layer DRAMs. Those signals should be routed as in the standard fly-by topology as defined in [Figure 30: Address, Command, and Control Fly-by Termination for DDR4 SDRAM](#) and the previous table.

The following table shows the DDR4 SDRAM clamshell impedance, length, and spacing guidelines for address, command, and control signals. Note the extra length of the L2 segment.

Table 109: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Address/Command/Control Signals

Parameter	L0 FPGA Breakout	L1	L2	L3	L4	L5	L6 V _{TT} Stub	Units
Layer (recommended)	Upper inner	Upper inner	Lower inner	Upper inner	Top	Bottom	Bottom	
Impedance Z ₀	50	50	50	50	50	50	39	Ω
Length	0.0~1.5 ¹	0.0~4.0	L3+0.2	0.45~0.85	For mirrored case, L4 = L5; for non-mirrored case, make as short as possible		≤ 0.95	inch
Width	4.0	4.0	4.0	4.0	4.0	4.0	6.0	mil
Spacing within group	4.0	8.0 ¹	8.0	8.0	8.0	8.0	8.0	mil
Spacing to clocks	8.0	20	20	20	20	20	20	mil
Spacing to other groups	8.0	30	30	30	30	30	30	mil

Notes:


1. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 12.0.

★ IMPORTANT! Add as many ground vias as possible to help avoid crosstalk issues. See Item 20 in [General Memory Routing Guidelines](#).

DDR4 SDRAM Clock Fly-By and Clamshell Termination

Inherent to fly-by topology, the timing skew between the clock and dq_s signals is deskewed by the write-leveling feature on DDR4 SDRAM (see the following figure).

Figure 106: Clock Fly-by Termination for DDR4 SDRAM



The following table shows the DDR4 SDRAM impedance, length, and spacing guidelines for clock signals.

Table 110: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals


Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	L3	L4 (To R _{TT})	Units
Trace type	Stripline	Stripline	Stripline	Stripline	Stripline	-
Clock differential impedance Z _{DIFF}	86±10%	76±10%	86±10%	90±10%	76±10%	Ω
Trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	4.0/5.0/4.0	6.0/6.0/6.0	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	0.0~0.1	0.35~0.75	0~1.0	inches
Spacing in address, command, and control signals (minimum)	8.0	20 ²	8.0	20	20	mil
Spacing to other group signals (minimum)	8.0	30	8.0	30	30	mil
Maximum PCB via count per signal	7					-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 30.

For DDR4 SDRAM clamshell topology, an alternating fly-by topology is recommended for clock signals. The alternating layer routing properly balances the signal loads at each memory device. As depicted in [Table 36: Impedance, Length, and Spacing Guidelines for One-Slot DIMM Address, Command, and Control Signals](#) with the FPGA located at the top layer, the inner layer routing to top layer devices 1, 3, 5, 7, and 9 is closer to the top layer, while the inner layer routing to bottom layer devices 2, 4, 6, and 8 is closer to the bottom layer.

Figure 107: Clock Topology for DDR4 SDRAM Clamshell



Note: The end-termination components can be located on either the top or bottom layer.

The following table shows the DDR4 SDRAM clamshell impedance, length, and spacing guidelines for clock signals. Note the extra length of the L2 segment.

Table 111: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Clock Signals

Parameter	L0 FPGA Breakout	L1	L2	L3	L4	L5	L6 V _{TT} Stub	Units
Layer (recommended)	Upper inner	Upper inner	Lower inner	Upper inner	Top	Bottom	Bottom	
Differential impedance Z ₀	86	76	90	90	86	86	76	Ω
Length	0.0~1.5 ¹	0.0~4.0	L3+0.2	0.45~0.85	For mirrored case, L4 = L5; for non-mirrored case, make as short as possible		≤ 1.0	inch
Width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/5.0/4.0	4.0/5.0/4.0	4.0/4.0/4.0	4.0/4.0/4.0	6.0/6.0/6.0	mil
Spacing to addr/cmd/ctrl	8.0	20 ¹	20	20	20	20	20	mil

Table 111: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Clock Signals (cont'd)

Parameter	L0 FPGA Breakout	L1	L2	L3	L4	L5	L6 V _{TT} Stub	Units
Spacing to other groups	8.0	30	30	30	30	30	30	mil

Notes:

1. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 30.




IMPORTANT! Add as many ground vias as possible to help avoid crosstalk issues. See item 20 in [General Memory Routing Guidelines](#).

DDR4 SDRAM Data Signals Point-to-Point for Fly-by and Clamshell Configurations

The following figure shows the data signals (DQ, DM, and DQS) point-to-point for DDR4 SDRAM. The guidelines are the same for fly-by and clamshell configurations.

Figure 108: Data Signals Point-to-Point for DDR4 SDRAM



The following table shows the DDR4 SDRAM impedance, length, and spacing guidelines for data signals.

Table 112: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
Trace type	Stripline	Stripline	Stripline	-
dq single-ended impedance Z_0	50±10%	39±10%	50±10%	Ω
dqs differential impedance Z_{DIFF}	95±10%	76±10%	95±10%	Ω
Trace width (nominal)	4.0	6.0	4.0	mil
Differential trace width/space/width	4.0/4.0/4.0	6.0/6.0/6.0	4.0/4.0/4.0	mil
Trace length	0.0~1.5 ^{1,2}	0.0~4.0	0.0~0.1	inches
Spacing in byte (minimum)	4.0	8.0 ²	4.0	mil
Spacing byte to byte (minimum)	4.0	20	4.0	mil

Table 112: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals (cont'd)

Parameter	L0 (Device Breakout)	L1 (Main PCB)	L2 (DRAM Breakout)	Units
dq to dq _s spacing (minimum)	4.0	20	8.0	mil
Spacing to other group signals (minimum)	8.0	30	30	mil
Maximum PCB via count	2			-

Notes:

1. See item 2 in [General Memory Routing Guidelines](#).
2. If greater than 1.5 inches is needed in the breakout area, change L1 spacing to 12.0.

DDR4 SDRAM Routing Constraints

There are two constraints requirements for each signal group in the DDR4 memory interface:

- Total length/delay constraints
- Skew constraints

The total length/delay constraints are shown in the following table.

Table 113: DDR4 SDRAM Total Length/Delay Constraints

Signal Group	Reference Figure	Maximum Length/Delay Constraints
address/command/control: x1/x2/x3/x4/x5/x6/x7/x8/x9 (includes package routing P0)	Figure 103: Address, Command, and Control Fly-by Termination for DDR4 SDRAM	7.15/7.9/8.65/9.4/10.15/10.9/11.65/12.4/13.15 inches 1211/1339/1466/1593/1720/1847/1974/2101/2228 ps
Data Signals P0+L0+L1+L2	Figure 108: Data Signals Point-to-Point for DDR4 SDRAM	6.0 inches, 1017 ps

The skew constraints are listed in the following tables.

Table 114: DDR4 SDRAM Data Group Skew Constraints

Signal Group	Skew Constraints (ps)	Skew Constraints (mil)
Data ¹ to DQS	±10	±58
dq _{s_p} and dq _{s_n}	2	12
CK to DQS	-149 to 1,796	-879 mil to 10.59 inches

Notes:

1. The data group includes dq and dm_n/dbi_n.
2. Delays are based on 169.5 ps/in. See item 8 in [General Memory Routing Guidelines](#).
3. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.

The data group skew constraints are listed in the following table.

Table 115: DDR4 SDRAM Address, Command, and Control Skew Constraints

Signal	Signal Segment	Skew Constraints (ps)	Skew Constraints (mil)
address/command/control ¹ to CK	UltraScale device to each memory device	±8 ps	±47
ck_p and ck_n	UltraScale device to each memory device	2	12

Notes:

1. The signal reset_n is not required to meet the skew constraints in this table.
2. For skew specifications, refer to the [General Memory Routing Guidelines](#) items 3–8.



IMPORTANT! FPGA package flight times must be included in both total length constraints and skew constraints. When minimum and maximum values are available for the package delay, use the midrange between the minimum and maximum values. Memory device package flight times do not need to be factored in because their variances have been accounted for in these guidelines.

PCB Routing Guidelines for LPDDR5 Interfaces

This section provides PCB design guidelines for LPDDR5 interfaces. Connections between the FPGA and LPDDR5 device are defined along with physical design rules and timing constraints.



IMPORTANT! All routing guidelines in this section must be followed to achieve the maximum data rates specified in *Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930)*. Customers could have unique or specific designs with particular violations of some rules. In these scenarios, design or routing trade-offs have to be taken in other routing parameters to mitigate the risk. System-level channel signal integrity simulations are required to evaluate such trade-offs. It is important to read [Required Memory Routing Guidelines for All Interfaces](#) before continuing with this section.

Note: LPDDR5x devices can be used in place of LPDDR5 devices but should only be operated at the LPDDR5 data rates as specified in *Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930)*.

Signals and Connections for LPDDR5 Interfaces

The required signals used in LPDDR5 applications are shown in the following table.

Table 116: LPDDR5 Signal Definitions


Signal	Description	Required PCB Termination
Clock Signals		
CK_T[C]_A, CK_T[C]_B	Address/Command clock	None, uses ODT 40Ω
Command/Address Signals		
CA[6:0]_A, CA[6:0]_B	Address	None, uses ODT 40Ω

Table 116: LPDDR5 Signal Definitions (cont'd)

Signal	Description	Required PCB Termination
Control Signals		
CS0_A, CS0_B	Chip Select	240Ω to GND
Data Signals		
DQ[15:0]_A, DQ[15:0]_B	Data	None, uses ODT 40Ω
DMI[1:0]_A, DMI[1:0]_B	Data Mask Inversion	None, uses ODT 40Ω
Data Clock Signals		
WCK[1:0]_T[C]_A, WCK[1:0]_T[C]_B	Differential Data Clocks	None, uses ODT 40Ω
Read Data Strobe Signals		
RDQS[1:0]_T[C]_A, RDQS[1:0]_T[C]_B	Differential Read Data Strobe	None, uses ODT 40Ω
Miscellaneous Signals		
RESET_n	Reset (one per LPDDR5 device)	4.99 kΩ to GND
LPDDR5 Device Only		
ZQ_A	Calibration Reference (per device)	240Ω to VDDQ


The following figures show the various supported connection options for LPDDR5 such as 1x16 single rank and 1x32 single rank.

Figure 109: LPDDR5 1x16 Single Rank



X29929-121725

Figure 110: LPDDR5 1x32 Single Rank



X29931-121725

Supply Voltages between FPGA and LPDDR5 Device

Users should be aware of the varying voltage levels for the FPGA and LPDDR5 device, as shown in the following table.

Table 117: Typical Supply Voltages for FPGA and LPDDR5¹

Interface	VCCO_xxx (FPGA)	VDDQ	VPP	VDD2H/VDD2L
LPDDR5	1.0 V	0.50 V	1.8 V	1.05 (Single Core Power Rail)

Notes:

1. Only I/O buffer specification range 1 (0.50V) is supported.

Physical Design Rules for LPDDR5 Signals

The following table defines routing rules for LPDDR5 signals.

Table 118: Physical Design Rules for LPDDR5 Signals

Parameter	Value
Impedance Rules	
Impedance for single-ended CAC ¹ and data signals	40Ω ² ±10%
Impedance for differential clock, data clock and read data strobe signals	75Ω ² ±10%
PCB Trace Length Rules (from FPGA to Furthest Device or Termination)	
Maximum PCB trace length for CAC/CK signals	2000 mil
Total Trace Length Rules (PCB + FPGA Package Routing to Furthest Device)	
Maximum PCB trace length for data signals	2000 mil
Spacing Rules for CAC and Clock Signals	
Minimum spacing between CAC signals within the same channel ³	2.5H ⁴ , except: 1H under FPGA or LPDDR5 device
Minimum spacing between CA signals and clock signals within the same channel ³	5H ⁴ , except: 2H under FPGA or LPDDR5 device
Minimum spacing between CAC/clock signals and data/strobe signals within the same channel ³	7H ⁴ , except: 2H under FPGA or LPDDR5 device
Spacing Rules for Data, Data Clock, and Read Data Strobe Signals	
Minimum spacing between data signals within the same byte	2.5H ⁴ , except: 1H under FPGA or LPDDR5 Device
Minimum spacing between data signals and WCK/RDQS signals within the same byte	5H ⁴ , except: 1H under FPGA or LPDDR5 device
Minimum spacing between data signals and WCK/RDQS signals between different bytes	7H ⁴ , except: 2H under FPGA or LPDDR5 device
Minimum spacing between data/WCK/RDQS signals and other signals within the same channel ³	7H ⁴ , except: 2H under FPGA or LPDDR5 device
Spacing Rules for Signals between Channels or Memory Interfaces	
Minimum spacing between signals in one memory interface to signals in another channel or memory interface	7H ⁴ , except: 2H under FPGA or LPDDR5 device
Maximum Vias per Signal	
Data, WCK, and RDQS signals	2 each
Other Physical Design Requirements	
Route all data/strobe signals within a byte on one internal signal layer	
Place FPGA and LPDDR5 devices on the top side and route all data/strobe signals on upper internal layers	
Maintain 2:1 signal to GND via ratio on the FPGA side of the PCB	

Notes:

1. CAC stands for command/address/control.
2. This can be up to 55Ω (single-ended) or 100Ω (differential) under adaptive SoC or DRAM devices, taking into account PCB manufacturing tolerances.
3. A channel is defined as either the “A” side of the LPDDR5 device or the “B” side of the LPDDR5 device.
4. H is the distance to the nearest ground return plane.

Timing Constraint Rules for LPDDR5 Signals

The following table defines timing constraints for various signal groups and their targets, similar to how they would be entered into PCB layout software tools. FPGA package delays should always be included for purposes of determining skews.

Table 119: Skew Constraint Rules for LPDDR5 Signals

Skew Constraint ¹	Pin Pair Set	Minimum (ps)	Maximum (ps)	Group	Target
Address to Clock A	FPGA to LPDDR5 Device	-50	+50	CA[6:0]_A	CK_T_A
Address to Clock B	FPGA to LPDDR5 Device	-50	+50	CA[6:0]_B	CK_T_B
Command to Clock A	FPGA to LPDDR5 Device	-20	+20	CS0_A CS1_A	CK_T_A
Command to Clock B	FPGA to LPDDR5 Device	-20	+20	CS0_B CS1_B	CK_T_B
Clock A ²	FPGA to LPDDR5 Device	0	2	CK_T_A CK_C_A	-
Clock B ²	FPGA to LPDDR5 Device	0	2	CK_T_B CK_C_B	-
Data to WCK0_A	FPGA to LPDDR5 Device	-50	+50	DQ[7:0]_A DMI0_A	WCK0_T_A
Data to WCK1_A	FPGA to LPDDR5 Device	-50	+50	DQ[15:8]_A DMI1_A	WCK1_T_A
Data to WCK0_B	FPGA to LPDDR5 Device	-50	+50	DQ[7:0]_B DMI0_B	WCK0_T_B
Data to WCK1_B	FPGA to LPDDR5 Device	-50	+50	DQ[15:8]_B DMI1_B	WCK1_T_B
Data to RDQS0_A	FPGA to LPDDR5 Device	-50	+50	DQ[7:0]_A DMI0_A	RDQS0_T_A
Data to RDQS1_A	FPGA to LPDDR5 Device	-50	+50	DQ[15:8]_A DMI1_A	RDQS1_T_A
Data to RDQS0_B	FPGA to LPDDR5 Device	-50	+50	DQ[7:0]_B DMI0_B	RDQS0_T_B
Data to RDQS1_B	FPGA to LPDDR5 Device	-50	+50	DQ[15:8]_B DMI1_B	RDQS1_T_B
WCK0_A ²	FPGA to LPDDR5 Device	0	2	WCK0_T_A WCK0_C_A	-
WCK1_A ²	FPGA to LPDDR5 Device	0	2	WCK1_T_A WCK1_C_A	-
WCK0_B ²	FPGA to LPDDR5 Device	0	2	WCK0_T_B WCK0_C_B	-
WCK1_B ²	FPGA to LPDDR5 Device	0	2	WCK1_T_B WCK1_C_B	-

Table 119: Skew Constraint Rules for LPDDR5 Signals (cont'd)

Skew Constraint ¹	Pin Pair Set	Minimum (ps)	Maximum (ps)	Group	Target
RDQS0_A ²	FPGA to LPDDR5 Device	0	2	RDQS0_T_A RDQS0_C_A	-
RDQS1_A ²	FPGA to LPDDR5 Device	0	2	RDQS1_T_A RDQS1_C_A	-
RDQS0_B ²	FPGA to LPDDR5 Device	0	2	RDQS0_T_B RDQS0_C_B	-
RDQS1_B ²	FPGA to LPDDR5 Device	0	2	RDQS1_T_B RDQS1_C_B	-
WCK0_A to Clock A	FPGA to LPDDR5 Device	-50	+50	WCK0_T_A	CK_T_A
WCK1_A to Clock A	FPGA to LPDDR5 Device	-50	+50	WCK1_T_A	CK_T_A
WCK0_B to Clock B	FPGA to LPDDR5 Device	-50	+50	WCK0_T_B	CK_T_B
WCK1_B to Clock B	FPGA to LPDDR5 Device	-50	+50	WCK1_T_B	CK_T_B

Notes:

1. Include FPGA package delays for all skew calculations.
2. It does not matter which signal is faster or slower, but the difference in time between the two should be no longer than specified.
3. There should be individual constraint sets for each byte/nibble/DQS pair.

PCB Routing Guidelines for LPDDR4x Interfaces

This section provides PCB design guidelines for LPDDR4x interfaces. Connections between the FPGA and LPDDR4x device are defined, along with physical design rules and timing constraints.

Note: To obtain maximum LPDDR4x speed, it is recommended to keep the board thickness less than 131 mil. Trace length specifications vary slightly based on board thickness. See [Physical Design Rules for LPDDR4x Signals](#) for trace length rules.



IMPORTANT! All routing guidelines in this section must be followed to achieve the maximum data rates specified in the Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930). Customers could have unique or specific designs with particular violations of some rules. In these scenarios, design or routing trade-offs have to be taken in other routing parameters to mitigate the risk. System-level channel signal integrity simulations are required to evaluate such trade-offs. It is important to read [Required Memory Routing Guidelines for All Interfaces](#) before continuing with this section.


Signals and Connections for LPDDR4x Interfaces

The required signals used in LPDDR4x applications are shown in the following table.

Table 120: LPDDR4x Signal Definitions

Signal	Description	Required PCB Termination
Clock Signals		
CK_T[C]_A, CK_T[C]_B	Address/Command clock	None, uses ODT 48Ω
Command/Address Signals		
CA[5:0]_A, CA[5:0]_B	Address	None, uses ODT 48Ω
Control Signals		
CS0_A, CS0_B	Chip Select	None, uses ODT 48Ω
CKE0_A, CKE0_B	Clock Enable	160Ω to VDDQ (0.6V) (see Figure 111: LPDDR4X CKE Termination) Note: When terminating to VDDQ, ensure VDDQ rail can sink current.
Data Signals		
DQ[15:0]_A, DQ[15:0]_B	Data	None, uses ODT 48Ω
DMI[1:0]_A, DMI[1:0]_B (if present)	Data Mask	None, uses ODT 48Ω
Data Strobe Signals		
DQS[1:0]_T[C]_A, DQS[1:0]_T[C]_B	Differential Data Strobe	None, uses ODT 48Ω
Miscellaneous Signals		
RESET_n	Reset	4.7 kΩ to GND
LPDDR4x Device		
ODT_A, ODT_B	On Die Termination Control	Direct to VDD2
ZQ0	Calibration Reference	240Ω to VDDQ
FPGA		
IO_VR_7xx IO_VR_8xx (if present)	Calibration Reference	240Ω to VCCO_7xx 240Ω to VCCO_8xx (if present)

Figure 111: LPDDR4X CKE Termination



X29906-101024

Note: When terminating CKE to VDDQ (0.6V), ensure that the VDDQ rail can sink current.

Supply Voltage between FPGA and LPDDR4x

Users should be aware of the varying voltages levels for the FPGA and LPDDR4x device as shown in the following table.

Table 121: Typical Supply Voltages for FPGA and LPDDR4x

Interface	VCCO_xxx (FPGA)	VDD2	VDDQ
LPDDR4x	1.2V	1.1V	0.6V

The following figures show the connections for 1x16 and 1x32 LPDDR4x interfaces.

Figure 112: Connections for a 1x16 LPDDR4x Interface





Figure 113: Connections for a 1x32 LPDDR4x Interface



Physical Design Rules for LPDDR4x Signals

The following table defines routing rules for LPDDR4x signals.

Table 122: Physical Design Rules for LPDDR4x Signals

Parameter	Value
Impedance Rules	
Impedance for single-ended CAC ¹ and data signals	45Ω ±10%
Impedance for differential clock and data strobe signals	82Ω ±10%
Trace Length Rules (from FPGA to Furthest Device or Termination)	
Maximum PCB trace length for CAC signals	3600 mil (route on inner signal layers only)
Maximum PCB trace length for data signals	3600 mil (route on inner signal layers no deeper than 45 mil from FPGA)
Spacing Rules for CAC and Clock Signals	
Minimum spacing between CAC signals within the same channel ²	2H ³ , except: 1H under FPGA or LPDDR4x device
Minimum spacing between CAC signals and clock signals within the same channel ²	5H ³ , except: 2H under FPGA or LPDDR4x device
Minimum spacing between CAC/clock signals and data/strobe signals within the same channel ²	7H ³ , except: 2H under FPGA or LPDDR4x device
Spacing Rules for Data and Data Strobe Signals	
Minimum spacing between data signals within the same byte	2H ³ , except: 1H under FPGA or LPDDR4x Device
Minimum spacing between data signals and data strobe signals within the same byte	5H ³ , except: 1H under FPGA or LPDDR4x device
Minimum spacing between data/strobe signals between different bytes	7H ³ , except: 2H under FPGA or LPDDR4x device
Minimum spacing between data/strobe signals and other signals within the same channel ²	7H ³ , except: 2H under FPGA or LPDDR4x device
Spacing Rules for Signals between Channels or Memory Interfaces	
Minimum spacing between signals in one memory interface to signals in another channel or memory interface	7H ³ , except: 2H under FPGA or LPDDR4x device
Maximum Vias per Signal	
CAC signals and clock signals	2 each, except: 4 for signals touching more than one channel, such as with pin-efficient pinouts
CKE signals	3 each
Data and strobe signals	2 each
Other Physical Design Requirements	
Do not route CAC/clock signals on more than one internal signal layer	
Route all data/strobe signals with a byte on one internal signal layer	

Table 122: Physical Design Rules for LPDDR4x Signals (cont'd)

Parameter	Value
Route data/strobe signals on internal layers as close to the memory devices as possible	

Notes:

1. CAC stands for command/address/control.
2. A channel is defined as either the "A" side of the LPDDR4x device or the "B" side of the LPDDR4x device.
3. H is the distance to the nearest ground return plane.

Timing Constraint Rules for LPDDR4x Signals

The following table defines timing constraints for various signal groups and their targets, similar to how they would be entered into PCB layout software tools. FPGA package delays should always be included for purposes of determining skews.

 Table 123: Skew Constraint Rules for LPDDR4x Signals¹

Skew Constraint	Pin Pair Set	Minimum (ps)	Maximum (ps)	Group	Target
Address to Clock A	FPGA to LPDDR4x Device	-100	+100	CA[5:0]_A	CK_T_A
Address to Clock B	FPGA to LPDDR4x Device	-100	+100	CA[5:0]_B	CK_T_B
Command to Clock A	FPGA to LPDDR4x Device	-20	+20	CKE0_A	CK_T_A
Command to Clock B	FPGA to LPDDR4x Device	-20	+20	CKE0_B	CK_T_B
Clock (A or B) ²	FPGA to LPDDR4x Device	0	2	CK_T_A/B CK_C_A/B	-
Data to DQS0	FPGA to LPDDR4x Device	-100	+100	DQ[7:0] DM0	DQS0_T
Data to DQS1	FPGA to LPDDR4x Device	-100	+100	DQ[15:8] DM1	DQS1_T
Data to DQS2	FPGA to LPDDR4x Device	-100	+100	DQ[23:16] DM2	DQS2_T
Data to DQS3	FPGA to LPDDR4x Device	-100	+100	DQ[31:24] DM3	DQS3_T
DQS0 ²	FPGA to LPDDR4x Device	0	2	DQS0_T DQS0_C	-
DQS1 ²	FPGA to LPDDR4x Device	0	2	DQS1_T DQS1_C	-
DQS2 ²	FPGA to LPDDR4x Device	0	2	DQS2_T DQS2_C	-
DQS3 ²	FPGA to LPDDR4x Device	0	2	DQS3_T DQS3_C	-
DQS0 to Clock A ²	FPGA to LPDDR4x Device	-150	+150	DQS0_T	CK_A

Table 123: Skew Constraint Rules for LPDDR4x Signals¹ (cont'd)

Skew Constraint	Pin Pair Set	Minimum (ps)	Maximum (ps)	Group	Target
DQS1 to Clock A ²	FPGA to LPDDR4x Device	-150	+150	DQS1_T	CK_A
DQS2 to Clock B ²	FPGA to LPDDR4x Device	-150	+150	DQS2_T	CK_B
DQS3 to Clock B ²	FPGA to LPDDR4x Device	-150	+150	DQS3_T	CK_B

Notes:

1. Include FPGA package delays for all skew calculations.
2. It does not matter which signal is faster or slower, but the difference in time between the two should be no longer than specified.

PCB Guidelines for Zynq UltraScale+ RFSocS

PDN Guidelines for RFSoc Digital Power Rails

AMD Zynq™ UltraScale+™ RFSocS contain new digital and analog circuitry, with dedicated power rails for each block. This section focuses on the PDN guidelines for the digital power rails in RFSocS, including the soft-decision forward error correction (SD-FEC) power rail $V_{CCSDFEC}$.

Power Delivery to the Programmable Logic Voltage Rails

The guidelines in this document are based on the following use case assumptions.

- 80% of CLB LUTs and registers running at 245 MHz
- 80% of block RAMs running at 491 MHz
- 80% of UltraRAMs running at 200 MHz
- 80% of DSPs running at 491 MHz
- 100% of MMCM and PLL running at 500 MHz
- 100% of GTY transceiver usage
- 8.02A step-load current on $V_{CCSDFEC}$

V_{CCINT} , V_{CCINT_IO} , and V_{CCBRAM} Tied Together (-1E, -1I, -2I, -2E)

For the -1E, -1I, -2I and -2E speed grades, AMD recommends connecting V_{CCINT} , V_{CCINT_IO} , and V_{CCBRAM} together because all run at the same voltage (0.85V or 0.90V (-3 only)). PCB plane areas are also easier to draw, because V_{CCINT}/V_{CCINT_IO} and V_{CCBRAM} can share the same plane design. Plane design for $V_{CCINT}/V_{CCINT_IO}/V_{CCBRAM}$ should be sized in such a way as to reliably provide the required amount of current/power with minimal IR drop to ensure that the voltage rails meet the data sheet specifications.


Separate V_{CCINT} and V_{CCBRAM}/V_{CCINT_IO} (-1LI, -2LI, -2LE)

For the -1LI, -2LI, and -2LE speed grades, V_{CCINT} can operate at 0.72V versus 0.85V/0.90V for V_{CCBRAM}/V_{CCINT_IO} . It is thus necessary to separate those rails onto different planes. Plane design for V_{CCINT} and V_{CCBRAM}/V_{CCINT_IO} should be sized in such a way as to provide the required amount of current/power with minimal IR drop to ensure that the voltage rails meet the data sheet specifications.

$V_{CCSDFEC}$

It is required to connect the $V_{CCSDFEC}$ rail to the plane that contains V_{CCBRAM} . $V_{CCSDFEC}$ and V_{CCBRAM} should have separate planes under the RFSoc, which are then merged when outside of the RFSoc shadow. The following figure shows the recommended plane methodology.

Figure 114: Recommended Plane Connection for $V_{CCSDFEC}$ and V_{CCBRAM}



X18895-061019

$V_{CCSDFEC}$ and Migration

To enable seamless migration between RFSoc devices, always connect the $V_{CCSDFEC}$ pins to V_{CCBRAM} , even if the $V_{CCSDFEC}$ circuitry is not going to be used. Extra decoupling capacitors for $V_{CCSDFEC}$ are not required when the circuitry is not being used.

PCB Decoupling Recommendations

The recommended decoupling capacitor quantities for Zynq UltraScale+ RFSocS can be found in the Power Design Manager (PDM) tool (download at www.amd.com/power). The capacitor quantities provided in the PDM tool are derived specifically for each design based on actual device utilization. It is for this reason that the decoupling tables in this user guide have been removed. However, the decoupling tables have been moved to [Answer Record 39033](#) to provide continuity as customers move to the PDM tool.

Power Delivery to the PS-GTR and GTY Transceivers

PCB guidelines for the PS-GTR transceivers can be found in [PS-GTR Transceiver Interfaces](#). The guidelines include decoupling capacitors and a design checklist.

PCB design guidelines for GTY transceivers can be found in the Board Design Guidelines chapter of *UltraScale Architecture GTY Transceivers User Guide* ([UG578](#)). The guidelines include decoupling capacitors and a design checklist.


ADC and DAC PCB Guidelines

This section provides routing guidelines for the AMD Zynq™ UltraScale+™ RFSocS with analog-to-digital (ADC) and digital-to-analog (DAC) functionality.

Zynq UltraScale+ RFSoc Device Organization and PCB Design Overview

To assist with keeping sensitive analog circuitry as separate as possible from digital circuitry, the Zynq UltraScale+ RFSoc is sectioned on the die such that the Analog-to-Digital (ADC) and Digital-to-Analog (DAC) portions are partitioned away from each other and from digital logic. The following figure shows the floorplan of the XCZU28DRFFVE1156 RFSoc with the ADC and DAC signals and pins located separately from the digital programmable logic (PL).

Figure 115: XCZU28DR-FFVE1156 Floorplan



The overall goal of a successful PCB design is to keep the analog portions as isolated as possible from both outside electromagnetic (EM) interference and potential on-chip EM interference by the digital logic and analog circuits.

Analog Ground to Digital Ground Connection

While the analog and digital grounds are isolated on the Zynq UltraScale+ RFSoc, eventually they must connect to each other on the PCB. The recommended method to ensure that the analog ground planes on the Zynq UltraScale+ RFSoc connect to the digital ground plane with as little noise transfer as possible is to use one common ground plane.

When connecting analog ground and digital ground as one common plane, adhere to the following guidelines:

- Do not add any slots in the ground plane.

- Place digital voltage regulator modules (VRMs) and capacitors near the digital power rails.
- Place RF VRMs and capacitors near RF power rails.
- Add ground stitching around groups of RF signals to 158 mil (4.0 mm) pitch. See [Figure 122: Close View of Guard Traces and Ground Stitching](#), [Figure 123: Cross-Section of Guard Traces](#) and [Figure 124: Rotated Panel to Compensate for Glass Non-Homogeneity](#).

Choosing the Appropriate Balun

A balun is typically used to interface to the differential input/output of the data converter for converting to a single-ended signal. The bandwidth of the chosen balun should be greater than the band of interest for the user application. This ensures that the measured RFSoc bandwidth is not limited by the balun bandwidth. Refer to *Zynq UltraScale+ RFSoc Data Sheet: DC and AC Switching Characteristics (DS926)* for exact values in the bandwidth column based on different generations.

The Zynq UltraScale+ RFSoc is designed to ensure that all the data converter inputs and outputs along with the clock inputs see a 100Ω DC resistance under normal operating conditions. The purpose of the balun is to convert the single-ended signal with a typical characteristic impedance of 50Ω /75Ω to a differential 100Ω system to meet the RFSoc requirements. Customers can either opt for a wire-wound or a stripline balun depending on the application requirements. The wire-wound baluns are generally used to generate differential signals for low-frequency applications on the order of hundreds of kHz to hundreds of MHz while the RF stripline baluns are more suited for GHz applications. The RF stripline baluns tend to exhibit a band-pass filter characteristic and thus care should be taken to pick a balun with the correct bandwidth requirements that matches the application requirements.

Customers should pay close attention to the balun specification recommendations in [Table 124: Balun Specification Recommendations](#) over the application bandwidth of interest when picking a balun. Several parts that can be considered are:

- Anaren BD1631J50100AHF (approximate bandwidth 1.4 GHz – 3.5 GHz)
- Anaren BD3150L50100AHF (approximate bandwidth 3.1 GHz – 6 GHz)
- Mini Circuits TCM2-33X+ (approximate bandwidth 10 MHz – 1.4 GHz)
- Mini Circuits TC1-1-13M-75X (cable applications)

Table 124: Balun Specification Recommendations

Specification	Typical Result
Impedance ratio	2:1
Bandwidth	Application specific
Insertion loss	-1 dB or better
Return loss	-15 dB or better
Common-mode rejection ratio (CMRR)	> 30 dB ¹

Table 124: Balun Specification Recommendations (cont'd)


Specification	Typical Result
Amplitude imbalance	< ~0.5 dB
Phase imbalance	< ~1.5°

Notes:

1. CMRR can be relaxed if the system is designed to avoid second harmonic distortion (HD2) by frequency planning.


The RF signal chain for each DAC and ADC, which consists of differential DC blocks, differential Pi attenuators (recommended for Gen1/Gen2 only), baluns, and filters, fits well on a 4 mm pitch including the RF shields which cover the RF signal chain components. The RF shield should be electrically attached to ground vias in the PCB at 3 mm spacing (same spacing as the guard stitching vias for signal traces). Failure to adhere to this recommendation results in violating crosstalk guidelines. While the attachment pitch is nominally 3 mm, each attachment point is allowed a tolerance of ± 1 mm from its nominal position. To help facilitate attaching the RF shields to the PCB in a tight space, AMD recommends coating the inside surface of the shield with liquid photo-imageable (LPI) solder mask. This avoids potential shorting to components located inside the RF shield. If soldering the RF shield to the PCB, AMD recommends plating the shield with either bright tin, solder plating, or other plating material that can bond to the solder. The ground vias to which the RF shield attaches should be through-hole vias connected to all inner ground plane layers. An example of this layout can be seen in the following figure.

Figure 116: Balun Isolation with GND Stitching and GND Plane



X21264-073018

Figure 117: RF Shield Solder Mask



X21265-073018

Recommended Clocking Options

The ADC and the DAC both have options to take in a direct RF sampling clock or a lower frequency reference clock to drive the internal PLLs. The built-in internal PLLs offer excellent phase noise and spurious performance and are likely to satisfy the requirements from the user application. The detailed performance of the PLL can be found in the *Zynq UltraScale+ RFSoc Data Sheet: DC and AC Switching Characteristics (DS926)* along with requirements for the reference clocks. Using the internal PLLs offers the most economical and highly integrated solution. In the case where very specific sampling clocking performance is required, an external RF PLL clock can be provided. Devices such as the Texas Instruments LMX259x family can be considered.

When using an external RF clock, particular care must be taken on the P to N skew of the differential input clock. This is very important if operating the DAC at near maximum sample rates greater than 9 GS/s. The P to N clock skew (electrical imbalance, not physical imbalance) presented at the BGA clock receiver balls should be of the order of lesser than 4 ps at the sample rate of interest. Due to the sensitivity and precision of this figure, AMD recommends that this figure be derived from a full 3D EM extraction of PCB. Higher P to N skews can lead to duty cycle distortion of the internal sampling clocks, which can limit the maximum achievable sample rate of the DAC. This is NOT a requirement if providing an input reference clock to the internal PLL and only for cases of providing an external RF clock at high sample rates.

In addition to P to N skew requirements on the DAC clock inputs, the duty cycle of the source clock should be kept to 49/51% when applying an RF clock input. Duty cycle on the input RF clock can again limit the achievable DAC maximum sample rates. A lower duty cycle of 48/52% can be tolerated if the DAC sample rate is lower than 9 GS/s. If the customer has difficulty in achieving the required specifications, further assistance should be requested from your local FAE.

Gen 3 RFSoc devices support a clock forwarding capability that allows either an internally generated RF clock or the external RF clock to be forwarded to other tiles. The clock should be sourced from the center two tiles and cannot be taken from outer tiles if forwarding an RF clock. If using the internal PLL, a reference rate clock can be sourced from the outer edge tiles. For more information on the rules and usage of clock forwarding, refer to *Zynq UltraScale+ RFSoc RF Data Converter LogiCORE IP Product Guide* ([PG269](#)).


If using the internal RF clock forwarding from either the external clock or internal PLL, there is a derating of 150 MHz (1.5%) of the maximum sample rate achieved without clock forwarding. This can be avoided by either using a separate RF clock input for each tile if available, or the internal PLL local to each tile.

RF PLL Placement and Routing

It is recommended to place the RF PLLs close to the RFSoc. The DAC RF PLLs should be placed next to the DAC side of the RFSoc while the ADC RF PLLs should be placed next to the ADC side of the RFSoc. Simulation analysis should be performed to ensure that the clock input amplitude at the RFSoc meets the specifications in the data sheet.

For PLLs that require pull-up resistors, ensure the resistor and decoupling capacitor layouts are symmetrical with respect to P and N. It is recommended to put the pull-up resistors and the decoupling capacitors for the buffer on the same layer as the RF PLL, typically layer one. The following figure shows a snapshot of a typical layout implementation.

Figure 118: PLL Layout



X20580-032618

Analog and Clock Pair Routing

Keeping the analog and clock signals as free as possible from any noise sources is essential, whether they be from the outside environment, or on the PCB itself. The following sections show various specifications and recommendations that result in optimal board and system performance.

To demonstrate the optimum data converter performance, AMD recommends that the PCB meet the signal integrity requirements set out in the following table. Refer to the sections below for more details regarding each type of specification.

Table 125: Signal Integrity Specifications for DAC/ADC Pairs and Clocks

Type	Metric	Maximum Frequency ¹
RFDC Signaling	Maximum Frequency	4 GHz (Gen 1) 5 GHz (Gen 2) 6 GHz (Gen 3) 7 GHz (DFE)
	DIFF FEXT	-70 dBc
	COMM FEXT	-60 dBc
	DIFF NEXT	-70 dBc
	COMM NEXT	-60 dBc
	CMRR	-38 dB
	DIFF RL	-18 dB
	COMM RL	-15 dB
	IL	-2.0 dB
RFDC Clocking	Maximum Frequency	6.4 GHz (Gen 1, Gen 2) 10 GHz (Gen 3)
	DIFF FEXT (DAC to ADC)	-80 dBc
	COMM FEXT (DAC to ADC)	-75 dBc
	DIFF NEXT (DAC to ADC)	-80 dBc
	COMM NEXT (DAC to ADC)	-75 dBc
	CMRR	-35 dB ²
	DIFF RL	-15 dB
	COMM RL	-12 dB
	Power at BGA footprint	2 dBm

Notes:

1. Refer to *Zynq UltraScale+ RFSoc Data Sheet: DC and AC Switching Characteristics (DS926)* for maximum bandwidth of target device.
2. Only applicable for external RF sampling clock.

Trace Routing Impedance Recommendation

Analog pairs should be loosely coupled and routed differentially to $100\Omega \pm 7\%$.

- Guard trace width should be 3.25H
- Gap between P and N traces should be 2.5H

AC/DC Coupling Guidelines

In applications that use AC coupling for interfacing the data converters, DAC outputs should be AC coupled using a 100 nF capacitor in package size 0402 or smaller. ADC can be AC coupled, if necessary, and should be 100 nF in package size 0402 or smaller. [Figure 129: -2 dB Pi Attenuation Network \(XCZU25DR/XCZU27DR/XCZU28DR/XCZU29DR\)](#) illustrates a representative placement of the AC coupling capacitors in a typical ADC interfacing circuit.

Capacitor values smaller than 100 nF can be used for AC coupling assuming the application does not need to operate to very low frequencies. The customer can simulate the S parameter model of the data converter with the AC coupling capacitor to establish the minimum size capacitance if necessary. The minimum AC coupling capacitance should not be below 100 pF.


DC coupling is also supported for the ADC and DAC. For DAC DC coupling, there are two constraints. First, the DAC requires the correct effective termination resistance of 100Ω differential. Second, the DAC outputs should be biased at the correct DC common mode point. The DC common mode point is 2.07V using DAC_AVTT = 3.0V with RFSoc Gen 3 or RFSoc DFE. It is possible to have an RFSoc Gen 1 compatibility mode but the VOP feature would not be available. Using RFSoc Gen 1 compatibility mode, the common mode point is 1.9V using DAC_AVTT = 2.5V. Refer to *Zynq UltraScale+ RFSoc RF Data Converter LogiCORE IP Product Guide (PG269)* for details and compatibility mode. It is also possible to enable DAC DC coupling with reduced VOP range using DAC_AVTT = 3.0V on RFSoc Gen3/DFE. See *Zynq UltraScale+ RFSoc Data Sheet: DC and AC Switching Characteristics (DS926)* for VOP range specification.

For ADC DC coupling, the ADC VCM output should be coupled with a 100 nF capacitor in package size 0402 or smaller. In addition, the input common-mode level should be taken from the ADC VCM output signal.

Note: The ADC common mode has changed from Gen 1 to Gen 3 from 1.25V to 0.7V. The ADC common mode is designed to drive a high impedance load of $\geq 10\text{ K}\Omega$. If a lower impedance load is required, a buffer should be used on the common mode signal.

A typical ADC DC coupling configuration is shown in the following figure.

Figure 119: Typical ADC DC Coupling Configuration



X20578-032618

Note: The VCM output buffer is only enabled in DC coupled mode.

Isolation Recommendations


Incorporating all the following PCB design techniques results in the highest isolation between traces on the PCB. If significant deviation from the guidelines is planned, AMD recommends performing 3D EM analysis of the board structure to verify that performance at the package launch traces is acceptable for the application.

Note: Separation between traces is not entirely sufficient to achieve the isolation recommendations listed below. The effect of crosstalk rolls off as $1/(1+(s/h)^2)$, where s is trace separation, and h is dielectric thickness. This function does not have a steep roll-off, and because space is limited, other methods must be used to achieve the recommended isolation.

Refer to [Table 125: Signal Integrity Specifications for DAC/ADC Pairs and Clocks](#) for isolation recommendations for ADC and DAC pairs. A number of methods can be used to achieve the best isolation:

- **Micro vias, 3 dB isolation improvement:** Micro vias are preferred over back drilling due to any residual stubbing and the presence of a resonance cavity.
- **Shorter vias, 1 dB isolation improvement per 5 mil (0.13 mm) length:** Vias are a source of crosstalk, which can be minimized by keeping lengths as short as possible.
- **Differential breakout, 3 dB isolation improvement:** Differential breakout, as shown in the following figure, is preferred over single-ended breakout. Single-ended breakout crosstalk is highest within 78 mil (2.0 mm) from the package edge and should be avoided, if possible.


Figure 120: Differential Breakout Example



X19470-122121


- **Guard traces, 10 dB isolation improvement:** Separate each P/N pair by a ground guard-band trace that is 2X the layer height. Ensure that the above and below ground planes extend at least 4H away from the edge of each signal trace. Stitch the guard-band trace to ground every 118 mil \pm 39 mil (3.0 mm \pm 1 mm). All ground stitching should be implemented with through-hole vias. The following figure shows an example of guard-band stitching around P/N pairs.

Figure 121: Example of Ground Stitching around Signal Pairs



The following figure shows closer detail of guard traces and ground stitching vias. Ensure to continue ground stitching around the pins of the RFSoc, as shown in [Sample Stackup](#).


Figure 122: Close View of Guard Traces and Ground Stitching



X19904-121217

The following figure shows a cross-section view of the guard traces.

Figure 123: Cross-Section of Guard Traces



X22756-042619

- **Via-In-Pad, Plated Over (VIPPO), 5 dB isolation improvement:** Plated via-in-pad is preferred over routing a dog-bone trace to a via, because a dog-bone on layer one is subject to more crosstalk from PCB and external sources.

P and N Skew Specifications

The total P to N skew should be no longer than 2000 fs (2.0 ps). This skew specification is broken down into two components:


- ± 1000 fs (1.0 ps) using simulation and taking into account RFSoc package delays and PCB trace length mismatch.
- ± 1000 fs (1.0 ps) due to manufacturing variations.

Simulations should include a complete channel s-parameter analysis that includes the RFSoc package and PCB S-parameter models. The PCB should be deskewed using the simulation analysis to meet the 1000 fs specification. You can request the package S-parameter models for the ADC/DAC channels by talking to your I/O/RF specialist.

Manufacturing variations can be minimized via the following methods:

- **1086 Mechanically Spread (MS) Glass, 700–900 fs skew per inch:** MS glass is spread in both directions to achieve greater homogeneity.
- **NE glass, 400–600 fs per inch:** Replacing E-glass with NE glass can provide for further skew reduction.
- **Image Rotation, 500–100 fs skew per inch:** Rotating the image on the PCB panel, as shown in the following figure, can help offset panel x-y non-homogeneity. The recommended rotation is 12 degrees.

Figure 124: Rotated Panel to Compensate for Glass Non-Homogeneity



X19472-070617

- **Trace Length Matching, ~150 fs:** Matching the P and N trace length to 1 mil (0.0254 mm) can help to achieve approximately 150 fs of skew. Trace length matching should be done as close as possible to any AC coupling capacitors. The following table for suggested specifications.

Table 126: Trace Length Matching


Component	Type of Skew	Specification (mil)	Comments
DAC/ADC data converters	Intra Pair	± 1	Imbalance here impacts the data converter HD2. Can be relaxed to ± 5 if not important on application.
ADC clock	Intra Pair	± 5	
DAC clock	Intra Pair	± 5 (or $\pm 2^1$)	
DAC data converters	Inter Pair	± 10	Matching only required if required by application.
ADC data converters	Inter Pair	± 10	Matching only required if required by application. No matching required between DACs and ADCs.
ADC clocks ²	Inter Pair	± 10	Clock matching is only required when using the MTS feature.
DAC clocks ³	Inter Pair	± 10	Do not need to match to ADC clocks.

Notes:

1. A tighter P to N skew is required on the DAC clock if using an external RF sampling clock. This is not required if using internal PLL.
2. Applied only when MTS is required for ADC group and sample clock feeds directly to tiles.
3. Applied only when MTS is required for DAC group and sample clock feeds directly to tiles.

When routing a 180-degree bend of a P/N pair, the minimum bend radius must be $3H$. This distance allows for a guard trace and ground stitching vias to be placed between the two segments of the inner trace, before and after the bend. Ground stitching vias must be placed on the guard trace in this inner region with the final ground via placed at the end of the guard trace near the signal trace bend. The following figure shows a properly routed P/N trace pair 180-degree bend with guard trace and ground stitching vias.

Figure 125: P/N Trace Pair with 180-Degree Bend




X21203-071918

If routing below 1 mm radius is required in tight areas, calculate the electrical length by following the path of the inner radius of the arc. Most EDA tools calculate the delay based on the center line radius of the arc segment. For example, for a 4 mil wide trace with an arc segment with a 10 mil center line radius, the signal actually follows the radius of $10 - (4/2) = 8$ mil. The effect of the signal following the inner radius of the arc begins at roughly 10 MHz and remains constant enough that no further optimization is needed for higher frequencies up to 10 GHz.

When P/N traces transition layers, two ground vias must be added per signal for impedance control, noise isolation, and P/N skew control. The ground vias must be within 1 mm of the signal vias. The following figure illustrates the via configuration for a signal layer transition.


Figure 126: Via Configuration for Signal Layer Transition



X21204-071918

P/N signal traces should be routed a minimum of $0.5H$ away from any anti-pad or cutout in the ground planes above or below the signal traces. The left side of the following figure illustrates an incorrectly routed signal trace. The right side of following figure illustrates the correct routing of a signal trace.

Figure 127: Signal Trace Routing



X21205-071918

Signal guard traces must be continuous with no breaks and must have proper ground stitching. An improperly implemented guard trace could result in additional P/N skew as well as increased noise susceptibility.

Insertion Loss Recommendations


Insertion loss is recommended to be better than -1.5 dB from the BGA ball to its balun (or similar interfacing component). Most board materials are capable of achieving better than -1.5 dB of insertion loss at frequencies up to around 6 GHz. Insertion loss mitigation does come with a potentially significant cost multiplier, so care should be taken in choosing the proper material. See [Appendix C: Material Properties and Insertion Losses](#) for insertion loss values based on trace widths and frequencies.

Return Loss Recommendations

Return loss is recommended to be better than -20 dB across the bandwidth of interest. Return loss can be mitigated with the following methods:

- **Circular ground antipad on layer one, 5 dB improvement:** A circular antipad on layer one should be used for P/N pairs, as shown in the following figure. The microvia drill diameter should be 5 mil, with a 25 mil pad diameter and 35 mil anti-pad diameter.

Figure 128: Circular Antipad



X20582-032618


- **Micro via, 7 dB:** Back drilling can leave residual metal and a resulting resonance cavity.

Note: AMD does not recommend using ground flooding on signal layers as this might cause the board to warp, causing strain on soldered components. For Gen 1 devices, AMD recommends adding a -2 dB differential resistive Pi attenuator to optimize the RL of the device. This Pi network can be placed anywhere on the PCB between the package BGA balls and the balun. See the following figure.

Pi Network for Improved Return Loss in Gen 1 Devices (XCZU25DR/XCZU27DR/XCZU28DR/XCZU29DR)

AMD ADCs and DACs have a nominal on-die 100Ω differential termination, which generally provides sufficient return loss performance for most RF applications. However, for applications requiring better than -10 dB return loss up to 4 GHz, AMD recommends placing -2 dB balanced Pi pad attenuators on the PCB. These Pi pad attenuators can be placed anywhere on the PCB between the ADC/DAC BGA balls of the package and the baluns. A representative -2 dB Pi pad attenuator for a differential ADC/DAC is shown in the following figure.

Figure 129: -2 dB Pi Attenuation Network (XCZU25DR/XCZU27DR/XCZU28DR/XCZU29DR)



X19905-121217

The component descriptions are listed below with recommended part numbers:

- R2 = R3 = 11.5Ω (Panasonic part number [ERJ-1GEF11R5C](#))
- R1 = R4 = 866Ω (Panasonic part number [ERJ-1GEF8660C](#))

AMD recommends the following characteristics for the discrete resistive components:

- 0201 body size or smaller to minimize parasitic inductance.
- $\pm 1\%$ resistor tolerance or better to minimize insertion loss variation
- 100 ppm/°C temperature coefficient or better to minimize variation of the resistance value over temperature

For purposes of PCB routing, there should be a minimum of four ground vias for each and every signal via. The length of the vias and trace routing between components should be minimized as much as possible. Placing the Pi pad attenuators underneath cans minimizes crosstalk between them.



IMPORTANT! For all generations of RFSoc later than Gen 1, there is no need for an attenuator due to improvements in return loss performance of both DAC and ADC.

SYSREF

In each Zynq UltraScale+ RFSoc there is one dedicated input SYSREF pin pair located in DAC tile 0. This SYSREF signal is used for multi-tile and multi-chip channel synchronization. This differential input pair is referred to as the Analog_SYSREF. To fully implement the synchronization features in the RFSoc, an additional SYSREF signal is needed to drive into the PL fabric through a pair of dedicated clock inputs. This pair is referred to as the PL_SYSREF. It


is possible to use the multi-tile sync feature even if not all available tiles in the device are being used. In such an event, DAC tile 0 should be one of the active tiles. ADC tile 0 should also be one of the active tiles if multi-tile sync for the ADCs is a requirement in the application. All DAC tiles are synchronized to DAC tile 0, and all ADC tiles are synchronized to ADC tile 0. That is why tile 0 is required to be active for the respective data converter types.

AMD recommends that the user prioritize the lower indexed tiles over the higher ones if not all tiles are to be enabled. For example, if only two tiles are used, tile 0 and tile 1 should be enabled, while tile 2 and 3 are disabled. Analog_SYSREF can also be used for synchronizing the phase of the digital up converter (DUC) numerically controlled oscillators (NCOs) and digital down converter (DDC) NCOs, respectively, in addition to synchronizing the overall delay of each tile (done with the MTS function). The NCO synchronization is a separate process from the MTS synchronization and only applicable if the fine complex mixer is used in the application. Refer to the Multi-Converter Synchronization section in *Zynq UltraScale+ RFSoc RF Data Converter LogiCORE IP Product Guide (PG269)* for the detailed theory of operations and connection diagram.

The Analog_SYSREF is internally terminated differentially at 100Ω. It can be used in AC coupled mode or DC coupled mode. For applications that require NCO synchronization across multiple devices, DC coupled mode is required along with the ability to generate the SYSREF signal in a pulse mode. Refer to *Zynq UltraScale+ RFSoc RF Data Converter LogiCORE IP Product Guide (PG269)* for additional guidance. Refer to *Zynq UltraScale+ RFSoc Data Sheet: DC and AC Switching Characteristics (DS926)* for further requirements.

The following figure shows the Analog_SYSREF receiver equivalent circuit. In the figure, the driver must be capable of sourcing and sinking common mode current through a 1 kΩ resistor. The driver must also set its own common mode bias.

Figure 130: Analog_SYSREF Receiver Equivalent Circuit



For AC coupling, the SYSREF input sets the common mode voltage internally as shown in the previous figure. For DC coupling, the SYSREF source driver must be capable of sourcing and sinking common mode current through the internal 1 kΩ resistor. The SYSREF driver must set its own common mode bias. AMD suggests using DC coupling to avoid glitches in the single shot or gapped clock scenarios.

If synchronizing ADC and DAC tiles with SYSREF, the frequency must be an integer sub-multiple of the following equation.

Equation 3: SYSREF Frequency

$$GCD\left(\frac{DAC_{SampleRate}}{16}, \frac{ADC_{SampleRate}}{16}\right)$$

An example calculation for SYSREF is shown below.

$$ADC_{SampleRate} = 3.93216 \text{ GHz}$$

$$DAC_{SampleRate} = 4.91520 \text{ GHz}$$


$$GCD\left(\frac{ADC_{SampleRate}}{16}, \frac{DAC_{SampleRate}}{16}\right) = 61.44 \text{ MHz}$$

$$\text{Integer Sub - Multiple AND} < 10 \text{ MHz} = 7.68 \text{ MHz}$$

For full functionality of the synchronization features, the PL_SYSREF should meet the following requirements:

- PL_SYSREF should be the same frequency as the Analog_SYSREF.
- PL_SYSREF must be a sub-multiple of the PL frequencies that interface to the DAC and ADC FIFOs.
- PL_SYSREF and PL clock must meet setup and hold to ensure the PL clock can be used to deterministically capture the rising edge of the PL_SYSREF.
- The same clock quality (if not the same clock) should be used for the PL_SYSREF as the Analog_SYSREF. Refer to the PL GPIO electrical requirements for coupling, voltage swing, and termination.
- If the MMCM is used to generate the final PL clocks for the DAC FIFO interface and ADC FIFO interface, the 0-delay (0 phase) feature of the MMCM should be used in addition to using flip-flops to synchronize the PL_SYSREF capture as shown in the example below.
- For multi-chip synchronization, it is very important to ensure alignment of Analog_SYSREF, PL_SYSREF, PL_Ref_clk, and data converter clocks between chips to ensure a common and matched timing reference distribution to all chips.

Figure 131: PL_SYSREF Capture



X20579-032618

Calibration Resistors

Two separate pins exist for ADC and DAC calibration:

- **ADC External Resistor:** Tie to GND through a 2.49 k Ω resistor as close to the RFSoc pin as possible.
- **DAC External Resistor:** Tie to GND through a 2.49 k Ω resistor as close to the RFSoc pin as possible.

Sample Stackup

This section shows a sample stackup for the top six layers of the board. The trace dimensions in the example are approximate trace geometries based on the sample stackup using a typical dielectric thickness between the signal/ground layers to achieve a 100 Ω differential impedance. Customers are requested to run their own SI analysis to arrive at these numbers based on their specific stackup information taking manufacturing tolerances into account.

Using Isola I-Tera material with these parameters:


- Pre-preg, $h = 6.0$ mils, $D_K = 3.17$, 72%, 2×10^6 MS
- Core, $h = 5.0$ mils, $D_K = 3.08$, 72%, 2×10^6 MS
- 2 μ m surface roughness for copper to minimize insertion loss

The pre-preg measures 6.0 mils before lamination. It will likely measure around 5.5 mils after lamination.

PCB Stackup

The PCB stackup is shown in the following figure.

Figure 132: PCB Stackup




X22697-042219

Route the DAC and ADC closest to the package edge (on right above) on layer 2. Route the DACs and ADCs toward the middle of the package on layer 4. Clock signals (layer 6 in this sample stackup) can be routed on the DAC/ADC layer (layer 2 or 4 in this sample stackup) if space allows.

Layer 1 (Top Layer) in BGA Area for DACs/ADCs

Figure 133: Layer 1 in BGA Area for DACs/ADCs



X22698-042219


Guidance:

- Land pads = 20 mils for signal and grounds for BGA balls

- Anti-pads for signals = 34 mils
- Anti-pads for grounds = 30 mils
- Micro via for layer 2 routes = 6 mil drill
- Via for layer 4 routes = 8 mil drill (back-drilled)
- Add extra horizontal ground vias between ground solder balls at 20 mil pitch (see the previous figure)
- Ground vias = 8 mil drill
- Land pad for extra grounds = 16 mils
- Thermal relief for ground pads (5 mil width vertical route and 9 mil width horizontal route)

Layer 2 in BGA Area for DACs/ADCs

Figure 134: Layer 2 in BGA Area for DACs/ADCs



X22699-042219

Signal P exists at 45° north-east. Signal N exists via the pad at due west. A neck down trace length of 78 mils with 4 mil width is needed to tune the return loss up to 6 GHz. Remove all unused pads from DACs and ADCs which are not routed on layer 2.


Air gap between the P and N signals should be 15 mil. The ground guard trace should be 18 mil wide. Air gap between signal and ground should be 15 mil. Trace routing in the BGA area should be located at the center line of the signal trace midway between the anti-pads on layer 1. In summary:

- Trace width = 5.8 mils
- Land pad for micro-via = 16 mils

- Land pad for all grounds = 18 mils

Layer 3 in BGA Area for DACs/ADCs

Figure 135: Layer 3 in BGA Area for DACs/ADCs




X22700-042219

Guidance:

- There should be no anti-pads underneath the micro-vias used for the layer 2 routes.
- Remove all unused pads.
- Anti-pads for vias should be 35 mils.
- Land pad for all grounds should be 18 mils.

Layer 4 in BGA Area for DACs/ADCs

Figure 136: Layer 4 in BGA Area for DACs/ADCs




X22701-043019

Guidance:

- Trace width = 5.8 mils. Need a neck down trace of 42 mil length with 4 mil width to tune the RL to -20 dB @ 6 GHz.
- Air gap between signal and ground = 15 mils. Air gap between P and N signals = 15 mils. Width of ground guard trace is 18 mils.
- Remove all unused pads.
- Land pads for all grounds = 18 mils.
- Land pad for signal via = 18 mils.

Layer 5 in BGA Area for DACs/ADCs

Figure 137: Layer 5 in BGA Area for DACs/ADCs




X22702-042219

Here, land pads for all grounds are 18 mils. Remove all unused signal pads.

Layer 6 for Clock Routes

Figure 138: Layer 6 for Clock Routes




X22703-042219

Guidance:

- Trace width of 5.8 mils. Neck down trace of 40 mils to achieve RL of -15 dB @10 GHz is required.
- Air gap between traces should be 15 mils.
- Signal gap to ground should be 15 mils.
- Ground guard trace width should be 18 mils.

Guard Traces along the Length of the Trace Routes

Figure 139: Guard Traces along Trace Route Length




- Guard traces of 18 mil width must be stitched every 4 mm width of 8 mil drilled vias through the PCB.
- Guard traces must be 15 mils away from signal traces.
- Each individual ground stitch via can have a true tolerance of ± 1 mm from nominal.
- Trace width = 5.8 mils.

Neck Down Trace

Neck down trace is used to achieve a return loss of 20 dB @ 6 GHz. The exact length of the neck down trace should be defined through simulation.

Layer 2 Routing for AC Coupling Capacitor Inputs and Outputs

Figure 140: Layer 2 Routing for AC Coupling Capacitor Inputs and Outputs




X22705-042219

A neck down trace of 48 mil length with 4 mil width for 0201 capacitors is required on the side between the data converter and capacitor when routing on layer 2.

Layer 4 Routing for AC Coupling Capacitor Inputs and Outputs

Figure 141: Layer 4 Routing for AC Coupling Capacitor Inputs and Outputs




X22706-042219

A neck down trace of 23 mil length with 4 mil width for 0201 capacitors is required on the side between the data converter and capacitor when routing on layer 4.

Layer 2 Between AC Capacitors and Baluns

Figure 142: Layer 2 Between AC Capacitors and Baluns




X22707-042219

A neck down trace of 52 mil length is required between the 0201 capacitors and Anaren 0404 baluns when routing on layer 2.

Layer 4 Between AC Capacitors and Baluns

Figure 143: Layer 4 Between AC Capacitors and Baluns



X22708-041919

A neck down trace of 52 mil length is required between the 0201 capacitors and Anaren 0404 balun when routing on layer 4. The neck down trace to the balun is the same for layer 2 and layer 4 so that the same size RF shield can be used.

Unused ADC & DAC Pins

- When an ADC channel is not enabled in the IP and the channel will never be in use over its lifetime, the ADC VIN pins should be tied to GND directly.
- When an ADC channel is enabled in the IP but is not being driven by anything on the PCB, the RF input can be left floating. The user needs to ensure AC coupled mode is chosen. This might be required for certain SYSREF distribution use cases where the ADC channels can be powered either up or down via software while the channel is enabled in the IP.
- All unused DAC outputs can be left floated (there is an internal pull-down on the outputs).
- All unused clock and SYSREF pins can be floated.


Power Regulation and Decoupling for ADC and DAC Supplies

AMD supports powering the RFDC analog Supplies from either low dropout (LDO) or higher efficiency switched mode regulator solutions that meet the ripple requirements. A typical LDO based power distribution is shown in the following figure. Later sections deal with powering off switched mode regulators.

AMD recommends separate power supplies be used for ADC and DAC. For certain applications where only ADCs are being used along with MTS, it is necessary to provide power to the DAC tiles but not enable the tiles. Under these circumstances, the DAC power supplies can be tied to the closest ADC power supply.

 **IMPORTANT!** *The guidelines provided as part of this document are for reference only. However, AMD recommends you do a detailed simulation analysis to validate the various specifications prior to releasing the board for fabrication. AMD can provide the necessary ADC/DAC S-parameter models for the die and package trace extracted using a full 3D field solver to facilitate this analysis. You can use these models in any EDA tool like ADS from Keysight to perform a detailed frequency-domain analysis to analyze insertion loss, return loss, crosstalk performance, and PN balance, respectively. You can request the die package models by talking to your local I/O/RF specialist.*

Figure 144: Power Regulation for Analog Supplies



X19474-050819

Recommended Linear Regulators

The following table lists regulators that AMD recommends using on the ADC and DAC RFSoc rails.

Table 127: Recommended Linear Regulators

Vendor	Devices
Intersil/Renesas	ISL80101-A ISL80102 ISL80103

ADC and DAC Voltage Supply Specifications

The ADC and DAC voltage supply specifications for Gen 1, Gen 2, and Gen 3 devices are listed in the following tables.

Table 128: ADC and DAC Voltage Supply Specifications for Gen 1 and Gen 2 Devices¹

Supply	Nominal Voltage (V)	Tolerance (%) ²	Frequency Range (MHz)	Maximum Supply Ripple (mVpp) ³
ADC_AVCC	0.925	±3	0.1–15	0.25
ADC_AVCCAUX	1.8	±3	0.1–15	11.03
DAC_AVCC	0.925	±3	0.1–15	0.40
DAC_AVCCAUX	1.8	±3	0.1–15	2.00
DAC_AVTT	2.5/3.0 ⁴	±3	0.1–15	8.94
VCCINT_AMS	0.85	±3	0.1–15	20.00

Notes:

1. For the maximum current, refer to the Power Design Manager (PDM) tool (download at www.amd.com/power).
2. The tolerance percentage is for the switching regulator that feeds the VRM.
3. Output of the VRM.
4. DAC_AVTT should be set to 2.5V if used in 20 mA mode, and 3.0V if used in 32 mA mode in Gen 1 and Gen 2. AMD recommends DAC_AVTT be set to 3.0V in Gen 3 to enable the VOP function. Refer to *Zynq UltraScale+ RFSoc RF Data Converter LogiCORE IP Product Guide* (PG269) for details and compatibility mode.

Table 129: ADC and DAC Voltage Supply Specifications for Gen 3 Devices¹

Supply	Nominal Voltage (V)	Tolerance (%) ²	Frequency Range (MHz)	Maximum Supply Ripple (mVpp) ³
ADC_AVCC	0.925 (1.01V DFE)	±3	0.1–15	0.25
ADC_AVCCAUX	1.8	±3	0.1–15	1.2
DAC_AVCC	0.925	±3	0.1–15	0.32
DAC_AVCCAUX	1.8	±3	0.1–15	1.0
DAC_AVTT	2.5/3.0 ⁴	±3	0.1–15	2.0

Table 129: ADC and DAC Voltage Supply Specifications for Gen 3 Devices¹ (cont'd)

Supply	Nominal Voltage (V)	Tolerance (%) ²	Frequency Range (MHz)	Maximum Supply Ripple (mVpp) ³
VCCINT_AMS	0.85	±3	0.1–15	20.0


Notes:

1. For the maximum current, refer to the Power Design Manager (PDM) tool (download at www.amd.com/power).
2. The tolerance percentage is for the switching regulator that feeds the VRM.
3. Output of the VRM.
4. DAC_AVTT should be set to 2.5V if used in 20 mA mode, and 3.0V if used in 32 mA mode in Gen 1 and Gen 2. AMD recommends DAC_AVTT be set to 3.0V in Gen 3 to enable the VOP function. Refer to *Zynq UltraScale+ RFSoc RF Data Converter LogiCORE IP Product Guide (PG269)* for details and compatibility mode.

Decoupling Capacitor Antipad to Reduce Insertion and Return Loss

It is recommended to place a rectangular antipad around each ADC and DAC decoupling capacitor on the top and bottom layer, as shown in the following figure.

Figure 145: Decoupling Capacitor Antipad



Voltage Sensing


To compensate for PCB IR losses, all of the ADC and DAC voltage rails need to properly use the sense features of their respective regulators, as listed below:

- ADC_AVCC
- DAC_AVCC
- ADC_AVCCAUX
- DAC_AVCCAUX

- DAC_AVTT
- VCCINT_AMS

Sense line connections should not use a dedicated spy hole pin on the RFSoc, because all power supply pins should be fully used and connected to their respective planes. Instead, a "Force" solution should be used, as shown in the following figure.

Figure 146: Recommended Sense Line Circuitry



X19544-080217

Notes related to the previous figure:

- R_{PDN} : Power distribution network resistance (PCB parasitics).
- R_{SENSE} : Discrete 2 mΩ precision sense resistor.
- R_{PRTC} : Discrete 10Ω regulator protection resistor.

Unused ADC and DAC Power Pins

- The DAC_AVTT supply can be tied to the DAC_AVCCAUX 1.8V supply when the DAC is not enabled.
- ADC and DAC tiles all share global power supplies, so all tiles are powered even if not enabled.
 - For example, all DAC tiles connect to one DAC_AVCC supply, one DAC_AVTT supply, and one DAC_AVCCAUX supply. All ADC tiles connect to one ADC_AVCC supply and one ADC_AVCCAUX supply.

Powering RFSocS with Switch Regulators

For higher power efficiency, AMD recommends using full switched mode power supply solutions. The RFSocS range does not require inefficient supply solutions based on low dropout (LDO) assuming the switched mode regulators can meet the ripple requirements. Refer to power solution reference designs on the ZCU216 and ZCU208 for examples of full switched mode regulator solutions that do not impact the SFDR of the data converters and offer superior power efficiency over traditional LDO based implementations.

Power Delivery Network Design for Time Division Duplex

For details on the time division duplex (TDD) feature, refer to *Zynq UltraScale+ RFSocS RF Data Converter LogiCORE IP Product Guide (PG269)*. Usage of the TDD mode requires a power distribution network (PDN) on the PCB with very fast step response settling times. The TDD feature puts some of the RFDC IP data converter supplies in a low-power state when in non-active mode, and switches back to full-power mode while in active mode in under 1 μ s. The PDN needs to be able to respond to this step response with a similar settling time to maximize the power saving offered by the TDD mode. Customers can use the Power Design Manager (PDM) tool (download at www.amd.com/power) to calculate the step change in current by using the duty cycle set at 100% and 0% to get the full power and low power currents. A typical fast settling reference design can be found on the ZCU208 or ZCU216 evaluation boards. TDD power saving mode is only supported in Gen 3 and DFE devices. The glitches in the PDN introduced by the TDD switching should be kept within the tolerance of the $\pm 3\%$ of the supplies.

PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC

This chapter lists layout guidelines specific to the PS MIO and PS-GTR interfaces in the Zynq UltraScale+ MPSoC.

Boot Mode

- Connect the boot mode pins to a 4.7 k Ω or lower pull-up resistor to VCCO_PSIO[3] or pull-down to ground depending on the desired setting.
 - If multiple switchable boot modes are desired, connect one pull-up/down resistor to the mode pin and place a jumper on the other side of the resistor to select between pull-up or pull-down.
 - An easily switchable boot mode configuration is recommended for debug ease-of-use.
-

CAN

- PCB and package skew between the TX/RX and clock should be within ± 100 ps.
 - A level shifter must be implemented if using a CAN PHY that operates at a voltage higher than VCCO_PSIO, for example, 5.0V.
-

DisplayPort

DisplayPort is a PS-GTR transceiver interface, and rules for PS-GTR transceiver connections can be found in [PS-GTR Transceiver Interfaces](#). In addition, the DP_OE, DP_AUX_IN, and DP_AUX_OUT signals need to be connected to a bidirectional LVDS buffer like the Fairchild FIN1019MTC. For additional details on connectivity, refer to the VESA DisplayPort Standard Version 1, Section 3.4, AUX Channel (see www.vesa.org).

eMMC

- Pull-up resistors are required on data[0-7] per JEDEC specification [JESD84-B451](#).
- A 10 k Ω pull-up resistor is recommended on CMD.
- A 30 Ω series resistor should be placed on the CLK, CMD, and DATA lines, as close to the MIO pins as possible.

Standard and High-Speed SDR Interfaces

PCB and package skew between data/cmd and clock should be within ± 100 ps.

HS200 (200 MHz) and High-Speed DDR Interfaces

PCB and package skew between data/cmd and clock should be within ± 50 ps.

Ethernet MAC RGMII

Delay skew for TXD/RXD[3:0] and CTL to clock delay should be within ± 50 ps including package time.

Ethernet MAC SGMII

Ethernet SGMII is a PS-GTR transceiver interface, and rules for PS-GTR transceiver connections can be found in [PS-GTR Transceiver Interfaces](#).

I2C

- Place 4.7 k Ω pull-up resistors at the far end of the SCL and SDA lines, furthest from the Zynq UltraScale+ MPSoC.
- A level-shifter/repeater might be required depending on the particular multiplexers used.

JTAG

- Place 4.7 k Ω pull-up resistors on the TMS, TCK, and TDI lines.
- All interconnected JTAG chains should be powered by the same voltage rail. Or, if different rails are used, they should all be powered at the same time.

Note: Not all third-party JTAG modules are supported by AMD.

NAND Memory Controller

SDR Mode

- PCB and package skew between ALE/CE/CLE/IO[7:0] and WE/RE should be within ± 100 ps.
- Place 4.7 k Ω pull-up resistors on CE and RB near the NAND device.

DDR Mode (100 MHz)

- PCB and package skew between ALE/CE/CLE and CLK should be within ± 50 ps.
 - To operate in mode 0 to mode 4, PCB and package skew between CLK and DQS should be within ± 50 ps.
 - To operate in mode 5 at 100 MHz, PCB and package skew between CLK and DQS should be a minimum of 400 ps + 50 ps. CLK should be skewed with respect to DQS by 400 ps.
 - PCB and package skew between DQ and DQS should be within ± 25 ps.
 - Place 4.7 k Ω pull-up resistors on CE and RB near the NAND device.
 - For devices without a DQS pin, the MPSoC MIO pin can be left unconnected or used for another purpose.
-

PCIe

PCIe[®] is a PS-GTR transceiver interface, and rules for PS-GTR transceiver connections can be found in [PS-GTR Transceiver Interfaces](#).

Note: Refer to *UltraScale Devices Gen3 Integrated Block for PCI Express LogiCORE IP Product Guide (PG156)* for detailed information regarding lane assignments and other PCIe-specific information.

PS_INIT_B, PS_PROG_B, and PS_DONE

- Connect PS_INIT_B to a 4.7 k Ω pull-up resistor to VCCO_PSIO[3].
 - PS_INIT_B is open drain and should not be driven during logic built-in self test (LBIST).
- Connect PS_PROG_B to a 4.7 k Ω pull-up resistor to VCCO_PSIO[3].
 - PS_PROG_B is open drain and should not be driven during LBIST.
- Connect PS_DONE to a 4.7 k Ω pull-up resistor to VCCO_PSIO[3].



IMPORTANT! PS_PROG_B is recommended to be independent of PS_POR_B. See Zynq UltraScale+ Device Technical Reference Manual ([UG1085](#)) for details regarding PS_PROG_B and PS_POR_B.



IMPORTANT! It is recommended to limit the external loading on PS_INIT_B so as not to excessively delay its rise time (see [Answer Record 70504](#)).

PS Reference Clock

- PS_REF_CLK is a single-ended LVCMOS signal.
- It can be terminated at the source with an appropriate series resistor or at the receiver side with a Thevenin split termination.
- Signal integrity simulation for this critical signal is recommended.

PS Reset (External System Reset and POR Reset)

- Connect PS_SRST_B to a 4.7 k Ω pull-up resistor to VCCO_PSIO[3] near the Zynq UltraScale+ MPSoC.
- Connect PS_POR_B to a 4.7 k Ω pull-up resistor to VCCO_PSIO[3] near the Zynq UltraScale+ MPSoC.
- Refer to [Power Supply Sequencing](#) in Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics ([DS925](#)) for rules regarding PS_POR_B during the power-on sequence.

QSPI

- The clock, data, and SS lines are recommended to have matched lengths to facilitate meeting setup and hold times.
- PCB and package delay skew for I/O[3:0] and SS lines relative to CLK should be within ± 50 ps.
- Keeping the clock and data lines equal provides greater immunity to undesirable setup and hold time effects.
- It is highly recommended to perform a signal integrity analysis on the clock line at the near (close to the Zynq UltraScale+ MPSoC) and far ends.
- For optimum performance, limit trace delays to less than 500 ps.
- Place 4.7 k Ω pull-up resistors on the HOLD, WP, and CS lines.
- Keep MIO[6] unconnected for higher FQSPICLK1 or FQSPICLK2 operating frequencies (>40 MHz). This allows the loopback feature to work properly.

Real-Time Clock

- See the PS RTC Crystal Requirements table in *Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)* for full specifications.
- If the real-time clock is not being used, connect PADI to ground and leave PADO floating.

SATA

SATA is a PS-GTR transceiver interface, and rules for PS-GTR transceiver connections can be found in [PS-GTR Transceiver Interfaces](#).

SD/SDIO

- A 30 Ω series resistor should be placed on the CLK, CMD, and DATA lines, as close to the MIO pins as possible.
- Ensure SD signals have a total delay of less than 1.3 ns (including package delays) to ensure manual tuning can find the valid input delay.
- Skew between SDx_DATA[3:0]/SDx_CMD and SDx_CLK should be within ± 50 ps.

- A level shifter might be required depending on the particular voltages used on the Zynq UltraScale+ MPSoC and SD chip.
 - When using an SD-compliant level shifter, the PCB trace is broken up into two segments. Match the skew of each segment as follows:
 - Segment 1: Zynq UltraScale+ MPSoC/RFSoc to SD level translator: SDx_DATA[3:0]/SDx_CMD to SDx_CLK: ± 25 ps
 - Segment 2: Level translator to SD card connector: SDx_DATA[3:0]/SDx_CMD to SDx_CLK: ± 25 ps
- Asynchronous signals CDn and WPn have no timing relationship to CLK.
- The CDn and WPn lines should both be pulled up with their own 4.7 k Ω resistors to the MIO I/O voltage. When using Micro-SD, WPn and CDn can be no connects.
- A 10 k Ω pull-up resistor should be added to DAT3 on the SD card side of the level shifter.
Note: An external pull-up resistor is not required if one is already present on the level shifter device.
- When using the SDIO 0/1 power control signal, use an external 1 k Ω to 10 k Ω pull-down resistor to get a 1 ms reset pulse according to the SD card 3.01 specification.
- For level shifters without a direction pin, the MPSoC "DIR" pin can be left floating or pulled-down with a 100 k Ω resistor. The DIR pin should not be used for any other purpose.
- For level shifters with 0.4 mm pitch, an adapter can be used, such as from Aries Electronics.

SPI

- PCB and package skew between the master-output, slave-input (MOSI)/master-input, slave-output (MISO) and clock should be within ± 100 ps.
- Place a 4.7 k Ω pull-up resistor on the SS pin near the serial peripheral interface (SPI) device.

Trace Port Interface Unit

When operating the trace port interface unit (TPIU) in MIO mode, the trace clock output should be delayed by approximately one half clock period. This can be done on the PCB, or by the debugging device (ARM_DSTREAM, Lauterbach, or Agilent).

Triple Time Counter

PCB and package skews between WAVE_OUT and CLK should be within ± 100 ps.

UART

Keep MIO trace delays below 1.30 ns.

USB 2.0

ULPI Interface (60 MHz)

- PCB and package delays should be kept to 1.30 ns or below.
 - PCB and package delay skews for DATA[7:0]/DIR/NXT/STP and CLK should be within ± 100 ps.
 - For optimum signal integrity, add a 30 Ω series resistor to the DATA and STP lines near the Zynq UltraScale+ MPSoC.
-

USB 3.0

USB 3.0 is a PS-GTR transceiver interface, and rules for PS-GTR transceiver connections can be found in [PS-GTR Transceiver Interfaces](#).

Watchdog Timer

PCB and package skews between WAVE_OUT and CLK should be within ± 100 ps.

PS-GTR Transceiver Interfaces

PS-GTR transceivers are analog circuits that require special attention when being designed for and implemented on a printed circuit board. Besides an understanding of the functionality of the device pins, a design that performs optimally requires attention to issues such as device interfacing, transmission line impedance and routing, power supply design filtering and distribution, component select, and PCB layout and stackup design.

These PS interfaces use PS-GTR transceivers:

- DisplayPort
- Ethernet SGMII
- PCIe®
- SATA
- USB3.0


Pin Description and Design Guidelines

The following table and figure detail the PCB guidelines for PS-GTR transceiver interfaces.

Table 130: PS-GTR Transceiver Quad Pin Descriptions

Pins	Direction	Descriptions
PS_MGTREFCLK[3:0]P PS_MGTREFCLK[3:0]N	In (Pad)	Differential clock input pin pair for the reference clock of the PS-GTR transceiver.
PS_MGTRRXP[3:0] PS_MGTRRXN[3:0]	In (Pad)	RXP and RXN are the differential input pairs for each of the PS-GTR receivers.
PS_MGTRTXP[3:0] PS_MGTRTXN[3:0]	Out (Pad)	TXP and TXN are the differential output pairs for each of the PS-GTR transmitters.
PS_MGTRREF	In (Pad)	Calibration input pin for the termination resistor calibration circuit.
PS_MGTRAVCC	In (Pad)	MGTRAVCC is the receiver, transmitter, and clock core power supply. The nominal voltage 0.85 VDC.
PS_MGTRAVTT	In (Pad)	MGTRAVTT is the I/O supply for the transmitter and receiver. The nominal voltage is 1.8 VDC.

Figure 147: PS-GTR External Power Supply and Calibration Resistor Connections



Note: The voltage values in the previous figure are nominal. See the *Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)* for values and tolerances.

Reference Clock

Overview

This section focuses on the selection of the reference clock source or oscillator. An oscillator is characterized by:

- Frequency range
- Output voltage swing
- Jitter (deterministic, random, peak-to-peak)
- Rise and fall times
- Supply voltage and current
- Noise specification
- Duty cycle and duty-cycle tolerance
- Frequency stability

These characteristics are selection criteria when choosing an oscillator for PS-GTR transceiver design. The following figure illustrates the convention for the single-ended clock input voltage swing, peak-to-peak, as used in the PS-GTR transceiver portion of the *Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)*. This figure is provided to show the contrast to the differential clock input voltage swing calculation shown in [Figure 149: Differential Clock Input Voltage Swing, Peak-to-Peak](#).

Figure 148: Single-Ended Clock Input Voltage Swing, Peak-to-Peak






Figure 149: Differential Clock Input Voltage Swing, Peak-to-Peak



The following figure shows the rise and fall time convention of the reference clock.


Figure 150: Rise and Fall Times



UG583_08_21_031716

The following figure illustrates the internal details of the MGTREFCLK input buffer. The dedicated differential reference clock input pin pair MGTREFCLKP/MGTREFCLKN is internally terminated with a 100Ω differential impedance. The common mode voltage of this differential reference clock input pair is ground or nominal 0.0V. See the *Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics* ([DS925](#)) for exact specifications.

Figure 151: MGTREFCLK Input Buffer Details




UG583_08_22_041516

Reference Clock Interface

LVDS

The following figure shows how an LVDS, current mode logic (CML), or high-speed current steering logic (HCSL) oscillator is connected to a reference clock input of the PS-GTR transceiver.


Figure 152: Interfacing an LVDS, CML, or HCSL Oscillator to the Zynq UltraScale+ MPSoC PS-GTR Transceiver Reference Clock Input



LVPECL

The following figure shows how a low-voltage positive emitter-coupled logic (LVPECL) oscillator is connected to a reference clock input of the PS-GTR transceiver.

Figure 153: Interfacing an LVPECL Oscillator to the Zynq UltraScale+ MPSoC PS-GTR Transceiver Reference Clock Input



UG583_08_24_041416

AC Coupled Reference Clock

AC coupling of the oscillator reference clock output to the PS-GTR transceiver reference clock input serves multiple purposes:

- Blocking a DC current between the oscillator and the PS-GTR transceiver reference clock dedicated reference clock input pins (which also reduces power consumption of both parts)
- Common-mode voltage independence
- The AC coupling capacitor forms a high-pass filter with the on-chip termination that attenuates a wander of the reference clock

To minimize noise and power consumption, external AC coupling capacitors between the sourcing oscillator and the PS-GTR transceiver dedicated reference clock input pins are required.

Unused Reference Clocks

If the reference clock input is not used, the reference clock input pins can be left unconnected.

Reference Clock Power

The PS-GTR transceiver reference clock input circuit is powered by PS_MGTRAVCC. Excessive noise on this supply has a negative impact on the performance of any PS-GTR transceivers that use the reference clock from this circuit.

Power Supply Noise

The PS-GTR transceiver power supplies PS_MGTRAVCC and PS_MGTRAVTT must not have noise greater than 10 mV_{pp} from 10 kHz to 80 MHz. Any noise on these power supplies that exceeds this requirement might affect the performance of the transceiver. If the noise on the power supply is less than or equal to this requirement, the performance of the transceiver is not negatively impacted.

PCB Design Checklist

The following table is a checklist of items that can be used to design and review any Zynq UltraScale+ MPSoC PS-GTR transceiver schematic and layout.

Table 131: PCB Design Checklist for PS-GTR

Pins	Recommendations
PS_MGTREFCLK0P PS_MGTREFCLK0N PS_MGTREFCLK1P PS_MGTREFCLK1N PS_MGTREFCLK2P PS_MGTREFCLK2N PS_MGTREFCLK3P PS_MGTREFCLK3N	<ul style="list-style-type: none"> • Use AC coupling capacitors for connection to oscillator. • For AC coupling capacitors, see Reference Clock Interface. The recommended value for LVDS is 10 nF. • Reference clock traces should be provided enough clearance to eliminate crosstalk from adjacent signals. • Reference clock oscillator output must comply with the minimum/maximum input amplitude requirements for these input pins. See the PS-GTR Transceiver Clock DC Input Level Specifications table in <i>Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)</i>. • PS-GTR reference clock PS_MGTREFCLKxP/N should be stable before releasing PS_POR_B. • If the reference clock input is not used, leave the associated pin pair unconnected or tie to ground. • Traces should be routed to 50Ω single-ended/100Ω differential.

Table 131: PCB Design Checklist for PS-GTR (cont'd)

Pins	Recommendations
PS_MGTRRX0/ PS_MGTRRXN0 PS_MGTRRX1/ PS_MGTRRXN1 PS_MGTRRX2/ PS_MGTRRXN2 PS_MGTRRX3/ PS_MGTRRXN3	<ul style="list-style-type: none"> Use AC coupling capacitors for connection to the transmitter. The recommended value for AC coupling capacitors is 100 nF (176 nf to 265 nf for PCI Express®). Receiver data traces should be provided enough clearance to eliminate crosstalk from adjacent signals. If a receiver is not used, leave the associated pin pair unconnected or tie to ground. Trace match P/N to 10% (or less) of the UI. Traces should be routed to 50Ω single-ended/100Ω differential¹.
PS_MGTRTX0/ PS_MGTRTXN0 PS_MGTRTX1/ PS_MGTRTXN1 PS_MGTRTX2/ PS_MGTRTXN2 PS_MGTRTX3/ PS_MGTRTXN3	<ul style="list-style-type: none"> The transmitter should be AC coupled to the receiver. The recommended value for the AC coupling capacitors is 100 nF (176 nf to 265 nf for PCI Express®). Transmitter data traces should be provided enough clearance to eliminate crosstalk from adjacent signals. If a transmitter is not used, leave the associated pin pair unconnected or tie to ground. Trace match P/N to 10% (or less) of the UI. Traces should be routed to 50Ω single-ended/100Ω differential¹.
PS_MGTRREF	<ul style="list-style-type: none"> Connect to a 500Ω 0.5% resistor to this pin with the other terminal of the resistor connected to ground. Can be left unconnected or tied to ground if the transceivers will not be used.
PS_MGTRAVCC	<ul style="list-style-type: none"> The nominal voltage is 0.85 VDC. See the <i>Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)</i> for power supply voltage tolerances. Not sharing this supply with other non-transceiver loads reduces the chance of exceeding the noise requirement for the power supply. The following filter is recommended: <ul style="list-style-type: none"> One of 10 μF 10% ceramic For optimal performance, power supply noise must be less than 10 mV_{pp}. If the noise on the power supply is less than or equal to this requirement, the performance of the transceiver will not be negatively impacted. If all of the PS-GTR transceivers in the Quad are not used, the associated power pins can be left unconnected or tied to ground. <p>Note: PS MGTs must be powered to perform zeroization of the FPD.</p> <ul style="list-style-type: none"> For power consumption, refer to the Power Design Manager (PDM) tool (download at www.amd.com/power).
PS_MGTRAVTT	<ul style="list-style-type: none"> The nominal voltage is 1.8 VDC. See <i>Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)</i> for power supply voltage tolerances. Not sharing this supply with other non-transceiver loads reduces the chance of exceeding the noise requirement for the power supply. The following filter is recommended: <ul style="list-style-type: none"> One of 10 μF 10% ceramic: If all of the PS-GTR transceivers in the Quad are not used, the associated power pins can be left unconnected or tied to ground. <p>Note: PS MGTs must be powered to perform zeroization of the FPD.</p> <ul style="list-style-type: none"> For power consumption, refer to the Power Design Manager (PDM) tool (download at www.amd.com/power).


Notes:

- The [USB 3.0/SuperSpeed specification](#) calls for differential trace routing of 90Ω ±15Ω .

PCB Routing Guidelines for MIPI D-PHY

MIPI D-PHY is intended for use in mobile devices including cameras, displays, and unified protocol interfaces. This standard is only supported in the HP I/O banks of the AMD Kintex™ UltraScale+™ and AMD Virtex™ UltraScale+™ FPGAs, and AMD Zynq™ UltraScale+™ MPSoCs. Support for this standard in UltraScale devices is in adherence to the MIPI alliance interface specifications. Typical MIPI D-PHY cores consist of four differential data lanes and one differential clock lane, as shown in the following figure.

Figure 154: MIPI D-PHY Transmitter and Receiver Paths



X22696-042219

MIPI D-PHY PCB guidelines for impedance, trace length, spacing, skew, and termination are shown in the previous figure.

Table 132: MIPI D-PHY PCB Trace and Skew Guidelines

Parameter	Specification
Trace type	Stripline
Differential impedance	95±10% Ω
Trace length	≤ 10.0 inches ^{1,2}
Spacing in pair	2X trace width
Spacing to other MIPI pairs	5X trace width
Spacing to other non-MIPI pairs	6X trace width
Skew between P/N per pair	±2 ps
Skew between clock and data	±8 ps
Skew between data and data	±8 ps
Termination at receiver	100 Ω differential ³

Notes:

1. For distances greater than 4 inches, receiver equalization must be used (CTLE). It is strongly recommended that the PCB (data-to-data and data-to-clock) be deskewed even if the deskew feature in the IP is used.
2. Operation at 2.5 Gb/s is supported as long as the channel loss is kept under 9 dB.
3. When FPGA/MPSoC is the receiver, 100 Ω on-chip termination is recommended via the DIFF_TERM attribute. See the MIPI_DPHY section in *UltraScale Architecture SelectIO Resources User Guide* ([UG571](#)).

Migration between UltraScale Devices and Packages

When migrating a design from one AMD UltraScale™ device to another, special attention must be paid to potential differences between the two devices. Refer to the checklist below, and the subsequent sections that explain each item in the checklist.

UltraScale Device Migration Checklist

1. Review footprint compatibility, including banks that are bonded vs. unbonded.
2. Review package dimensions, including height and overhang.
3. Review I/O and transceiver Quad numbering and locations.
4. Review HPIO to HRIO and HRIO to HPIO migration.
5. Review GTH to GTY and GTY to GTH transceiver migration.
6. Review the super logic region (SLR) boundaries.
7. Review monolithic to stacked silicon migration as it pertains to configuration banks.
8. Review memory interface locations to maintain bank grouping.
9. Review DCI cascade and internal V_{REF} requirements.
10. Review system monitor requirements and pin differences.
11. Review decoupling capacitor requirements.
12. Review PCI Express® requirements and block locations.
13. Review integrated 100G Ethernet requirements and block locations.
14. Review Interlaken requirements and block locations.
15. Review power supply and thermal requirements.
16. Review pin flight time differences between devices.

1. Footprint Compatibility between Packages

Any two packages with the same footprint identifier code are footprint compatible. The footprint identifier code consists of a package designator code and BGA pin count information (such as A2104, B2104, or C2104). A customer that wants pin compatibility across any two devices in a given package should select a footprint compatible package such that they can be designed to be electrically compatible. If migrating from a bigger device to a smaller device within the same package, some I/O and transceiver banks and their pins might be unbonded or not exist on the smaller device. Refer to the Footprint Compatibility, I/O Bank Migration, and Transceiver Quad Migration tables in *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)* to determine which UltraScale devices are footprint compatible, as well as the extent to which the various banks might become unbonded or not exist.

Examples

The following figure shows a snippet of the Footprint Compatibility table from *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)*. The XCKU035 in the A1156 package can only be compatible with the XCKU040 and XCKU060 in the A1156 package, but not with any of the devices in the A1517. A design created with the XCKU035 in the A1156 package, for example, can be footprint compatible with any bigger A1156 device.

Note: For purposes of understanding footprint compatibility and migration, only the last letter in the package name is relevant. For example, the FFVA2104 and FLVA2104 are essentially the same as far as footprint compatibility and migration due to the last letter "A" being the same.

Figure 155: Portion of Footprint Compatibility Table

A1156	XCKU035	XCKU040	XCKU060
A1517	XCKU060	XCKU085	XCKU115

UG583_c3_14_022515

If migrating from a bigger device to a smaller device within the same package, there exists the potential for some I/O and transceiver banks to not be bonded out or present on the smaller device. The following figure shows a snippet of the Transceiver Quad Migration table from *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)*. Transceiver banks 120, 121, 122, 220, 221, and 222 are present on the XCVU160 and XCVU190 C2104 devices, but those transceiver banks are not available on the XCVU080, XCVU095, and XCVU125 C2104 devices.

Figure 156: Transceiver Quad Availability within C2104 Packages

Power Supply Group:	LS	LS	LS	LC	LC	LC	LC	LC	LN	LN	LN	LN	LN	RS	RS	RS	RC	RC	RC	RC	RC	RN	RN	RN	RN	RN	
FFVC2104	XCVU080			124	125	126	127	128	129	130	131						224	225	226	227	228	229	230	231			
	XCVU095								RCAL									RCAL									
FLVC2104	XCVU125			124	125	126	127	128	129	130	131	132	133				224	225	226	227	228	229	230	231	232	233	
					RCAL					RCAL								RCAL					RCAL				
FLGC2104	XCVU160	120	121	122	124	125	126	127	128	129	130	131	132	133	220	221	222	224	225	226	227	228	229	230	231	232	233
	XCVU190	RCAL				RCAL				RCAL					RCAL			RCAL					RCAL				

UG583_c9_15_022515


2. Package Dimensions

Package dimensions for UltraScale devices can vary from one package to another, so care must be taken in system design to ensure that changes in length, width, or height do not interfere with other components in the system. Some packages are also “overhang” packages in that their lengths and widths extend further beyond the pin array than other devices with similar pin counts. If migrating from a smaller package to a bigger package, ensure that an appropriate keep-out area is in place so that no capacitors or other components interfere with the bigger outline. Refer to the Mechanical Drawings chapter in *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)* for the dimensions of the various UltraScale device packages.

Example


According to the Mechanical Drawings chapter in *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)*, the FFVA1156 package used for the XCKU035 and XCKU040 (see the following figure) has a smaller height than the FFVA1156 package used for the XCKU060 (see [Figure 158: Portion of Package Drawing for FFVA1156 \(XCKU060\)](#)).

Figure 157: Portion of Package Drawing for FFVA1156 (XCKU035 and XCKU040)



UG583_c3_16_022515

Figure 158: Portion of Package Drawing for FFVA1156 (XCKU060)



UG583_c3_17_022515

3. I/O Bank and Transceiver Quad Numbers

Certain UltraScale devices that are footprint compatible in a package might have different I/O bank and transceiver Quad numbers associated with the same package pins. Depending on where these banks are physically located on the die, certain multi-bank interfaces (such as memory or transceiver) could be impacted in regards to migrating from one device to the other. If a bank number change also involves the bank moving from one column on the die to another, a multi-bank interface that requires both banks to be in the same column could be impacted.

Refer to the I/O Bank Migration and Transceiver Quad Migration tables in *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)* for the bank and transceiver Quad numbering differences. To determine bank locations on a die, refer to the Die Level Bank Numbering Overview section of *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)*.

The example below shows a successful migration path despite both bank number changes and banks being in different columns between the two different devices.

Example: Planning Ahead for Bank Number Changes

As shown in the following figure, banks 69, 70, and 71 in the XCVU095-FFVD1517 connect to the same pins as banks 71, 72, and 73 in the XCVU125-FLVD1517. The next figure confirms that the multi-bank interface remains intact despite the differing locations on each die. All three banks remain contiguous on each respective die. This would result in a successful migration for interfaces that use these three banks, with pin flight times being the only potential differences.

Figure 159: Portion of I/O Bank Migration Table

FFVD1517	XCVU080 XCVU095	Right								
		84	94	65	66	67	69	70	71	
FLVD1517	XCVU125	Right-South						Right-North		
		84	94	65	66	67	71	72	73	
	XCKU115	Right-South						Right-North		
		84	94	65	66	67	71	72	73	

UG583_c3_18_022515

Figure 160: XCVU095 (L) and XCVU125 (R) Bank Locations

GTU Quad 131 X0Y28–X0Y31	CMAC X0Y3	Bank 51 HP I/O	PL10(415) CMT MMCM071	Bank 71 HP I/O	PL10(91) CMT MMCM053	PCIe X0Y3	GTH Quad 231 X1Y28–X1Y31	GTU Quad 133 X0Y36–X0Y39	CMAC X0Y5	Bank 53 HP I/O	PL10(635) CMT MMCM059	Bank 73 HP I/O	PL10(59) CMT MMCM059	PCIe X0Y3	GTH Quad 233 X1Y36–X1Y39
GTU Quad 130 X0Y24–X0Y27	ILKN X0Y4	Bank 50 HP I/O	PL10(213) CMT MMCM051	Bank 70 HP I/O	PL10(29) CMT MMCM045	ILKN X1Y4	GTH Quad 230 X1Y24–X1Y27	GTU Quad 132 X0Y32–X0Y35	CMAC X0Y4	Bank 52 HP I/O	PL10(617) CMT MMCM058	Bank 72 HP I/O	PL10(571) CMT MMCM058	ILKN X1Y5	GTH Quad 232 X1Y32–X1Y35
GTU Quad 129 X0Y20–X0Y23	CMAC X0Y2	Bank 49 HP I/O	PL10(111) CMT MMCM051	Bank 69 HP I/O	PL10(277) CMT MMCM058	PCIe X0Y2	GTH Quad 229 X1Y20–X1Y23	GTU Quad 131 X0Y28–X0Y31	ILKN X0Y4	Bank 51 HP I/O	PL10(433) CMT MMCM057	Bank 71 HP I/O	PL10(531) CMT MMCM057	SYSTEM CFG	GTH Quad 231 X1Y28–X1Y31
GTU Quad 128 X0Y16–X0Y19	ILKN X0Y3	Bank 48 HP I/O	PL10(895) CMT MMCM054	Bank 68 HP I/O	PL10(423) CMT MMCM052	PCIe X0Y1	GTH Quad 228 X1Y16–X1Y19	GTU Quad 130 X0Y24–X0Y27	CMAC X0Y3	Bank 50 HP I/O	PL10(433) CMT MMCM056	Bank 70 HP I/O	PL10(231) CMT MMCM057	CFG	GTH Quad 230 X1Y24–X1Y27
GTU Quad 127 X0Y12–X0Y15	CMAC X0Y1	Bank 47 HP I/O	PL10(597) CMT MMCM033	Bank 67 HP I/O	PL10(223) CMT MMCM031	ILKN X1Y2	GTH Quad 227 X1Y12–X1Y15	GTU Quad 129 X0Y20–X0Y23	ILKN X0Y3	Bank 49 HP I/O	PL10(131) CMT MMCM053	Bank 69 HR I/O	PL10(231) CMT MMCM057	PCIe X0Y2	GTH Quad 229 X1Y20–X1Y23
GTU Quad 126 X0Y8–X0Y11	ILKN X0Y1	Bank 46 HP I/O	PL10(495) CMT MMCM032	Bank 66 HP I/O	PL10(203) CMT MMCM035	SYSTEM CFG	GTH Quad 226 X1Y8–X1Y11	SLR Crossing							
GTU Quad 125 X0Y4–X0Y7	CMAC X0Y0	Bank 45 HP I/O	PL10(293) CMT MMCM031	Bank 65 HP I/O	PL10(183) CMT MMCM059	CFG	GTH Quad 225 X1Y4–X1Y7	GTU Quad 128 X0Y16–X0Y19	CMAC X0Y2	Bank 48 HP I/O	PL10(895) CMT MMCM054	Bank 68 HP I/O	PL10(29) CMT MMCM054	PCIe X0Y1	GTH Quad 228 X1Y16–X1Y19
GTU Quad 124 X0Y0–X0Y3	ILKN X0Y0	Bank 44 HP I/O	PL10(91) CMT MMCM030	Bank 84/94 HR I/O	PL10(161) CMT MMCM058	PCIe X0Y0 (tandem)	GTH Quad 224 X1Y0–X1Y3	GTU Quad 127 X0Y12–X0Y15	CMAC X0Y1	Bank 47 HP I/O	PL10(617) CMT MMCM033	Bank 67 HP I/O	PL10(271) CMT MMCM054	ILKN X1Y2	GTH Quad 227 X1Y12–X1Y15
								GTU Quad 126 X0Y8–X0Y11	ILKN X0Y1	Bank 46 HP I/O	PL10(495) CMT MMCM032	Bank 66 HP I/O	PL10(293) CMT MMCM031	SYSTEM CFG	GTH Quad 226 X1Y8–X1Y11
								GTU Quad 125 X0Y4–X0Y7	CMAC X0Y0	Bank 45 HP I/O	PL10(293) CMT MMCM031	Bank 65 HP I/O	PL10(183) CMT MMCM059	CFG	GTH Quad 225 X1Y4–X1Y7
								GTU Quad 124 X0Y0–X0Y3	ILKN X0Y0	Bank 44 HP I/O	PL10(91) CMT MMCM030	Bank 84/94 HR I/O	PL10(231) CMT MMCM057	PCIe X0Y0	GTH Quad 224 X1Y0–X1Y3

UG583_c3_19_022515

4. HP/HR Migration

In a small number of cases, an HR bank in one device might become an HP bank in an otherwise footprint-compatible package, and vice-versa. If migration is desired at the initial design stage, it is advisable to use only 1.8V I/O standards so that any move from HR to HP banks is not affected. If this is not possible, avoid HR banks that can become HP banks. Refer to the I/O Bank Migration table in *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)* to determine which banks could be affected, or refer to the following table.

Table 133: Interchanged HP and HR Banks

Package	Bank(s)	HP	HR
A2104	70	XCVU125	XCKU115

In addition, there are functional differences between HP and HR I/Os, such as whether DCI is available (HP only) and I/O standard availability. Refer to the Supported Features in the HR and HP I/O Banks table in *UltraScale Architecture SelectIO Resources User Guide (UG571)* to determine the capabilities and I/O standards available with each type of I/O.

When it is not possible to avoid migrating from an HR bank to an HP bank, the new HP I/Os might need to be level translated due to their maximum 1.8V output. Refer to *Interfacing 7 Series FPGAs High-Performance I/O Banks with 2.5V and 3.3V I/O Standards (XAPP520)*, which gives many methodologies for interfacing HP I/Os with higher voltage I/Os.

5. GTH/GTY Transceiver Migration

GTH and GTY transceivers are designed to support pin migration and have identical differential signal and reference clock bond-out. GTY transceivers offer higher performance than GTH transceivers, and thus have different operating characteristics. For example, GTY transceivers are capable of supporting line rates up to 30.5 Gb/s, while GTH transceivers can support up to 16.375 Gb/s. If moving to/from GTH/GTY transceivers, these differences and any others must be accounted for. Refer to *Virtex UltraScale FPGAs Data Sheet: DC and AC Switching Characteristics (DS893)*, *UltraScale Architecture GTH Transceivers User Guide (UG576)*, and *UltraScale Architecture GTY Transceivers User Guide (UG578)* to learn more about the capabilities of each type of transceiver.

To determine if a migration path will result in any swapping between GTH and GTY transceivers (or vice versa), refer to the Transceiver Quad Migration table in *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)*.

Example

The following figure shows a snippet of the Transceiver Quad Migration table from *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)*. Transceiver banks 126, 127, 128, 131, 132, and 133 are shown as GTH transceivers in the XCKU115 A2104 device, while the corresponding transceiver banks are GTY transceivers in the XCVU080, XCVU095, and XCVU125 A2104 devices. Note how the transceiver numbers are different, indicating that these are different transceiver Quads that bond to the same pins (see [3. I/O Bank and Transceiver Quad Numbers](#)).

Figure 161: Snippet of Transceiver Quad Migration Table (GTY in bold)

Power Supply Group:		LS	LS	LS	LN	LN	LN	RS	RS	RS	RS	RN	RN	RN
FFVA2104	XCVU080	125	126	127	128	129	130	224	225	226	227	228	229	230
	XCVU095					RCAL			RCAL					
FLVA2104	XCVU125	125	126	127	130	131	132	224	225	226	227	231	232	233
	XCKU115	126	127	128	131	132	133	224	225	226	227	231	232	233
						RCAL			RCAL			RCAL		

UG583_c3_20_022515

MGTRREF and MGTAVTTRCAL

Some MGTRREF and MGTAVTTRCAL pins might have slightly different labels in different devices within the same package footprint. For example, pin A39 in the B1760 package is either labeled MGTAVTTRCAL_L or MGTAVTTRCAL_LN depending on the device. The functionality of the pin does not change, but the change in label should be noted.

6. SLR Migration

Some UltraScale devices are implemented with stacked silicon interconnect (SSI) technology, and certain implementations can only be contained within one SLR. SLRs are divided by a passive silicon interposer. Migrations that span multiple SLRs are not recommended. For example, DDR4 interfaces running at the maximum data rate must be contained within one SLR. See [8. Memory Interface Migration](#) for a specific example of where the interposer can break an existing implementation.

The bank drawings in *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)* identify SLRs connected by interposers.

Example

The following figure shows three SLRs for the XCVU190 device, each separated from the next by a silicon interposer.

Figure 162: Bank Drawing for the XCVU190 Showing SLRs

GTY Quad 133 X0Y56–X0Y59	CMAC X0Y8	Bank 53 HP I/O	PLL[28:29] CMT MMCM[14]	Bank 73 HP I/O	PLL[58:59] CMT MMCM[29]	PCIe X0Y5	GTH Quad 233 X1Y56–X1Y59
GTY Quad 132 X0Y52–X0Y55	CMAC X0Y7	Bank 52 HP I/O	PLL[26:27] CMT MMCM[13]	Bank 72 HP I/O	PLL[56:57] CMT MMCM[28]	ILKN X1Y8	GTH Quad 232 X1Y52–X1Y55
GTY Quad 131 X0Y48–X0Y51	ILKN X0Y7	Bank 51 HP I/O	PLL[24:25] CMT MMCM[12]	Bank 71 HP I/O	PLL[54:55] CMT MMCM[27]	SYSMON CFG	GTH Quad 231 X1Y48–X1Y51
GTY Quad 130 X0Y44–X0Y47	CMAC X0Y6	Bank 50 HP I/O	PLL[22:23] CMT MMCM[11]	Bank 70 HP I/O	PLL[52:53] CMT MMCM[26]	CFG	GTH Quad 230 X1Y44–X1Y47
GTY Quad 129 X0Y40–X0Y43	ILKN X0Y6	Bank 49 HP I/O	PLL[20:21] CMT MMCM[10]	Bank 69 HP I/O	PLL[50:51] CMT MMCM[25]	PCIe X0Y4	GTH Quad 229 X1Y40–X1Y43
SLR Crossing							
GTY Quad 128 X0Y36–X0Y39	CMAC X0Y5	Bank 48 HP I/O	PLL[18:19] CMT MMCM[09]	Bank 68 HP I/O	PLL[48:49] CMT MMCM[24]	PCIe X0Y3	GTH Quad 228 X1Y36–X1Y39
GTY Quad 127 X0Y32–X0Y35	CMAC X0Y4	Bank 47 HP I/O	PLL[16:17] CMT MMCM[08]	Bank 67 HP I/O	PLL[46:47] CMT MMCM[23]	ILKN X1Y5	GTH Quad 227 X1Y32–X1Y35
GTY Quad 126 X0Y28–X0Y31	ILKN X0Y4	Bank 46 HP I/O	PLL[14:15] CMT MMCM[07]	Bank 66 HP I/O	PLL[44:45] CMT MMCM[22]	SYSMON CFG	GTH Quad 226 X1Y28–X1Y31
GTY Quad 125 X0Y24–X0Y27	CMAC X0Y3	Bank 45 HP I/O	PLL[12:13] CMT MMCM[06]	Bank 65 HP I/O	PLL[42:43] CMT MMCM[21]	CFG	GTH Quad 225 X1Y24–X1Y27
GTY Quad 124 X0Y20–X0Y23	ILKN X0Y3	Bank 44 HP I/O	PLL[10:11] CMT MMCM[05]	Bank 84/94 HR I/O	PLL[40:41] CMT MMCM[20]	PCIe X0Y2 (tandem)	GTH Quad 224 X1Y20–X1Y23
SLR Crossing							
GTY Quad 123 X0Y16–X0Y19	CMAC X0Y2	Bank 43 HP I/O	PLL[08:09] CMT MMCM[04]	Bank 63 HP I/O	PLL[38:39] CMT MMCM[19]	PCIe X0Y1	GTH Quad 223 X1Y16–X1Y19
GTY Quad 122 X0Y12–X0Y15	CMAC X0Y1	Bank 42 HP I/O	PLL[06:07] CMT MMCM[03]	Bank 62 HP I/O	PLL[36:37] CMT MMCM[18]	ILKN X1Y2	GTH Quad 222 X1Y12–X1Y15
GTY Quad 121 X0Y8–X0Y11	ILKN X0Y1	Bank 41 HP I/O	PLL[04:05] CMT MMCM[02]	Bank 61 HP I/O	PLL[34:35] CMT MMCM[17]	SYSMON CFG	GTH Quad 221 X1Y8–X1Y11
GTY Quad 120 X0Y4–X0Y7	CMAC X0Y0	Bank 40 HP I/O	PLL[02:03] CMT MMCM[01]	Bank 60 HP I/O	PLL[32:33] CMT MMCM[16]	CFG	GTH Quad 220 X1Y4–X1Y7
GTY Quad 119 X0Y0–X0Y3	ILKN X0Y0	Bank 39 HP I/O	PLL[00:01] CMT MMCM[00]	Bank 59 HP I/O	PLL[30:31] CMT MMCM[15]	PCIe X0Y0	GTH Quad 219 X1Y0–X1Y3

UG583_c3_21_022515

7. Monolithic to Stacked Silicon Migration

All UltraScale devices have pins dedicated to configuration functions contained in I/O bank 0. There are also I/O pins in bank 65 (multifunction configuration bank) known as multifunction or multipurpose pins that can be used for configuration, but then convert to normal I/O pins after configuration is complete. Additionally, in SSI devices, pins in bank 60 and bank 70 (as present) have restrictions during configuration similar to multifunction pins. However, pins in these banks do not have any configuration functions. See the Configuration Details chapter in *UltraScale Architecture Configuration User Guide (UG570)* to learn about how the behavior of these pins might affect your design.

Example: Migration from Monolithic XCKU060 to SSI XCKU115 in A1517 Package

The following figure shows the bank drawing for the monolithic XCKU060 device. The multifunction configuration bank (65) is circled. The next figure shows the bank drawing for the SSI XCKU115 device. The multifunction configuration bank (65) is circled along with bank 70. While bank 65 is the actual multifunction configuration bank, bank 70 shares the same restrictions during configuration.

Figure 163: Bank Drawing for XCKU060 Showing Multifunction Configuration Bank 65

GTH Quad 128 X0Y16–X0Y19	Bank 48 HP I/O	PLL[12:13] CMT MMCM[06]	Bank 68 HP I/O	PLL[22:23] CMT MMCM[11]	PCIe X0Y1	GTH Quad 228 X1Y16–X1Y19
GTH Quad 127 X0Y12–X0Y15	Bank 47 HP I/O	PLL[10:11] CMT MMCM[05]	Bank 67 HP I/O	PLL[20:21] CMT MMCM[10]	CMAC X0Y0	GTH Quad 227 X1Y12–X1Y15
GTH Quad 126 X0Y8–X0Y11	Bank 46 HP I/O	PLL[08:09] CMT MMCM[04]	Bank 66 HP I/O	PLL[18:19] CMT MMCM[09]	SYSMON CFG	GTH Quad 226 X1Y8–X1Y11
Bank 25 HP I/O	PLL[02:03] CMT MMCM[01]	Bank 45 HP I/O	Bank 65 HR I/O	PLL[16:17] CMT MMCM[08]	CFG	GTH Quad 225 X1Y4–X1Y7
Bank 24 HP I/O	PLL[00:01] CMT MMCM[00]	Bank 44 HP I/O	Bank 64 HR I/O	PLL[14:15] CMT MMCM[07]	PCIe X0Y0 (tandem)	GTH Quad 224 X1Y0–X1Y3

UG583_c3_22_022515

Figure 164: Bank Drawing for XCKU115 Showing Multifunction Configuration Bank 65, along with Bank 70

GTH Quad 133 X0Y36–X0Y39		Bank 53 HP I/O	PLL[26:27] CMT MMCM[13]	Bank 73 HP I/O	PLL[46:47] CMT MMCM[23]	PCIe X0Y3	GTH Quad 233 X1Y36–X1Y39
GTH Quad 132 X0Y32–X0Y35		Bank 52 HP I/O	PLL[24:25] CMT MMCM[12]	Bank 72 HP I/O	PLL[44:45] CMT MMCM[22]	CMAC X0Y1	GTH Quad 232 X1Y32–X1Y35
GTH Quad 131 X0Y28–X0Y31		Bank 51 HP I/O	PLL[22:23] CMT MMCM[11]	Bank 71 HP I/O	PLL[42:43] CMT MMCM[21]	SYSMON CFG	GTH Quad 231 X1Y28–X1Y31
Bank 30 HP I/O	PLL[06:07] CMT MMCM[03]	Bank 50 HP I/O	PLL[20:21] CMT MMCM[10]	Bank 70 HR I/O	PLL[40:41] CMT MMCM[20]	CFG	GTH Quad 230 X1Y24–X1Y27
Bank 29 HP I/O	PLL[04:05] CMT MMCM[02]	Bank 49 HP I/O	PLL[18:19] CMT MMCM[09]	Bank 69 HR I/O	PLL[38:39] CMT MMCM[19]	PCIe X0Y2	GTH Quad 229 X1Y20–X1Y23
SLR Crossing							
GTH Quad 128 X0Y16–X0Y19		Bank 48 HP I/O	PLL[16:17] CMT MMCM[08]	Bank 68 HP I/O	PLL[36:37] CMT MMCM[18]	PCIe X0Y1	GTH Quad 228 X1Y16–X1Y19
GTH Quad 127 X0Y12–X0Y15		Bank 47 HP I/O	PLL[14:15] CMT MMCM[07]	Bank 67 HP I/O	PLL[34:35] CMT MMCM[17]	CMAC X0Y0	GTH Quad 227 X1Y12–X1Y15
GTH Quad 126 X0Y8–X0Y11		Bank 46 HP I/O	PLL[12:13] CMT MMCM[06]	Bank 66 HP I/O	PLL[32:33] CMT MMCM[16]	SYSMON CFG	GTH Quad 226 X1Y8–X1Y11
Bank 25 HP I/O	PLL[02:03] CMT MMCM[01]	Bank 45 HP I/O	PLL[10:11] CMT MMCM[05]	Bank 65 HR I/O	PLL[30:31] CMT MMCM[15]	CFG	GTH Quad 225 X1Y4–X1Y7
Bank 24 HP I/O	PLL[00:01] CMT MMCM[00]	Bank 44 HP I/O	PLL[08:09] CMT MMCM[04]	Bank 64 HR I/O	PLL[28:29] CMT MMCM[14]	PCIe X0Y0 (tandem)	GTH Quad 224 X1Y0–X1Y3

UG583_c3_23_022515

8. Memory Interface Migration

When migrating a memory interface from one device to another, it is important to verify that all pinout and banking restrictions remain valid. For example, most memory interfaces require the use of two or three consecutive banks in one column. In addition, memory interfaces in UltraScale devices should not cross SLR boundaries. Refer to *UltraScale Architecture-Based FPGAs Memory IP LogiCORE IP Product Guide* (PG150) for a more comprehensive guide regarding the proper generation and verification of UltraScale memory interfaces.

Example

A memory interface in the XCVU095 that spans banks 47, 48, and 49 (see the following figure) would NOT migrate over to the XCVU125 because banks 48 and 49 are separated by an SLR in the XCVU125 (see the next figure). A better option for the memory interface would be banks 44, 45, and 46, because both banks remain contiguous on both devices.

Figure 165: Bank Drawing of XCVU095 Showing Consecutive Banks 47, 48, and 49

GTY Quad 131 X0Y28–X0Y31	CMAC X0Y3	Bank 51 HP I/O	PLL[14:15] CMT MMCM[07]	Bank 71 HP I/O	PLL[30:31] CMT MMCM[15]	PCle X0Y3	GTH Quad 231 X1Y28–X1Y31
GTY Quad 130 X0Y24–X0Y27	ILKN X0Y4	Bank 50 HP I/O	PLL[12:13] CMT MMCM[06]	Bank 70 HP I/O	PLL[28:29] CMT MMCM[14]	ILKN X1Y4	GTH Quad 230 X1Y24–X1Y27
GTY Quad 129 X0Y20–X0Y23	CMAC X0Y2	Bank 49 HP I/O	PLL[10:11] CMT MMCM[05]	Bank 69 HP I/O	PLL[26:27] CMT MMCM[13]	PCle X0Y2	GTH Quad 229 X1Y20–X1Y23
GTY Quad 128 X0Y16–X0Y19	ILKN X0Y3	Bank 48 HP I/O	PLL[08:09] CMT MMCM[04]	Bank 68 HP I/O	PLL[24:25] CMT MMCM[12]	PCle X0Y1	GTH Quad 228 X1Y16–X1Y19
GTY Quad 127 X0Y12–X0Y15	CMAC X0Y1	Bank 47 HP I/O	PLL[06:07] CMT MMCM[03]	Bank 67 HP I/O	PLL[22:23] CMT MMCM[11]	ILKN X1Y2	GTH Quad 227 X1Y12–X1Y15
GTY Quad 126 X0Y8–X0Y11	ILKN X0Y1	Bank 46 HP I/O	PLL[04:05] CMT MMCM[02]	Bank 66 HP I/O	PLL[20:21] CMT MMCM[10]	SYSMON CFG	GTH Quad 226 X1Y8–X1Y11
GTY Quad 125 X0Y4–X0Y7	CMAC X0Y0	Bank 45 HP I/O	PLL[02:03] CMT MMCM[01]	Bank 65 HP I/O	PLL[18:19] CMT MMCM[09]	CFG	GTH Quad 225 X1Y4–X1Y7
GTY Quad 124 X0Y0–X0Y3	ILKN X0Y0	Bank 44 HP I/O	PLL[00:01] CMT MMCM[00]	Bank 84/94 HR I/O	PLL[16:17] CMT MMCM[08]	PCle X0Y0 (tandem)	GTH Quad 224 X1Y0–X1Y3

UG583_c3_24_022515

Figure 166: Bank Drawing of XCVU125 Showing SLR between Banks 47, 48, and 49

GTY Quad 133 X0Y36–X0Y39	CMAC X0Y5	Bank 53 HP I/O	PLL[18:19] CMT MMCM[09]	Bank 73 HP I/O	PLL[38:39] CMT MMCM[19]	PCIe X0Y3	GTH Quad 233 X1Y36–X1Y39
GTY Quad 132 X0Y32–X0Y35	CMAC X0Y4	Bank 52 HP I/O	PLL[16:17] CMT MMCM[08]	Bank 72 HP I/O	PLL[36:37] CMT MMCM[18]	ILKN X1Y5	GTH Quad 232 X1Y32–X1Y35
GTY Quad 131 X0Y28–X0Y31	ILKN X0Y4	Bank 51 HP I/O	PLL[14:15] CMT MMCM[07]	Bank 71 HP I/O	PLL[34:35] CMT MMCM[17]	SYSMON CFG	GTH Quad 231 X1Y28–X1Y31
GTY Quad 130 X0Y24–X0Y27	CMAC X0Y3	Bank 50 HP I/O	PLL[12:13] CMT MMCM[06]	Bank 70 HP I/O	PLL[32:33] CMT MMCM[16]	CFG	GTH Quad 230 X1Y24–X1Y27
GTY Quad 129 X0Y20–X0Y23	ILKN X0Y3	Bank 49 HP I/O	PLL[10:11] CMT MMCM[05]	Bank 69 HR I/O	PLL[30:31] CMT MMCM[15]	PCIe X0Y2	GTH Quad 229 X1Y20–X1Y23
SLR Crossing							
GTY Quad 128 X0Y16–X0Y19	CMAC X0Y2	Bank 48 HP I/O	PLL[08:09] CMT MMCM[04]	Bank 68 HP I/O	PLL[28:29] CMT MMCM[14]	PCIe X0Y1	GTH Quad 228 X1Y16–X1Y19
GTY Quad 127 X0Y12–X0Y15	CMAC X0Y1	Bank 47 HP I/O	PLL[06:07] CMT MMCM[03]	Bank 67 HP I/O	PLL[26:27] CMT MMCM[13]	ILKN X1Y2	GTH Quad 227 X1Y12–X1Y15
GTY Quad 126 X0Y8–X0Y11	ILKN X0Y1	Bank 46 HP I/O	PLL[04:05] CMT MMCM[02]	Bank 66 HP I/O	PLL[24:25] CMT MMCM[12]	SYSMON CFG	GTH Quad 226 X1Y8–X1Y11
GTY Quad 125 X0Y4–X0Y7	CMAC X0Y0	Bank 45 HP I/O	PLL[02:03] CMT MMCM[01]	Bank 65 HP I/O	PLL[22:23] CMT MMCM[11]	CFG	GTH Quad 225 X1Y4–X1Y7
GTY Quad 124 X0Y0–X0Y3	ILKN X0Y0	Bank 44 HP I/O	PLL[00:01] CMT MMCM[00]	Bank 84/94 HR I/O	PLL[20:21] CMT MMCM[10]	PCIe X0Y0	GTH Quad 224 X1Y0–X1Y3

UG583_c3_25_022515

9. DCI Cascade and Internal V_{REF} Features

DCI Cascade and Internal V_{REF} are features that span multiple banks in an HPIO column. When migrating, if a bank (or banks) moves to another column or if it crosses an SLR/interposer boundary, it is no longer allowed as part of the original DCI cascade or internal V_{REF} group. Each SLR contains its own DCI and V_{REF} circuitry. See the DCI Cascading and Internal V_{REF} sections in *UltraScale Architecture SelectIO Resources User Guide (UG571)* to learn more about these features.

Example

The following figure highlights an HPIO column in the XCVU095. DCI cascade or internal V_{REF} can be used in this column among some or all of these banks. If migrating to the XCVU125 device, the interposer does not allow DCI cascade or internal V_{REF} to cross between banks 48 and 49, as shown in the next figure.

Figure 167: XCVU095 Showing HPIO Column

GTY Quad 131 X0Y28-X0Y31	CMAC X0Y3	Bank 51 HP I/O	PLL[14:15] CMT MMCM[07]	Bank 71 HP I/O	PLL[30:31] CMT MMCM[15]	PCIe X0Y3	GTH Quad 231 X1Y28-X1Y31
GTY Quad 130 X0Y24-X0Y27	ILKN X0Y4	Bank 50 HP I/O	PLL[12:13] CMT MMCM[06]	Bank 70 HP I/O	PLL[28:29] CMT MMCM[14]	ILKN X1Y4	GTH Quad 230 X1Y24-X1Y27
GTY Quad 129 X0Y20-X0Y23	CMAC X0Y2	Bank 49 HP I/O	PLL[10:11] CMT MMCM[05]	Bank 69 HP I/O	PLL[26:27] CMT MMCM[13]	PCIe X0Y2	GTH Quad 229 X1Y20-X1Y23
GTY Quad 128 X0Y16-X0Y19	ILKN X0Y3	Bank 48 HP I/O	PLL[08:09] CMT MMCM[04]	Bank 68 HP I/O	PLL[24:25] CMT MMCM[12]	PCIe X0Y1	GTH Quad 228 X1Y16-X1Y19
GTY Quad 127 X0Y12-X0Y15	CMAC X0Y1	Bank 47 HP I/O	PLL[06:07] CMT MMCM[03]	Bank 67 HP I/O	PLL[22:23] CMT MMCM[11]	ILKN X1Y2	GTH Quad 227 X1Y12-X1Y15
GTY Quad 126 X0Y8-X0Y11	ILKN X0Y1	Bank 46 HP I/O	PLL[04:05] CMT MMCM[02]	Bank 66 HP I/O	PLL[20:21] CMT MMCM[10]	SYSMON CFG	GTH Quad 226 X1Y8-X1Y11
GTY Quad 125 X0Y4-X0Y7	CMAC X0Y0	Bank 45 HP I/O	PLL[02:03] CMT MMCM[01]	Bank 65 HP I/O	PLL[18:19] CMT MMCM[09]	CFG	GTH Quad 225 X1Y4-X1Y7
GTY Quad 124 X0Y0-X0Y3	ILKN X0Y0	Bank 44 HP I/O	PLL[00:01] CMT MMCM[00]	Bank 84/94 HR I/O	PLL[16:17] CMT MMCM[08]	PCIe X0Y0 (tandem)	GTH Quad 224 X1Y0-X1Y3

UG583_c3_26_022515

Figure 168: XCVU125 Showing HPIO Column

GTY Quad 133 X0Y36–X0Y39	CMAC X0Y5	Bank 53 HP I/O	PLL[18:19] CMT MMCM[09]	Bank 73 HP I/O	PLL[38:39] CMT MMCM[19]	PCIe X0Y3	GTH Quad 233 X1Y36–X1Y39
GTY Quad 132 X0Y32–X0Y35	CMAC X0Y4	Bank 52 HP I/O	PLL[16:17] CMT MMCM[08]	Bank 72 HP I/O	PLL[36:37] CMT MMCM[18]	ILKN X1Y5	GTH Quad 232 X1Y32–X1Y35
GTY Quad 131 X0Y28–X0Y31	ILKN X0Y4	Bank 51 HP I/O	PLL[14:15] CMT MMCM[07]	Bank 71 HP I/O	PLL[34:35] CMT MMCM[17]	SYSMON CFG	GTH Quad 231 X1Y28–X1Y31
GTY Quad 130 X0Y24–X0Y27	CMAC X0Y3	Bank 50 HP I/O	PLL[12:13] CMT MMCM[06]	Bank 70 HP I/O	PLL[32:33] CMT MMCM[16]	CFG	GTH Quad 230 X1Y24–X1Y27
GTY Quad 129 X0Y20–X0Y23	ILKN X0Y3	Bank 49 HP I/O	PLL[10:11] CMT MMCM[05]	Bank 69 HR I/O	PLL[30:31] CMT MMCM[15]	PCIe X0Y2	GTH Quad 229 X1Y20–X1Y23
SLR Crossing							
GTY Quad 128 X0Y16–X0Y19	CMAC X0Y2	Bank 48 HP I/O	PLL[08:09] CMT MMCM[04]	Bank 68 HP I/O	PLL[28:29] CMT MMCM[14]	PCIe X0Y1	GTH Quad 228 X1Y16–X1Y19
GTY Quad 127 X0Y12–X0Y15	CMAC X0Y1	Bank 47 HP I/O	PLL[06:07] CMT MMCM[03]	Bank 67 HP I/O	PLL[26:27] CMT MMCM[13]	ILKN X1Y2	GTH Quad 227 X1Y12–X1Y15
GTY Quad 126 X0Y8–X0Y11	ILKN X0Y1	Bank 46 HP I/O	PLL[04:05] CMT MMCM[02]	Bank 66 HP I/O	PLL[24:25] CMT MMCM[12]	SYSMON CFG	GTH Quad 226 X1Y8–X1Y11
GTY Quad 125 X0Y4–X0Y7	CMAC X0Y0	Bank 45 HP I/O	PLL[02:03] CMT MMCM[01]	Bank 65 HP I/O	PLL[22:23] CMT MMCM[11]	CFG	GTH Quad 225 X1Y4–X1Y7
GTY Quad 124 X0Y0–X0Y3	ILKN X0Y0	Bank 44 HP I/O	PLL[00:01] CMT MMCM[00]	Bank 84/94 HR I/O	PLL[20:21] CMT MMCM[10]	PCIe X0Y0	GTH Quad 224 X1Y0–X1Y3

UG583_c3_27_022515

10. System Monitor

Each I/O bank provides access to up to sixteen of the system monitor differential auxiliary analog input pairs. Some banks do not contain full access, so choose banks carefully so that any desired system monitor analog input pairs are available for device migration.

Refer to the following table for a summary of UltraScale banks that do not contain full access to the sixteen differential auxiliary analog input pairs. If a particular device/bank is not listed, then full access should be assumed for that bank.

Table 134: Unavailable Analog Input Pairs

Package	Limited System Monitor Analog Access (Bank ¹)
D1517	XCKU115 (67) XCVU095 (67) XCVU125 (67)

Table 134: Unavailable Analog Input Pairs (cont'd)

Package	Limited System Monitor Analog Access (Bank ¹)
B1760	XCKU115 (53) XCVU095 (51) XCVU125 (53)
B2104	XCKU115 (68) XCVU095 (68) XCVU125 (68) XCVU190 (68)
A2577	XCVU190 (66 ² , 68)

Notes:

1. No access to pairs 0, 1, 2, 3, 8, 9, 10, 11 unless noted.
2. No access to any pairs.

In addition, analog auxiliary pair access can be found within each device's package file. Refer to the Package Files chapter in *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)* to obtain links to the UltraScale device package files. Analog auxiliary pins are in the form `_ADxP_` in the pin name.

Example

The following figure shows a portion of a package file. The pin name portions within the red boxes denote which analog auxiliary pair access those particular pins have access to.

Figure 169: Package File Showing Analog Pair Access

```
Device/Package xcku035fbva676 6/16/2014 16:20:56
```

Pin	Pin Name	Memory Byte Group	Bank	I/O Type	Super Logic Region	No-Connect
T11	DXN	NA	NA	NA	NA	NA
N12	VCCADC	NA	NA	NA	NA	NA
N11	GNDADC	NA	NA	NA	NA	NA
T12	DXP	NA	NA	NA	NA	NA
R12	VREFP	NA	NA	NA	NA	NA
P11	VREFN	NA	NA	NA	NA	NA
P12	VP	NA	NA	NA	NA	NA
R11	VN	NA	NA	NA	NA	NA
K7	M0_0	NA	0	CONFIG	NA	NA
H7	M1_0	NA	0	CONFIG	NA	NA
G7	INIT_B_0	NA	0	CONFIG	NA	NA
L7	M2_0	NA	0	CONFIG	NA	NA
Y7	CFGBVS_0	NA	0	CONFIG	NA	NA
J8	PUDC_B_0	NA	0	CONFIG	NA	NA
F7	POR_OVERRIDE	NA	NA	NA	NA	NA
V7	DONE_0	NA	0	CONFIG	NA	NA
M7	PROGRAM_B_0	NA	0	CONFIG	NA	NA
L9	TDO_0	NA	0	CONFIG	NA	NA
N9	TDI_0	NA	0	CONFIG	NA	NA
N7	RDWR_FCS_B_0	NA	0	CONFIG	NA	NA
R7	D02_0	NA	0	CONFIG	NA	NA
U7	D00_MOSI_0	NA	0	CONFIG	NA	NA
P7	D03_0	NA	0	CONFIG	NA	NA
T7	D01_DIN_0	NA	0	CONFIG	NA	NA
P9	TMS_0	NA	0	CONFIG	NA	NA
U9	CCLK_0	NA	0	CONFIG	NA	NA
R9	TCK_0	NA	0	CONFIG	NA	NA
AA7	VBATT	NA	NA	NA	NA	NA
U20	IO_L24P_T3U_N10_44	3U	44	HP	NA	NA
U21	IO_L24N_T3U_N11_44	3U	44	HP	NA	NA
V21	IO_T3U_N12_44	3U	44	HP	NA	NA
P20	IO_L23P_T3U_N8_44	3U	44	HP	NA	NA
R20	IO_L23N_T3U_N9_44	3U	44	HP	NA	NA
T19	IO_L22P_T3U_N6_DBC_AD0P_44	3U	44	HP	NA	NA
T20	IO_L22N_T3U_N7_DBC_AD0N_44	3U	44	HP	NA	NA
P18	IO_L21P_T3L_N4_AD8P_44	3L	44	HP	NA	NA
P19	IO_L21N_T3L_N5_AD8N_44	3L	44	HP	NA	NA
P21	IO_L20P_T3L_N2_AD1P_44	3L	44	HP	NA	NA
R21	IO_L20N_T3L_N3_AD1N_44	3L	44	HP	NA	NA

UG583_c3_28_022515

11. Decoupling Capacitors

If the design is not expected to change when migrating to the new device, the same decoupling scheme can be used. If adding logic, AMD recommends using a decoupling scheme appropriate for the new device and design. Refer to [Chapter 1: Power Distribution System in UltraScale Devices](#) for decoupling capacitor guidelines for the various UltraScale devices.

Example

The following figure shows a portion of the Kintex UltraScale decoupling recommendations. If beginning a design in the XCKU040-FFVA1156 and migrating to the XCKU060-FFVA1156, use the XCKU040 capacitor scheme if no design changes are anticipated in the migration. If more logic is added during the migration to the XCKU060, AMD recommends using the decoupling scheme for that device.

Figure 170: Portion of Kintex Decoupling Capacitor Recommendations

Table 1-1: Kintex UltraScale Devices Power Supply Decoupling Capacitors

	V_{CCINT}/V_{CCINT_IO} ⁽¹⁾			V_{CCBRAM}		V_{CCAUX}/V_{CCAUX_IO} ⁽²⁾		$HRIO$ ⁽³⁾	$HPIO$ ⁽³⁾
	680 μF ⁽⁴⁾	100 μF	4.7 μF	47 μF	4.7 μF	47 μF	4.7 μF	47 μF	47 μF
XCKU035-FBVA676	1	1	3	1	1	1	3	1	1
XCKU035-FBVA900	1	1	3	1	1	2	4	1	1
XCKU035-FFVA1156	1	1	3	1	1	2	5	1	1
XCKU040-FBVA676	1	2	3	1	1	1	3	1	1
XCKU040-FBVA900	1	2	3	1	1	2	4	1	1
XCKU040-FFVA1156	1	2	3	1	1	2	5	1	1
XCKU060-FFVA1156	1	2	4	1	1	2	5	1	1
XCKU060-FFVA1517	1	2	4	1	1	3	6	1	1

UG583_c3_28_0226
15

12. PCI Express Migration

PCI Express lane assignments might differ between devices. In addition, the block that is used as tandem boot might also be different from device to device. Refer to the Core Pinouts appendix in *UltraScale Devices Gen3 Integrated Block for PCI Express LogiCORE IP Product Guide* (PG156) for UltraScale device core pinouts.

13. Integrated 100G Ethernet Migration

The 100G Ethernet IP core is a dedicated block that provides 100 Gb/s MAC and PCS logic capability. The block consists of a number of dedicated locations that are located at different points on different devices. There are specific rules for connecting these dedicated blocks to specific transceiver locations, which might impact migration to another device. In addition, the CAUI-4 and switchable CAUI-10/CAUI-4 modes require GTY transceivers.

Refer to *UltraScale Devices Integrated 100G Ethernet LogiCORE IP Product Guide (PG165)* for a more comprehensive guide regarding designing with the 100G Ethernet core, with special attention to the Transceiver Selection Rules section.

14. Interlaken Migration

The AMD UltraScale architecture integrated IP core for Interlaken is a highly configurable integrated IP core that can support an overall bandwidth up to 150 Gb/s for protocol logic transmission. The core connects to serial transceivers at defined rates up to 12.5 Gb/s with GTH transceivers, and up to 25.78125 Gb/s with GTY transceivers. The block consists of a number of dedicated locations that are located at different points on different devices. There are specific rules for connecting these dedicated blocks to specific transceiver locations, which might impact migration to another device. Refer to *Integrated Interlaken up to 150G LogiCORE IP Product Guide (PG169)* for a more comprehensive guide regarding the proper generation and verification of Interlaken interfaces in UltraScale devices with special attention to the Transceiver Interface section.

15. Power Supplies and Thermal Considerations

Migration across devices, especially from smaller to bigger devices, can result in higher power consumption requiring improved thermal management and power delivery. Consult the Power Design Manager (PDM) tool (download at www.amd.com/power) and select the largest device being considered to estimate power consumption and thermal requirements.



RECOMMENDED: AMD recommends use of the Delphi thermal model during thermal modeling of a package. AMD does not recommend using a two-resistor model for thermal simulation and design due to lack of precision and accuracy. If more accuracy is desired than the Delphi model, AMD can provide a detailed model representation of the package by request. However, this might consume more simulation memory and runtime in its use. The user of the thermal model needs to include a consideration of thermal sensor accuracy, thermal interface material parameters, and manufacture variation on the thermal solution. Also, other examples of manufacture variations include the tolerance in airflow from a fan, the tolerance on performance of the heat pipe and vapor chamber, and the manufacture variation of the attachment of fins to the heat-sink base and the flatness of the surface.

16. Pin Flight Times across Packages

Even if two devices are footprint compatible from a design and PCB standpoint, FPGA package flight times will be different across different devices, so they will need to be accounted for. Flight time information can be found within the AMD Vivado™ tools in the Package Pins tab during the I/O planning stage or after synthesis. The ideal strategy to account for pin flight time differences is to deskew the printed circuit board when migrating to the new device. If this is not

Migration between UltraScale and UltraScale+ FPGAs

When migrating a design from one AMD UltraScale™ FPGA to another, special attention must be paid to potential differences between the two FPGAs. [Chapter 7: Migration between UltraScale Devices and Packages](#) contains a comprehensive guide towards general migration concerns within the UltraScale families. This section addresses additional concerns specific to migration from UltraScale to UltraScale+ FPGAs, as well as from UltraScale+ to UltraScale+ FPGAs.

The underlying architecture of UltraScale+ FPGAs is the same as that of UltraScale FPGAs, meaning that most IP that contains CLBs, routing resources, DSP slices, and block RAM is the same. *UltraScale Architecture and Product Data Sheet: Overview (DS890)* contains a complete listing of the available resources found in UltraScale and UltraScale+ FPGAs.

The following checklist and sections contain the necessary items to review for optimal UltraScale FPGAs migration.

UltraScale+ FPGA Migration Checklist

Review [Chapter 7: Migration between UltraScale Devices and Packages](#) for general migration concerns within the UltraScale family.

1. Review the package migration table from UltraScale to UltraScale+ FPGAs.
2. Review power supply voltage differences between UltraScale and UltraScale+ FPGAs.
3. Review VCCINT_IO connection requirements.
4. Review I/O differences between UltraScale to UltraScale+ FPGAs.
5. Review transceiver differences between UltraScale to UltraScale+ FPGAs.
6. Review configuration details that pertain to UltraScale+ FPGAs.
7. Review memory interface PCB routing guidelines.
8. Review new SMBAlert signal for UltraScale+ FPGAs.
9. Review block RAM/UltraRAM capacities between UltraScale and UltraScale+ FPGAs.
10. Review ESD requirements pertaining to UltraScale+ FPGAs.

11. Review PCI Express differences between UltraScale+ FPGAs.
12. Review migration examples between UltraScale and UltraScale+ FPGAs.
13. Review migration scenarios to better understand the migration process.

1. Package Migration from UltraScale to UltraScale+ FPGAs

The packages listed in the following table support UltraScale to UltraScale+ FPGA migration.

Table 135: UltraScale to UltraScale+ FPGA Migration Paths by Package

Package	UltraScale FPGA	UltraScale+ FPGA
A676	KU035 KU040	KU3P ¹ KU5P ¹
A1156	KU025 KU035 KU040 KU060 KU095	KU11P ² KU15P ²
C1517	KU095 ³ VU065 ³ VU080 ³ VU095 ³	VU3P ³
F1924	KU085 ³ KU115	VU11P ⁴
A2104	KU115 VU080 ³ VU095 ³ VU125	VU5P VU7P VU9P VU13P
B2104	KU095 ³ KU115 VU080 ³ VU095 ³ VU125 VU160 VU190	VU5P VU7P VU9P VU11P VU13P
C2104	VU095 ³ VU125 VU160 VU190	VU5P VU7P VU9P VU11P VU13P

Table 135: UltraScale to UltraScale+ FPGA Migration Paths by Package (cont'd)

Package	UltraScale FPGA	UltraScale+ FPGA
A2577	VU190	VU9P VU11P VU13P

Notes:

1. KU3P and KU5P GTY transceivers up to 16.3 Gb/s in A676 package.
2. KU11P and KU15P GTY transceivers up to 16.3 Gb/s in A1156 package.
3. Monolithic device, otherwise SSI device.
4. VU11P GTY transceivers up to 16.3 Gb/s in F1924 package.

2. Voltage Differences between UltraScale to UltraScale+ FPGAs

Pay special attention to the voltage differences between V_{CCINT} , V_{CCINT_IO} , V_{CCBRAM} , and $AVCC$ between UltraScale and UltraScale+ FPGAs (see the following table).

Table 136: Voltages Differences between UltraScale and UltraScale+ FPGAs

Voltage	UltraScale FPGA (V)	UltraScale+ FPGA (V)
AVCC	1.0	0.9
AVTT	1.2	1.2
AVCCPLL	1.8	1.8
VCCINT (-3)	1.0	0.90
VCCINT (-2, -1)	0.95	0.85
VCCINT (-2L)	N/A	0.72 or 0.85
VCCINT (-1L)	0.90	0.72 or 0.85
VCCINT_IO (-3)	1.0	0.90
VCCINT_IO (-2, -1)	0.95	0.85
VCCINT_IO (-2L)	N/A	0.85
VCCINT_IO (-1L)	0.90	0.85
VCCBRAM (-3)	1.0	0.90
VCCBRAM (-2, -1)	0.95	0.85
VCCBRAM (-2L)	N/A	0.85
VCCBRAM (-1L)	0.95	0.85
VCCAUX	1.80	1.80
VCCAUX_IO	1.80	1.80

3. Power Supply Voltage Levels and VCCINT_IO Connection

For UltraScale and UltraScale+ FPGAs in the -3, -2, or -1 speed grades, it is required to connect VCCINT_IO to VCCINT and VCCBRAM. However, for UltraScale+ -2L and -1L FPGAs, it is required to connect VCCINT_IO to VCCBRAM.


To support this requirement, the system shown in the following table is recommended with the following figure. The VCCBRAM plane must be sized appropriately to support the additional load for VCCINT_IO when using -2L or -1L devices.

It is possible to connect VCCINT, VCCINT_IO, and VCCBRAM together when using -2 or -1 devices because the respective voltages are the same. It is also possible to connect VCCINT, VCCINT_IO, and VCCBRAM together when using -2L or -1L devices if VCCINT is always operated at 0.85V.

Table 137: VCCINT/VCCINT_IO/VCCBRAM Connection Matrix

	-3, -2, -1	-2L, -1L
VCCINT_IO connection	VCCINT/VCCBRAM	VCCBRAM
VCCINT sense line connection	Averaged with VCCINT_IO	Single connection to VCCINT
VCCBRAM sense line connection	Single connection to VCCBRAM	Averaged with VCCINT_IO
R1	0Ω 0603	0Ω 0603
R2	0Ω 0603	Do not populate
R3	10Ω 0603	10Ω 0603
R4	10Ω 0603	Do not populate
R5	Do not populate	10Ω 0603
R6	10Ω 0603	10Ω 0603

Figure 172: Reference Layout to Support VCCINT/VCCINT_IO/VCCBRAM Connection Matrix



4. I/O Changes from UltraScale to UltraScale+ FPGAs

HRIO banks in UltraScale FPGAs become HPIO banks in Virtex UltraScale+ FPGAs.

★ IMPORTANT! There are a number of cases in which two “half” banks in an UltraScale device correspond to one “full” bank in an UltraScale+ device, and vice versa. When migrating from a device with half banks to a device with full banks, ensure that the voltage level on all V_{CCO} pins for the half banks match the intended voltage of the corresponding full bank. Refer to the following table and to the I/O bank migration table in UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575) to determine if this situation applies to your intended migration path.

Table 138: Corresponding Half-Banks and Full-Banks

Package	Half Banks	Full Bank
A676	XCKU3P XCKU5P (Banks 84/85)	XCKU035 XCKU040 (Bank 64)
A1156	XCKU11P (Banks 88/89)	XCKU025 XCKU035 XCKU040 XCKU060 XCKU095 (Bank 64)
A1156	XCKU15P (Banks 90/91)	XCKU025 XCKU035 XCKU040 XCKU060 XCKU095 (Bank 64)

Table 138: Corresponding Half-Banks and Full-Banks (cont'd)


Package	Half Banks	Full Bank
C1517	XCVU065 XCVU080 XCVU095 (Banks 84/94)	XCVU3P (Bank 64)
A2104	XCKU115 XCVU080 XCVU095 XCVU125 (Banks 84/94)	XCVU5P XCVU7P XCVU9P XCVU13P (Bank 64)
B2104	XCKU095 XCKU115 XCVU080 XCVU095 XCVU125 XCVU160 XCVU190 (Banks 84/94)	XCVU5P XCVU7P XCVU9P XCVU11P XCVU13P (Bank 64)
C2104	XCVU095 XCVU125 (Banks 84/94)	XCVU5P XCVU7P XCVU9P XCVU11P XCVU13P (Bank 64)

5. Transceiver Changes from UltraScale FPGAs to UltraScale+ FPGAs

- All GTH transceivers in UltraScale FPGAs become GTY transceivers in Virtex UltraScale+ FPGAs.
- Some GTY transceiver voltage levels change from UltraScale to UltraScale+ FPGAs, as noted in [Table 136: Voltages Differences between UltraScale and UltraScale+ FPGAs](#).
- While still maintaining footprint compatibility, some devices might have one or more extra RCAL and RREF pins than smaller dies in the same package. If planning to migrate from a smaller die to a bigger die in the same package, AMD recommends reserving any extra RCAL and RREF pins for later use, because they are NC pins in the smaller devices.
- Special exceptions for the F1924 package are:
 - VU11P-FLVF1924 GTY transceivers run up to 16.3 Gb/s versus the standard GTY transceiver speed of 32.75 Gb/s.

- Two ground pins (G9 and G10) in the UltraScale KU085/115 F1924 device later give way to becoming RREF (G9) and RCAL (G10) pins in the UltraScale+ VU11P F1924 device. The recommended way to design for this case would be by using the schematic shown in the following table and figure. When using the KU085/115 F1924, the G9 and G10 pins are grounded, and the resistors can be left unpopulated. When the RCAL and RREF pins are present when using the VU11P, both resistors should be populated.

Figure 173: Schematic for Reserving Future RCAL and RREF Pins for F1924



UG583_c7_02_081315

Table 139: Reserving RCAL and RREF Pins for F1924

F1924	G9	G10	Resistors
KU085/KU115	GND	GND	Do not populate
VU11P	MGTRREF	MGTAVTTRCAL	Populate

6. Configuration

- UltraScale+ FPGAs do not support 3.3V operation on bank 0:
 - Refer to *Interfacing 7 Series FPGAs High-Performance I/O Banks with 2.5V and 3.3V I/O Standards* (XAPP520) for methods of level translation between high-performance I/O banks and 2.5V/3.3V I/O standards.
- UltraScale+ FPGAs do not support master SelectMAP or master serial configuration modes:
 - Slave SelectMAP with an externally sourced clock can be used in place of master SelectMAP.

- SPIx1 or slave serial (with an externally sourced clock) can be used in place of master serial mode.

7. Memory Interface PCB Routing

The following table illustrates the expected PL DDR4 SDRAM performance for UltraScale+ FPGAs in the presence of migration for a single-rank component.

Table 140: DDR4 SDRAM Performance in UltraScale+ FPGAs in the Presence of Migration for Single-Rank Component

DDR4 Data Rate in UltraScale FPGAs	$V_{CCINT} = 0.9V$	UltraScale+ $V_{CCINT} = 0.85V$		UltraScale+ $V_{CCINT} = 0.72V$
	-3 (Max = 2667 Mb/s)	-2I, -2E, -2LE (Max = 2667 Mb/s)	-1E, -1I (Max = 2400 Mb/s)	-2LE (Max = 2400 Mb/s)
-3 @ 2400	2400	2400	2133	2133
-3 @ 2133	2133	2133	2133	2133
-3 @ 1866	1866	1866	1866	1866
-2E, -2I @ 2400	2400	2400	2133	2133
-2E, -2I @ 2133	2133	2133	2133	2133
-2E, -2I @ 1866	1866	1866	1866	1866
-1, -1I @ 2133	2133	2133	2133	2133
-1, -1I @ 1866	1866	1866	1866	1866

Notes:

1. This table applies only to PL memory interfaces. PS interfaces do not derate in performance due to speed grade.

8. SMBAlert

The SMBAlert pin is new to UltraScale+ FPGAs. If this feature is used in a migration from UltraScale to UltraScale+ FPGAs, the corresponding UltraScale FPGA pin must be reserved.

9. Block RAM

Typically, as device densities get larger, the amount of block RAM also keeps increasing. However, in certain UltraScale+ devices, the amount of block RAM might be lesser than in devices of smaller density. This is due to the addition of UltraRAM resources to UltraScale+ FPGAs and MPSoCs. With UltraRAM, there is an increase in overall bit count compared to UltraScale FPGAs but with fewer block RAM resources available. AMD advises coding the memory to use features present in both block RAM and UltraRAM, and the tools target UltraRAM as applicable.

10. ESD Requirements

UltraScale FPGA ESD handling requirements are based on JEDEC JEP155, available at www.jedec.org/standards-documents/docs/jep-155, as well as Industry Council ESD-CDM Target Levels, JEP157, available at www.esdindustrycouncil.org/ic/en.

11. PCI Express

Some UltraScale+ FPGAs only contain one PCI reset pin (PERSTN0). PERSTN1 is not available.


12. Migration Examples

This section describes two examples of differing complexity to illustrate the procedures for assessing migration.

Example One: VU125 to VU7P in C2104 Package

The following figure shows a die-level floorplan for a design using the VU125 in the C2104 package. This design includes several memory interfaces, a 100G Ethernet interface, an Interlaken 12x12.5G interface, and a Tandem PCIe interface. The bank(s) that each interface uses are shown color-coordinated, along with the relevant block locations.


Figure 174: Example VU125 Design Floorplan



UG583_c7_03_081315

The following figure shows how the various interfaces can migrate into a VU7P device that is also in the C2104 package. All interfaces can and must map to the same banks (and corresponding pins), as well as adhering to all block location requirements such as MAC, Interlaken, and PCIe banks needing to be within one row of their respective locations on the die. This particular example is reasonably straightforward in that the two device floorplans are very close to equal, with similar bank and block locations along with similar SLR boundaries.

Figure 175: VU7P Floorplan for Migration from VU125




UG583_c7_04_081315

Example Two: VU190 to VU13P in A2577 Package

The following figure shows a die-level floorplan for a design using the VU190 device in the A2577 package. This design is similar to the design in the previous example in that it includes several memory interfaces, a 100G Ethernet interface, an Interlaken 12x12.5G interface, and a Tandem PCIe interface.


Figure 176: Example VU190 Design Floorplan



UG583_c7_05_081315

The following figure shows how this design could migrate into a VU13P device. Careful bank selections for the various interfaces are necessary due to the notable differences in the device floorplans. The key differences in the device floorplans are such that there are more available banks in the VU13P, the SLRs cross different bank boundaries, and the block locations are noticeably different in the two devices relative to their neighboring banks.

Figure 177: VU13P Floorplan for Migration from VU190



UG583_c7_06_081315

13. Migration Scenarios

Not all migration scenarios have the same level of complexity in the 80 paths shown in [Table 135: UltraScale to UltraScale+ FPGA Migration Paths by Package](#). The following table outlines various scenarios in various degrees of complexity.

Migration from a Kintex or Virtex UltraScale FPGA into a Virtex UltraScale+ FPGA built with a five row tall SLR is seamless compared to migrating into a device made up of SLRs that are four rows tall. However, all migration paths are supported as long as the customer plans upfront on the extra considerations needed to make this switch successful.

An example of an easy and seamless migration scenario is from the VU095 to the V3P where both devices are monolithic. Similarly, migrating from the KU115, VU125, VU160, or VU190 into the VU5P, VU7P, or VU9P is also seamless and straightforward because they are based on SLRs that are five rows tall. Extra consideration is needed to migrate to the VU11P or VU13P because these devices consist of SLRs that are four rows tall.

Table 141: UltraScale Device Migration Scenarios

Package	UltraScale FPGAs	UltraScale+ FPGAs	Package Migration	Comment
C1517	KU095, VU065, VU080, VU095	VU3P	Supported	Easy and seamless because all the devices are monolithic.
F1924	KU085, KU115	VU11P	Supported	Extra consideration needed because migration path is from devices with 2 SLRs (5 rows) to 3 SLRs (4 rows) tall.
A2104	VU080, VU095	VU5P, VU7P, VU9P, VU13P	Supported	Extra consideration needed because migration path is from monolithic device to SLR devices.
A2104	KU115, VU125	VU5P, VU7P, VU9P	Supported	Easy and seamless because all the devices are made up of 5 row tall SLRs.
A2104	KU115, VU125	VU13P	Supported	Extra consideration needed because migration path is from devices with 2 SLRs (5 rows) to 3 SLRs (4 rows) tall.
B2104	KU095, VU080, VU095	VU5P, VU7P, VU9P, VU11P, VU13P	Supported	Extra consideration needed because migration path is from monolithic device to SLR devices.
B2104	KU115, VU125, VU160, VU190	VU5P, VU7P, VU9P	Supported	Easy and seamless because all the devices are made up of 5 row tall SLRs.
B2104	KU115, VU125, VU160, VU190	VU11P, VU13P	Supported	Extra consideration needed because migration path is from devices with 2/3 SLRs (5 rows) to 3/4 SLRs (4 rows) tall.
C2104	VU095	VU5P, VU7P, VU9P, VU11P, VU13P	Supported	Extra consideration needed because migration path is from a monolithic device to SLR devices.

Table 141: UltraScale Device Migration Scenarios (cont'd)

Package	UltraScale FPGAs	UltraScale+ FPGAs	Package Migration	Comment
C2104	VU125, VU160, VU190	VU5P, VU7P, VU9P	Supported	Easy and seamless because all the devices are made up of 5 row tall SLRs.
C2104	VU125, VU160, VU190	VU11P, VU13P	Supported	Extra consideration needed because migration path is from devices with 2/3 SLRs (5 rows) to 3/4 SLRs (4 rows) tall.
A2577	VU190	VU9P	Supported	Easy and seamless because both the devices are made up of 3 to 5 row tall SLRs.
A2577	VU190	VU11P, VU13P	Supported	Extra consideration needed because migration path is from device with 3 SLRs (5 rows) to 3/4 SLRs (4 rows) tall.

Migration between Zynq UltraScale+ MPSoCs and Packages

When migrating a design from one Zynq UltraScale+ MPSoC to another, special attention must be paid to potential differences between the two devices. Refer to the checklist below, and the subsequent sections that explain each item.

Zynq UltraScale+ MPSoC Migration Checklist

1. Review footprint compatibility, including banks that are bonded vs. unbonded.
2. Review I/O and transceiver Quad numbering and locations.
3. Review power supply differences across speed and temperature grades.
4. Review -3 speed grade migration.
5. Review video codec unit (VCU) migration.
6. Review graphics processing unit (GPU) migration.
7. Review decoupling capacitor requirements.
8. Review PCI Express® requirements and block locations.
9. Review integrated 100G Ethernet requirements and block locations.
10. Review Interlaken requirements and block locations.
11. Review power supply and thermal requirements.
12. Review pin flight time differences between devices.

1. Footprint Compatibility between Packages

Any two packages with the same footprint identifier code are footprint compatible. The footprint identifier code consists of a package designator code and BGA pin count information (such as B1156 or C1156, for example). A customer that wants pin compatibility across any two devices in a given package should select a footprint compatible package such that they can be designed to be electrically compatible. If migrating from a bigger device to a smaller device within the

same package, some I/O and transceiver banks and their pins might be unbonded or not exist on the smaller device. Refer to the Footprint Compatibility, I/O Bank Migration, and Transceiver Quad Migration tables in *Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)* to determine which Zynq UltraScale+ MPSoCs are footprint compatible, as well as the extent to which the various banks might become unbonded or not exist.

Example 1

The following figure shows a snippet of the Footprint Compatibility table from *Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)*. The XCZU2CG/EG SBVA484 devices are footprint compatible with the XCZU3CG/EG A484 devices, but the XCZU2CG/EG A484 devices are NOT compatible with any of the A625 devices, even though they share the same die.

Figure 178: Example of Footprint Compatibility

Table 1-5: Footprint Compatibility

Packages	Footprint Compatible Devices		
SBVA484	XCZU2CG and XCZU2EG	XCZU3CG and XCZU3EG	
SFVA625	XCZU2CG and XCZU2EG	XCZU3CG and XCZU3EG	

X18640-011017

Example 2

The following figure shows a snippet of the I/O Bank Migration table from *Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)*. I/O banks 72, 73, and 74 are present on the XCZU17 and ZCZU19 B1517 devices, but those I/O banks are not available on the XCZU11 B1517 device.

Figure 179: Example I/O Bank Migration within B1517 Devices

Table 1-6: I/O Bank Migration (HD Banks are Shaded) (Cont'd)

Package	Device	Package to Device I/O Mapping ⁽¹⁾																Unbonded I/O Banks									
		A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P		Q	R	S	T	U	V	W	X	Y
FFVB1517	XCZU11			65	64	66										88	89	90			71	70	69	68	67		91
	XCZU17			65	64	66										90	91	93	74	73	72	71	70	69	68	67	94
	XCZU19			65	64	66										90	91	93	74	73	72	71	70	69	68	67	94

X18641-011017

Example 3

The following figure shows a snippet of the Transceiver Quad Migration table from *Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)*. Transceiver Quads 230 and 231 are present on the XCZU11 F1517 device, but those transceiver Quads are not available on the XCZU7 F1517 device.

Figure 180: Example Transceiver Quad Migration within F1517 Devices

Power Supply Group		RS						RN		
FFVF1517	XCZU7	223	224	225	226	227	228			
	XCZU11	224	225	226	227	228	229	230	231	
									131, 130, 129, 128, 127	

X18642-022218

2. I/O Bank and Transceiver Quad Numbers

Certain Zynq UltraScale+ MPSoCs that are footprint compatible in a package might have different I/O bank and transceiver Quad numbers associated with the same package pins. Depending on where these banks are physically located on the die, certain multi-bank interfaces (such as memory or transceiver) could be impacted in regards to migrating from one device to the other. If a bank number change also involves the bank moving from one column on the die to another, a multi-bank interface that requires both banks to be in the same column could be impacted.

Refer to the I/O Bank Migration and Transceiver Quad Migration tables in *Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)* for the bank and transceiver Quad numbering differences. To cross-reference bank locations on a die, refer to the Die Level Bank Numbering Overview section of *Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)*.

The following example shows a successful migration path despite both bank number changes and banks being in different columns between the two different devices.

Example: Planning Ahead for Bank Number and Location Changes

As shown in the following figure, banks 88, 89, and 90 in the XCZU11-FFVB1517 device connect to the same pins as banks 90, 91, and 93 in the XCZU17-FFVB1517 and XCZU19-FFVB1517 devices. [Figure 182: XCZU11 Bank Locations for B1517](#) and [Figure 183: ZCZU17 and XCZU19 Bank Locations for B1517](#) show how the banks are arranged on each particular device. For the XCZU11, the banks are contiguous and could support a three-bank interface that requires the banks to be contiguous. However, the three banks that connect to the same pins on the XCZU17 and ZXCUI9 are not contiguous, and this might result in timing differences and possible design rule issues for interfaces that require the banks to be contiguous, such as memory interfaces.

Figure 181: I/O Bank Migration for B1517

Table 1-6: I/O Bank Migration (HD Banks are Shaded) (Cont'd)

Package	Device	Package to Device I/O Mapping ^[1]																										Unbonded I/O Banks	
		A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z		
FFVB1517	XCZU11			65	64	66											88	89	90			71	70	69	68	67		91	
	XCZU17			65	64	66											90	91	93	74	73	72	71	70	69	68	67		94
	XCZU19			65	64	66											90	91	93	74	73	72	71	70	69	68	67		94

X18643-081518

Figure 182: XCZU11 Bank Locations for B1517

GTU Quad 131 X0Y16-X0Y19	PCIE4 X0Y4	HP I/O Bank 71 T	HD I/O Bank 91	GTH Quad 231 X0Y28-X0Y31
GTU Quad 130 X0Y12-X0Y15	CMAC X0Y1	HP I/O Bank 70 U	HD I/O Bank 90 P	GTH Quad 230 X0Y24-X0Y27
GTU Quad 129 X0Y8-X0Y11 (RCAL)	ILKN X0Y0	HP I/O Bank 69 V	HD I/O Bank 89 O	GTH Quad 229 X0Y20-X0Y23
GTU Quad 128 X0Y4-X0Y7	PCIE4 X0Y2	HP I/O Bank 68 W	HD I/O Bank 88 N	GTH Quad 228 X0Y16-X0Y19
GTU Quad 127 X0Y0-X0Y3	CMAC X0Y0	HP I/O Bank 67 X	PCIE4 X1Y1	GTH Quad 227 X0Y12-X0Y15 D [R]
PS GTR 505	PS MIO 502	HP I/O Bank 66 E	SYSMON Configuration	GTH Quad 226 X0Y8-X0Y11 C [R] (RCAL)
PS DDR 504	PS MIO 501	HP I/O Bank 65 C	Configuration	GTH Quad 225 X0Y4-X0Y7 B [R]
PS CONFIG 503	PS MIO 500	HP I/O Bank 64 D	PCIE4 X1Y0 (tandem)	GTH Quad 224 X0Y0-X0Y3 A [R]

X18627-010917

Figure 183: ZCZU17 and XCZU19 Bank Locations for B1517

GTY Quad 134 X0Y28-X0Y31	CMAC X0Y3	HP I/O Bank 74 Q	HD I/O Bank 94	GTH Quad 234 X0Y40-X0Y43
GTY Quad 133 X0Y24-X0Y27	ILKN X0Y3	HP I/O Bank 73 R	HD I/O Bank 93 P	GTH Quad 233 X0Y36-X0Y39
GTY Quad 132 X0Y20-X0Y23	CMAC X0Y2	HP I/O Bank 72 S	ILKN X1Y2	GTH Quad 232 X0Y32-X0Y35
GTY Quad 131 X0Y16-X0Y19	PCIE4 X0Y4	HP I/O Bank 71 T	HD I/O Bank 91 O	GTH Quad 231 X0Y28-X0Y31
GTY Quad 130 X0Y12-X0Y15	CMAC X0Y1	HP I/O Bank 70 U	HD I/O Bank 90 N	GTH Quad 230 X0Y24-X0Y27
GTY Quad 129 X0Y8-X0Y11 (RCAL)	ILKN X0Y0	HP I/O Bank 69 V	ILKN X1Y1	GTH Quad 229 X0Y20-X0Y23
GTY Quad 128 X0Y4-X0Y7	PCIE4 X0Y2	HP I/O Bank 68 W	PCIE4 X1Y3	GTH Quad 228 X0Y16-X0Y19
GTY Quad 127 X0Y0-X0Y3	CMAC X0Y0	HP I/O Bank 67 X	PCIE4 X1Y1	GTH Quad 227 X0Y12-X0Y15 D [R]
PS GTR 505	PS MIO 502	HP I/O Bank 66 E	SYSMON Configuration	GTH Quad 226 X0Y8-X0Y11 C [R] (RCAL)
PS DDR 504	PS MIO 501	HP I/O Bank 65 C	Configuration	GTH Quad 225 X0Y4-X0Y7 B [R]
PS CONFIG 503	PS MIO 500	HP I/O Bank 64 D	PCIE4 X1Y0 (tandem)	GTH Quad 224 X0Y0-X0Y3 A [R]

X18628-010917

3. Power Supply Differences across Speed Grades and Temperature Grades

For programmable logic (PL), VCCINT, VCCINT_IO, and VCCBRAM have voltage limits that can differ across the speed and temperature grades as shown in the following table.

For the processing system (PS), VCC_PSINTFP, VCC_PSINTPL, VCC_PSINTFP_DDR, and VPS_MGTRAVCC can have voltage limits that differ across the speed and temperature grades as shown in [Table 143: MPSoC PS Voltage Matrix by Speed/Temperature Grade](#).

Table 142: MPSoC PL Voltage Matrix by Speed/Temperature Grade

Speed/Temperature Grade	VCCINT (V)	VCCINT_IO (V)	VCCBRAM (V)
-1E	0.85	0.85	0.85
-2E	0.85	0.85	0.85
-2LE	0.72 or 0.85	0.85	0.85
-3E	0.90	0.90	0.90
-1I	0.85	0.85	0.85
-1LI	0.72 or 0.85	0.85	0.85
-2I	0.85	0.85	0.85

Notes:

- VCCINT_IO must connect to VCCBRAM.

Table 143: MPSoC PS Voltage Matrix by Speed/Temperature Grade

Speed/Temperature Grade	VCC_PSINTFP (V)	VCC_PSINTLP (V)	VCC_PSINTFP_DDR (V)	VPS_MGTRAVCC (V)
-1E	0.85	0.85	0.85	0.85
-2LE	0.85	0.85	0.85	0.85
-3E	0.90	0.90	0.90	0.85
-1I	0.85	0.85	0.85	0.85
-1LI	0.85	0.85	0.85	0.85
-2I	0.85	0.85	0.85	0.85

Notes:

- VCC_PSINTFP_DDR must connect to VCC_PSINTFP.

4. -3 Speed Migration

CG devices and the XCZU2EG/XCZU3EG are not available in the -3 speed grade. CG devices can only migrate to EG or EV devices in non -3 speed grades to maintain similar performance. Likewise, EG and EV devices can only migrate to a CG device in non -3 speed grades to maintain similar performance. EG to EV devices can migrate to any speed grades, as well as EV to EG.

5. VCU Migration

The VCU is only available in EV devices. EV devices have a number of VCCINT_VCU power pins that correspond to ground pins on the CG and EG devices. If moving from a CG or EG to an EV device, the corresponding pins (VCCINT_VCU) should remain grounded if the VCU is not used.




IMPORTANT! A feasible migration path only exists for moving from an EV device to a CG/EG device. To go to an EV device from a CG/EG device, the PCB would need to be modified substantially to add the regulator for VCCINT_VCU.

Example: Schematic/PCB Options to Enable VCU Migration from EV to CG/EG Devices

The following figure shows a schematic representation of a connection option to connect VCCINT_VCU to ground when moving from an EV device to a CG or EG device. The migration path would be to DNP the VCCINT_VCU regulator and populate the two previously DNP resistors to ground the VCCINT_VCU planelet. To connect to ground (CG/EG devices), populate the resistors on the side with zero ohms. Two resistors are used to minimize inductance and increase current carrying ability.

Figure 184: VCCINT_VCU Connection Options for CG/EG/EV Migration (Schematic)



X18629-020718

6. GPU Migration

The GPU is only available in EG and EV devices. If migrating from an EG or EV device to a CG device, the GPU should not be used.

7. Decoupling Capacitors

If the design is not expected to change when migrating to the new device, the same decoupling scheme can be used. If adding logic, AMD recommends using a decoupling scheme appropriate for the new device and design. Refer to the decoupling capacitor guidelines in the Power Design Manager (PDM) tool (download at www.amd.com/power).

8. PCI Express Migration

PCI Express lane assignments might differ between devices. In addition, the block that is used as tandem boot might also be different from device to device. Refer to the GT Locations appendix in *UltraScale+ Devices Integrated Block for PCI Express LogiCORE IP Product Guide (PG213)* for GT locations.

9. Integrated 100G Ethernet Migration

The 100G Ethernet IP core is a dedicated block that provides 100 Gb/s MAC and PCS logic capability. The block consists of a number of dedicated locations that are located at different points on different devices. There are specific rules for connecting these dedicated blocks to specific transceiver locations, which might impact migration to another device. In addition, the CAUI-4 and switchable CAUI-10/CAUI-4 modes require GTY transceivers.

Refer to *UltraScale+ Devices Integrated 100G Ethernet Subsystem Product Guide (PG203)* for a more comprehensive guide regarding designing with the 100G Ethernet core, with special attention to the Transceiver Selection Rules section.

10. Interlaken Migration

The AMD UltraScale architecture integrated IP core for Interlaken is a highly configurable integrated IP core that can support an overall bandwidth up to 150 Gb/s for protocol logic transmission. The core connects to serial transceivers at defined rates up to 12.5 Gb/s with GTH transceivers, and up to 25.78125 Gb/s with GTY transceivers. The block consists of a number of dedicated locations that are located at different points on different devices. There are specific rules for connecting these dedicated blocks to specific transceiver locations, which might impact migration to another device. Refer to *Interlaken 150G LogiCORE IP Product Guide (PG212)* for a more comprehensive guide regarding the proper generation and verification of Interlaken interfaces.

11. Power Supplies and Thermal Considerations

Migration across devices, especially from smaller to bigger devices, can result in higher power consumption requiring improved thermal management and power delivery. Consult the Power Design Manager (PDM) tool (download at www.amd.com/power) and select the largest device being considered to estimate power consumption and thermal requirements.



RECOMMENDED: AMD recommends use of the Delphi thermal model during thermal modeling of a package. The Delphi thermal model includes consideration of the thermal interface material parameters and the manufacture variation on the thermal solution. Examples of manufacture variations include the tolerance in airflow from a fan, the tolerance on performance of the heat pipe and vapor chamber, and the manufacture variation of the attachment of fins to the heat sink base and the flatness of the surface. Further information regarding thermal considerations can be found in the Thermal Management Strategy chapter of Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide ([UG1075](#)).

12. Pin Flight Times across Packages

Even if two devices are footprint compatible from a design and PCB standpoint, FPGA package flight times will be different across different devices, so they will need to be accounted for. Flight time information can be found within the AMD Vivado™ tools during the I/O planning stage or after synthesis. The ideal strategy to account for pin flight time differences is to deskew the printed circuit board when migrating to the new device. If this is not possible, AMD recommends laying out the printed circuit board with the final device in mind to maximize system performance for the long term. System performance might have to be derated when using the initial device. As a last option, choosing the midpoint of the range of flight times and routing the board based on that value can act as a compromise, though maximum system performance might not be achievable with this method. Pin flight time information can be obtained via the Package Pins tab within the Vivado tools, both in the I/O planning stage and after synthesis.

Example 1: Obtaining Pin Flight Times During I/O Planning

Open the Vivado tools and follow the following steps to create an I/O planning project:


1. Double-click **Create New Project**, then click **Next**.
2. Choose a project name and location, then click **Next**.
3. Specify **I/O Planning Project**, then click **Next**.
4. Click **Do not import I/O ports at this time**, then click **Next**.
5. Select the device for which you wish to desire pin delays, then click **Next**.
6. Click **Finish**.
7. When the project loads, click **Export I/O Ports** from the Project Manager pane.
8. Make sure **CSV** is checked, select a filename and location, then click **OK**.

The resultant CSV file lists all I/Os and their associated pin delays under the Min Trace and Max Trace columns.

Example 2: Obtaining Pin Flight Times after Synthesis

After synthesis has completed, open the synthesized design (**Synthesis** → **Open Synthesized Design**), then click the **Package Pins** tab, as shown in the following figure. Pin delays are found in the Min Trace Dly and Max Trace Dly columns.

Figure 185: Pin Delays shown in the Vivado Tools



X18631-010917


Alternatively, the Tcl command `write_csv {filename.csv}` can be used in the Tcl Console tab to export I/O information to a file that includes the pin delays.

Using the Enable Migration Feature for DDR4 Memory Designs (Vivado Tools 2016.3)

A new feature for DDR4 memories in the Memory Interface Generator (MIG) tool allows users to prepare for migration from one device to another by specifying the differences in address/command/control pin delays from the initial device to the migrated device.

To enable the DDR4 migration feature in the MIG tool, select the **Enable Migration** checkbox in the Advanced Options tab, as shown in the following figure.


Figure 186: Selecting the Enable Migration Feature for DDR4 Designs



X18632-010917

After enabling migration, select the **Migration Options** tab and enter the relevant skew values as outlined in *UltraScale Architecture-Based FPGAs Memory IP LogiCORE IP Product Guide (PG150)* between the two devices (see the following figure). The pin delays can be found via [Example 1: Obtaining Pin Flight Times During I/O Planning](#) or [Example 2: Obtaining Pin Flight Times after Synthesis](#). The delay values entered should be positive and between 0 and 75 ps.

Figure 187: Entering the Pin Delay Differences for Migration



X18633-010917

Migration between Virtex UltraScale+ VU13P Devices and VU27P/VU29P Devices

VU27P and VU29P devices contain faster GTM transceivers (supporting 58 Gb/s PAM4 technology) in place of some GTY transceiver locations in the VU13P. The addition of the GTM tiles in the VU27P/VU29P results in some key areas that affect the PCB migration path between these devices. The following sections contain the necessary items to review for optimal migration between the VU13P and VU27P/VU29P devices.

Power Supply Differences

New VCCINT_GT Rail

VU27P/VU29P devices have two additional rails VCCINT_GT_R and VCCINT_GT_L (hereby referred to as VCCINT_GT), in addition to the other rails present on the VU13P device. Some of the existing VCCINT balls are converted to the new VCCINT_GT balls. VCCINT_GT is powered at 0.9V for -3 speed grades, and 0.85V for all other speed grades. The VCCINT_GT rails can be tied to VCCINT/VCCINT_IO/VCCBRAM for all speed grades other than -2LE. For -2LE, with VCCINT at 0.72V, VCCINT_GT should connect to VCCINT_IO/VCCBRAM while VCCINT is powered separately.

The following table shows a connection matrix regarding VCCINT and VCCINT_IO/VCCBRAM/VCCINT_GT.

Table 144: VCCINT/VCCINT_IO/VCCBRAM/VCCINT_GT Connection Matrix

VCCINT	Combined Connection	Standalone Connection
0.85V/0.90V	VCCINT VCCINT_IO VCCBRAM VCCINT_GT	-

Table 144: VCCINT/VCCINT_IO/VCCBRAM/VCCINT_GT Connection Matrix (cont'd)

VCCINT	Combined Connection	Standalone Connection
0.72V (-2LE)	VCCINT_IO VCCBRAM VCCINT_GT	VCCINT

There are three migration scenarios of concern, and each is explained below:

1. Migration from VU13P -1/-2/-3 to VU27P/VU29P -1/-2/-3.

For this scenario, no accommodation needs to be made for VCCINT migration, because VCCINT, VCCINT_IO, VCCBRAM, and VCCINT_GT remain at the same voltage level (0.85V or 0.90V).

2. Migration from VU13P -2LE to VU27P/VU29P -1/-2/-3

For this scenario, two regulators are required. One regulator powers VCCINT, and another regulator powers VCCBRAM/VCCINT_IO. When using the VU13P -2LE, the VCCINT regulator is powered at 0.72V, and the VCCBRAM/VCCINT_IO regulator is powered at 0.85V or 0.90V. When migrating to the VU27P/VU29P -1/-2/-3, VCCINT is now powered at 0.85V or 0.90V just as VCCBRAM/VCCINT_IO and VCCINT_GT is. AMD recommends designing the VCCINT regulator such that it can be easily changed from 0.72V to 0.85V/0.90V via an external interface (such as PMBus) or with a BOM change such as resistor values.


3. Migration from VU13P (any speed grade) to VU27P/VU29P -2LE

The following figure shows a layout example that supports this scenario. Two regulators are required, one for VCCINT and one for VCCBRAM/VCCINT_IO/VCCINT_GT. When using a VU13P device, resistors R9–R16 should be populated while resistors R1–R8 should be unpopulated. When using a VU27P/VU29P -2LE device, resistors R1–R8 should be populated to properly connect VCCBRAM/VCCINT_IO to the VCCINT_GT pins, while resistors R9–R16 should be unpopulated to keep VCCINT at a separate 0.72V.

Note: This solution requires separate PCB planes for VCCINT, VCCINT_GT, and VCCBRAM/VCCINT_IO.

The number of 0Ω resistors required on the PCB is dependent on the current consumption and capabilities of the resistor device and layout implementation.

Figure 188: Layout Migration Example for Migration to VU27P/VU29P-2LE



X20555-032218

VRM Sizing and Decoupling Capacitor Quantities

VRM Sizing

For non -2LE devices, where VCCINT, VCCBRAM, VCCINT_IO, and VCCINT_GT are all on one common plane, the voltage regulator module (VRM) power output does not need to change versus the VU13P device. However, for -2LE devices VCCINT_GT consumes up to 12W of power. This extra 12W must be accounted for on the common VCCBRAM/VCCINT_IO/VCCINT_GT plane. Likewise, the VCCINT VRM can be downward sized by approximately 12W.

Decoupling Capacitor Quantities

For non -2LE devices, the decoupling for the VU29P is exactly the same as for the VU13P. The VU27P in non -2LE has slightly less decoupling due to fewer available resources. For -2LE devices, however, the VU27P and VU29P have fewer VCCINT decoupling capacitors and more VCCBRAM/VCCINT_IO/VCCINT_GT. The following table lists decoupling capacitor requirements for non -2LE devices, while the next table lists capacitors for -2LE devices.

For migration planning, include pads for the highest decoupling count that will be required between the VU13P and VU27P/VU29P. This allows for selective population depending on the actual device used on the PCB.

Table 145: Decoupling for Non-2LE Devices

	VCCINT/VCCBRAM/VCCINT_IO/ VCCINT_GT ¹			VCCAUX/VCCAUX_IO		HPIO (per bank) ²
	680 μ F	100 μ F	10 μ F	47 μ F	10 μ F	47 μ F
XCVU13P-D2104	7	11	21	3	5	1
XCVU13P-A2577	7	11	21	2	3	1
XCVU27P-D2104	5	8	14	3	4	1
XCVU27P-A2577	5	8	14	2	3	1
XCVU29P-D2104	7	11	21	3	5	1
XCVU29P-A2577	7	11	21	2	3	1

Notes:

1. VCCINT_GT not present on VU13P devices.
2. One 47 μ F capacitor can be shared with up to four connected HPIO banks.

Table 146: Decoupling for -2LE Devices

	VCCINT			VCCBRAM/VCCINT_IO/ VCCINT_GT ¹			VCCAUX/ VCCAUX_IO		HPIO (per bank) ²
	680 μ F	100 μ F	10 μ F	680 μ F	100 μ F	10 μ F	47 μ F	10 μ F	47 μ F
XCVU13P-D2104	6	10	19	1	1	2	3	5	1
XCVU13P-A2577	6	10	19	1	1	2	2	3	1
XCVU27P-D2104	4	7	13	1	1	3	3	4	1
XCVU27P-A2577	4	7	13	1	1	3	2	3	1
XCVU29P-D2104	6	10	19	1	1	4	3	5	1
XCVU29P-A2577	6	10	19	1	1	4	2	3	1

Notes:

1. VCCINT_GT not present on VU13P devices.
2. One 47 μ F capacitor can be shared with up to four connected HPIO banks.


Package Height and Thermal Design Differences

Package Height

The D2104 package height does not change between VU13P and VU27P/VU29P, so no design accommodations for migration are necessary.

The VU13P A2577 is offered in lidded (FLGA2577) and unlidded versions (FSGA2577). The VU27P/VU29P are only offered unlidded (FSGA2577). If planning to migrate from a VU13P to VU27P/VU29P, it is recommended to use the unlidded version of the VU13P (FSGA2577). However, if going from a lidded VU13P (FLGA2577) to a VU27P/VU29P-FSGA2577, accommodations for the differing package dimensions must be accounted for. The following figure shows the FLGA2577 and FSGA2577 package drawings. The package height for the FSGA2577 is noticeably higher than the FLGA2577.

Figure 189: Package Dimension Differences between FLGA2577 and FSGA2577



X20556-032218


Thermal Design

A new heat sink must only be designed for the VU27P/VU29P-FSGA2577 if migrating from a lidded VU13P-FLGA2577. If migrating from an unlidded VU13P-FSGA2577 to VU27P/VU29P-FSGA2577, no extra thermal accommodations need to be made.

Package Flight Time Differences

Propagation delay for a given net across two devices in a footprint compatible package are not identical. Skews within a data group are controlled, however, such that the maximum skew in a data byte is ± 30 ps. The following figure illustrates these potential differences. For migration purposes, a one speed grade performance hit should be assumed by default to account for this additional skew. See [Table 140: DDR4 SDRAM Performance in UltraScale+ FPGAs in the Presence of Migration for Single-Rank Component](#) for an example of speed grade performance impact.

Figure 190: Skew Illustration Across Devices



X20557-032218

Bonded/Unbonded Bank/Quad Differences

The VU13P and VU29P contain an SLR that is not available on the VU27P. This must be accounted for when planning to migrate from the VU13P to the VU27P. The following figure shows an example of the unavailable super logic region (SLR) on the VU27P versus the VU13P.

Figure 191: Unavailable SLR on VU27P

GTU Quad 135 X0Y60-X0Y63 AA [LN]	CMAC X0Y11	HP I/O Bank 75 K	ILKN X1Y7	GTU Quad 235 X1Y60-X1Y63 K [RN]
GTU Quad 134 X0Y56-X0Y59 Z [LN]	CMAC X0Y10	HP I/O Bank 74 J	SYSMON Configuration	GTU Quad 234 X1Y56-X1Y59 J [RN]
GTU Quad 133 X0Y52-X0Y55 Y [LN] (RCAL)	ILKN X0Y6	HP I/O Bank 73 I	Configuration	GTU Quad 233 X1Y52-X1Y55 I [RN] (RCAL)
GTU Quad 132 X0Y48-X0Y51 X [LN]	CMAC X0Y9	HP I/O Bank 72	PCIE4 X0Y3	GTU Quad 232 X1Y48-X1Y51 H [RN]
SLR Crossing				
GTU Quad 131 X0Y44-X0Y47 W [LUC]	CMAC X0Y8	HP I/O Bank 71 H (Partial)	ILKN X1Y5	GTU Quad 231 X1Y44-X1Y47 G [RUC]
GTU Quad 130 X0Y40-X0Y43 V [LUC]	CMAC X0Y7	HP I/O Bank 70 G	SYSMON Configuration	GTU Quad 230 X1Y40-X1Y43 F [RUC]
GTU Quad 129 X0Y36-X0Y39 U [LUC] (RCAL)	ILKN X0Y4	HP I/O Bank 69	Configuration	GTU Quad 229 X1Y36-X1Y39 E [RUC] (RCAL)
GTU Quad 128 X0Y32-X0Y35 T [LUC]	CMAC X0Y6	HP I/O Bank 68	PCIE4 X0Y2	GTU Quad 228 X1Y32-X1Y35 D [RUC]
SLR Crossing				
GTU Quad 127 X0Y28-X0Y31 S [LLC]	CMAC X0Y5	HP I/O Bank 67	ILKN X1Y3	GTU Quad 227 X1Y28-X1Y31 C [RLC]
GTU Quad 126 X0Y24-X0Y27 R [LLC]	CMAC X0Y4	HP I/O Bank 66 B (Partial)	SYSMON Configuration	GTU Quad 226 X1Y24-X1Y27 B [RLC]
GTU Quad 125 X0Y20-X0Y23 Q [LLC] (RCAL)	ILKN X0Y2	HP I/O Bank 65 C	Configuration	GTU Quad 225 X1Y20-X1Y23 A [RLC] (RCAL)
GTU Quad 124 X0Y16-X0Y19 AF [LLC]	CMAC X0Y3	HP I/O Bank 64	PCIE4 X0Y1 (tandem)	GTU Quad 224 X1Y16-X1Y19 P [RLC]
SLR Crossing				
GTU Quad 123 X0Y12-X0Y15 AE [LS]	CMAC X0Y2	HP I/O Bank 63 F	ILKN X1Y1	GTU Quad 223 X1Y12-X1Y15 O [RS]
GTU Quad 122 X0Y8-X0Y11 AD [LS]	CMAC X0Y1	HP I/O Bank 62 E	SYSMON Configuration	GTU Quad 222 X1Y8-X1Y11 N [RS]
GTU Quad 121 X0Y4-X0Y7 AC [LS] (RCAL)	ILKN X0Y0	HP I/O Bank 61 D	Configuration	GTU Quad 221 X1Y4-X1Y7 M [RS] (RCAL)
GTU Quad 120 X0Y0-X0Y3 AB [LS]	CMAC X0Y0	HP I/O Bank 60	PCIE4 X0Y0	GTU Quad 220 X1Y0-X1Y3 L [RS]

VU13P_A2577

GTM Dual 135 X0Y7 AA [LN]	CMAC X0Y8	HP I/O Bank 75 K	ILKN X1Y5	GTM Dual 235 X1Y7 K [RN]
GTM Dual 134 X0Y6 Z [LN]	CMAC X0Y7	HP I/O Bank 74 J	SYSMON Configuration	GTM Dual 234 X1Y6 J [RN]
GTM Dual 133 X0Y5 Y [LN] (RCAL)	ILKN X0Y4	HP I/O Bank 73 I	Configuration	GTM Dual 233 X1Y5 I [RN] (RCAL)
GTM Dual 132 X0Y4 X [LN]	CMAC X0Y6	HP I/O Bank 72	CMAC X0Y6	GTM Dual 232 X1Y4 H [RN]
SLR Crossing				
GTM Dual 131 X0Y3 W [LUC]	CMAC X0Y5	HP I/O Bank 71 H (Partial)	ILKN X1Y3	GTM Dual 231 X1Y3 G [RUC]
GTM Dual 130 X0Y2 V [LUC]	CMAC X0Y4	HP I/O Bank 70 G	SYSMON Configuration	GTM Dual 230 X1Y2 F [RUC]
GTM Dual 129 X0Y1 U [LUC] (RCAL)	ILKN X0Y2	HP I/O Bank 69	Configuration	GTM Dual 229 X1Y1 E [RUC] (RCAL)
GTM Dual 128 X0Y0 T [LUC]	CMAC X0Y3	HP I/O Bank 68	CMAC X1Y3	GTM Dual 228 X1Y0 D [RUC]
SLR Crossing				
GTU Quad 127 X0Y12-X0Y15 S [LLC]	CMAC X0Y2	HP I/O Bank 67	ILKN X1Y1	GTU Quad 227 X1Y12-X1Y15 C [RLC]
GTU Quad 126 X0Y8-X0Y11 R [LLC]	CMAC X0Y1	HP I/O Bank 66 B (Partial)	SYSMON Configuration	GTU Quad 226 X1Y8-X1Y11 B [RLC]
GTU Quad 125 X0Y4-X0Y7 Q [LLC] (RCAL)	ILKN X0Y0	HP I/O Bank 65 C	Configuration	GTU Quad 225 X1Y4-X1Y7 A [RLC] (RCAL)
GTU Quad 124 X0Y0-X0Y3 AF [LLC]	CMAC X0Y0	HP I/O Bank 64	PCIE4 (tandem)	GTU Quad 224 X1Y0-X1Y3 P [RLC]


VU27P_A2577

X20558-032218

CMAC/PCIe Location Differences

PCIe® hard blocks are replaced with CMAC blocks on any SLR that contains GTM duals. GTM duals do not offer PCIe support, though PCIe is still supported in SLRs with GTY transceivers. The following figure illustrates these differences.

Figure 192: VU13P Floorplan with PCIe/GTY versus VU29P Floorplan with CMAC/GTM



X20559-032218

GTY/GTM Channel and Clock Mapping Differences

GTY and GTM channels and clocks map from device to device are shown in the following figures. Key design items to note about GTY/GTM reference clocks are:

- Reference clock rates between 25G and 50G are not identical
 - Assume a different clock is necessary at 50G
- RefClk0 must always be used
- RefClk1 migrates to NC on VU27P/VU29P
- Reference clocks should not be shared across GTM duals
- Reference clocks should not be shared across GTY Quads
- Provide local reference clocks to each GTY Quad that is migrated into a GTM Quad

Figure 193: VU13P/VU27P D2104 GTY/GTM TX/RX Channel/Clock Mapping




Figure 194: VU13P/VU29P D2104 GTY/GTM TX/RX Channel/Clock Mapping




Figure 195: VU13P/VU27P A2577 GTY/GTM TX/RX Channel/Clock Mapping






Figure 196: VU13P/VU29P A2577 GTY/GTM TX/RX Channel/Clock Mapping



100G Ethernet Layout Guidelines

Four channels across two GTY Quads instead of all four channels in a given Quad should be used to connect to the QSFP optics to ensure migration compatibility. The unbonded GTY transceivers disappear when moving to the dual architecture of the VU27P/VU29P. The following figure illustrates the GTY to GTM transceiver mapping concept.

Figure 197: GTY to GTM Transceiver Mapping for 100G Ethernet



X20564-032318

SelectIO Signaling

The AMD UltraScale™ architecture SelectIO resources are the general-purpose I/O and its various settings. With numerous I/O standards and hundreds of variants within these standards, these SelectIO resources offer a flexible array of choices for designing I/O interfaces.

This chapter provides some strategies for choosing I/O standard, topography, and termination, and offers guidance on simulation and measurement for more detailed decision making and verification. In many cases, higher-level aspects of the system (other device choices or standards support) define the I/O interfaces to be used. In cases where such constraints are not defined, it is up to the system designer to choose I/O interface standards and optimize them according to the purpose of the system.

Interface Types

To better address the specifics of the various interface types, it is necessary to first break interfaces into categories. Two relevant divisions are made:

- [Single-Ended versus Differential Interfaces](#)
- [SDR versus DDR Interfaces](#)

Single-Ended versus Differential Interfaces

Traditional digital logic uses single-ended signaling – a convention that transmits a signal and assumes a GND common to the driver and receiver. In single-ended interfaces, a signal's assertion (whether it is High or Low) is based on its voltage level relative to a fixed voltage threshold that is referenced to GND. When the voltage of the signal is higher than the V_{IH} threshold, the state is considered High. When the voltage of the signal is lower than the V_{IL} threshold, the state is considered Low. TTL is one common example of a single-ended I/O standard.

To reach higher interface speeds and increase noise margin, some single-ended I/O standards rely on a precise dedicated local reference voltage other than GND. HSTL and SSTL are examples of I/O standards that rely on a V_{REF} to resolve logic levels. V_{REF} can be thought of as a fixed comparator input.

Higher-performance interfaces typically make use of differential signaling – a convention that transmits two complementary signals referenced to one another. In differential interfaces, a signal's assertion (whether it is High or Low) is based on the relative voltage levels of the two complementary signals. When the voltage of the P signal is higher than the voltage of the N signal, the state is considered High. When the voltage of the N signal is higher than the voltage of the P signal, the state is considered Low. Typically the P and N signals have similar swing, and have a common-mode voltage above GND (although this is not always the case). LVDS is one common example of a differential I/O standard.

SDR versus DDR Interfaces

The difference between Single Data Rate (SDR) and Double Data Rate (DDR) interfaces has to do with the relationship of the data signals of a bus to the clock signal of that bus. In SDR systems, data is only registered at the input flip-flops of a receiving device on either the rising or the falling edge of the clock. One full clock period is equivalent to one bit time. In DDR systems, data is registered at the input flip-flops of a receiving device on both the rising and falling edges of the clock. One full clock period is equivalent to two bit times. The distinction of SDR and DDR has nothing to do with whether the I/O standard carrying the signals is single-ended or differential. A single-ended interface can be SDR or DDR, and a differential interface can also be SDR or DDR.

Single-Ended Signaling

A variety of single-ended I/O standards are available in the UltraScale architecture I/O. For a complete list of supported I/O standards and detailed information about each one, refer to the *SelectIO Resources* chapter of the *UltraScale Architecture SelectIO Resources User Guide* ([UG571](#)). Tables at the end of this chapter summarize for each supported I/O standard which ones support DRIVE and SLEW attributes, bidirectional buffers, and the DCI options. It also describes which I/O standards are supported in the high-performance (HP) and high-range (HR) I/O banks.

Modes and Attributes

Some I/O standards can be used only in unidirectional mode, while some can be used in bidirectional mode or unidirectional mode.

Some I/O standards have attributes to control drive strength and slew rate, as well as the presence of weak pull-up or pull-down and weak-keeper circuits (not intended for use as parallel termination). Drive strength and slew rate can be used to tune an interface for adequate speed while not overdriving the signals. Weak pull-ups, weak pull-downs, and weak keepers can be used to ensure a known or steady level on a floating or 3-stated signal. The *SelectIOResources* chapter of the *UltraScale Architecture SelectIO Resources User Guide* ([UG571](#)) describes which standards support these attributes. Refer to this user guide for more information.

LVC MOS, when set to 6 mA DRIVE and FAST slew, has an approximate output impedance close to 50Ω , allowing it to be used as a crude approximation of a controlled-impedance driver. The impedance match of the weak driver to the transmission line is only approximate and varies with voltage and temperature. LVDCI and HSLVDCI, true controlled-impedance drivers, are adaptive, maintain a much closer impedance match, and remain constant over voltage and temperature.

Input Thresholds

The input circuitry of the single-ended standards fall into two categories: those with fixed input thresholds and those with input thresholds set by the V_{REF} voltage. The use of V_{REF} has three advantages:

- It allows for tighter control of input threshold levels
- It removes dependence on die GND for the threshold reference
- It allows for input thresholds to be closer together, which reduces the need for a large voltage swing of the signal at the input receiver

Two 1.8V I/O standards that illustrate this are LVC MOS18 and SSTL18 Class 1. The thresholds for 1.8V LVC MOS are set at 0.63V and 1.17V (necessitating that the signal at the receiver swing a full 540 mV at minimum to make a logic transition). The thresholds for SSTL18 Class 1 are set at $V_{REF} - 0.125V$ and $V_{REF} + 0.125V$, or for a nominal V_{REF} of 0.9V, set at 0.775V and 1.025V (necessitating that the signal at the receiver only swing 250 mV at minimum to make a logic transition). This smaller required swing allows for higher frequency of operation in the overall link. A smaller swing at the driver means reduced DC power is required with less transient current. In the UltraScale architecture-based devices, the reference voltage can either be provided using the dedicated V_{REF} pins, or optionally generated internally using the Internal V_{REF} feature. See the *SelectIO Resources* chapter of the *UltraScale Architecture SelectIO Resources User Guide* (UG571) for more details on Internal V_{REF} . For more information on V_{REF} decoupling and decoupling of all other supplies, see [Chapter 1: Power Distribution System in UltraScale Devices](#).

Topographies and Termination

Topography generally refers to the arrangement of drivers, receivers, interconnect and terminations in an interface. The techniques used in unidirectional topographies are different from those used in bidirectional topographies, so these are treated separately.

The SelectIO standards can be used in countless topographies depending on the requirements of the system. SelectIO drivers and receivers adhering to a standard (SSTL, LVCMOS, etc.) either can be used according to the letter of the standard (published by a standards body such as EIA/TIA or JEDEC) or they can be mixed and matched with drivers or receivers from another standard or hybrid I/O. An I/O standard specification might define something as limited as the V_{IL} and V_{IH} of the receiver, or it might define every aspect of the interface, including driver impedance and slew rate, PCB trace length and topography, value and position of passive termination, the maximum input capacitance of a receiving device, and even the maximum number of receivers.

It is up to the designer to apply the standard in question to the system in which it is working. There are many decisions to make with respect to topographies and termination, which determine the signal integrity of the interface. It is of utmost importance that the signal integrity of each interface be verified through both simulation and measurement.

Termination generally refers to impedance-matching or impedance-compensating devices that are used to maintain signal integrity in an interface. While many types of elements can be used as terminators (such as, resistors, capacitors, diodes), this discussion is limited to resistive termination. In general, capacitor and diode termination techniques are more complicated.

Unidirectional Topographies and Termination

The two basic subsets of unidirectional topographies are point-to-point and multi-drop. A point-to-point topography has one driver and one receiver, while a multi-drop topography has one driver and many receivers. Whether or not a topography is point-to-point or multi-drop defines important aspects of the interface that determine which termination strategies are appropriate and which are not.

Unidirectional Point-to-Point Topographies

The simplest unidirectional topography is point-to-point. That is, there is one driver and one receiver. Termination, if present, can consist of parallel termination at the receiver (see [Figure 198: Parallel-Terminated Unidirectional, Point-to-Point Topography](#)), series termination at the driver (see [Figure 199: Series-Terminated Unidirectional, Point-to-Point Topography](#)), or a controlled-impedance driver ([Figure 200: DCI-Controlled Impedance Driver Unidirectional, Point-to-Point Topography](#) and [Figure 201: “Weak Driver” Unidirectional, Point-to-Point Topography](#)). Always use IBIS simulation to determine the optimal resistor values, V_{TT} voltage level, and VRP reference resistors for these terminations.

Figure 198: Parallel-Terminated Unidirectional, Point-to-Point Topography




Figure 199: Series-Terminated Unidirectional, Point-to-Point Topography




Figure 200: DCI-Controlled Impedance Driver Unidirectional, Point-to-Point Topography





Figure 201: "Weak Driver" Unidirectional, Point-to-Point Topography




In general, parallel resistive termination (R_p) has a value equal to the characteristic impedance (Z_0) of the transmission line it is terminating. Series resistive terminations (R_s) have a value equal to the characteristic impedance of the transmission line (Z_0) minus the output impedance of the driver (R_O) to which they are connected. Controlled-impedance drivers are tuned such that the driver output impedance (R_O) is equal to the characteristic impedance (Z_0) of the transmission line it is terminating.

Assuming transmission lines with 50Ω characteristic impedance and a driver output impedance (R_O) of 25Ω , a 25Ω series termination (Figure 199: Series-Terminated Unidirectional, Point-to-Point Topography) or a 50Ω parallel termination (Figure 198: Parallel-Terminated Unidirectional, Point-to-Point Topography) is appropriate. Controlled-impedance drivers, whether implemented with DCI or with weak LVCMOS drivers, should be sized to have an output impedance (R_O) of 50Ω . This corresponds to a VRP resistor equal to 50Ω for DCI. Weak LVCMOS drivers of 6 mA to 8 mA drive strength have an output impedance approximately equal to 50Ω (Figure 200: DCI-Controlled Impedance Driver Unidirectional, Point-to-Point Topography).

Typically, parallel terminations have best performance when V_{TT} (the voltage source connected to the parallel termination resistor) is equal to half of the signaling voltage. For 2.5V signals ($V_{CC0} = 2.5V$), V_{TT} is ideally 1.25V. In cases where this voltage is not available, it is possible to use a Thevenin parallel termination. Thevenin parallel termination consists of a voltage divider with a parallel equivalent resistance (R_{PEQ}) equal to the characteristic impedance of the transmission line (50Ω in most cases). The divided voltage point is designed to be at V_{TT} . The following figure illustrates a Thevenin parallel termination powered from 2.5V V_{CC0} , made up of two 100Ω resistors, resulting in a V_{TT} of 1.25V and a parallel equivalent resistance (R_{PEQ}) of 50Ω .

Parallel termination can be less desirable than series termination or controlled-impedance drivers because it dissipates more power. This trade-off must be weighed against other trade-offs to determine the optimum termination topography for an interface.

Figure 202: Thevenin Parallel Termination



The following table lists example I/O interface types that can be used with the unidirectional point-to-point topography.

Table 147: Example I/O Interface Type for Unidirectional Point-to-Point Topographies

I/O Interface
LVTTTL
LVC MOS
LVDCI
SSTL Class I
HSTL Class I

LVTTTL and LVC MOS do not specify any canonical termination method. Series termination at the driver or parallel termination at the receiver are both appropriate considerations.

LVDCI implicitly uses controlled-impedance driver termination. No form of termination is needed at the receiver.

Every I/O standard can have different requirements for termination techniques. In some cases the specification for the I/O standard can rigidly define the termination topology. Other standards might not have any hard requirements, but rather might simply provide examples of termination topologies. An example of a standard with specific termination requirements is HSTL. HSTL Class I is a unidirectional I/O standard that recommends a parallel termination at the receiver. In the case of HSTL Class I, the termination voltage V_{TT} is defined as half of the supply voltage V_{CC} . The designer can ultimately elect either not to use termination at all or to use a different termination, such as series termination at the driver. There are a number of reasons why this selection might be advantageous in a given system. It is up to the designer to verify through simulation and measurement that the signal integrity at the receiver is adequate.

The SSTL standards tend to not have rigid requirements for termination topology. Rather, the JEDEC specifications provide example termination techniques that tend to be the commonly used topographies. The *SelectIO Resources* chapter of the *UltraScale Architecture SelectIO Resources User Guide* (UG571) provides example termination techniques for each of the I/O standards, including the SSTL standards, for the purpose of providing a good starting point for consideration. Similar to HSTL, it is ultimately up to the designer to verify through simulation and measurement that the signal integrity at the receiver is adequate.

Unidirectional Multi-Drop Topographies

In more complex topographies, a single driver can drive multiple receivers. The receivers represent loads that must be fed by individual transmission line stubs. From a signal integrity standpoint, the best topography to use in this case is a single long transmission line with the driver at one end and parallel termination at the other, with receivers connected to the main trace by short stubs in between. This type of topography is often referred to as a *flyby multi-drop* topography.


There are two critical aspects of this topography. The first is the presence of a single parallel termination at the far end of the transmission line. Series termination at the driver or a controlled impedance driver must *never* be used. Parallel termination is the only applicable termination type for this topography. The second critical aspect is the length of the connecting stubs at each receiver. These must remain short: no more than a fraction of a signal rise time in length. With a typical signal rise time of 600 ps, a stub no longer than $600 \text{ ps}/4 = 150 \text{ ps}$, or 0.9 inches (22.86 mm) should be used. As the stubs become longer, they present a larger impedance discontinuity to the signal travelling down the transmission line, and can support significant reflections. These impedance discontinuities corrupt the signal. With increasing numbers of loads and increasing length of stubs, the signal is corrupted to the point where it is no longer usable.

Star topographies are not recommended. The constraints involved in designing a star topography with good signal integrity are beyond the scope of this document.

As stated in [Unidirectional Point-to-Point Topographies](#), ideal parallel resistive termination has a value equal to the characteristic impedance of the transmission line it is terminating. The best performance is achieved when V_{TT} is equal to half of the signaling voltage, and when this voltage is not available, a Thevenin parallel termination is recommended, as defined in the previous section.

The following figure illustrates a Thevenin parallel termination powered from V_{CC0} , made up of two 100Ω resistors, resulting in a V_{TT} of $V_{CC0}/2$ and a parallel equivalent resistance of 50Ω . This figure shows a topography with one driver (an LVCMOS driver) and four receivers. The driver is on the left side, the receivers are spaced at interim points across the 50Ω transmission line, and the Thevenin parallel termination of two 100Ω resistors is on the right side.

Figure 203: Basic Multi-Drop Topography



UG583_c2_06_112113

The main transmission line should be kept as short as possible. Lengths up to 20 inches or more are practical for most I/O standards as long as precise trace impedance is maintained and crosstalk sources are avoided. The lengths of interim segments of the main transmission line need not be equal. Their relative lengths can be arbitrary. Receivers at different points along the main transmission line receive the signal with varying amounts of delay, but all signal rise times are similar.

Stubs stretching from the main transmission line to the individual receivers must be kept as short as possible. The longer these stubs become, the more corrupted the received waveforms are. Simulation and measurement are required to assess signal integrity at the individual receivers.

The following table lists example I/O interface types that can be used with the unidirectional multi-drop topography.

Table 148: Example I/O Interface Types for Unidirectional Multi-Drop I/O Topographies

I/O Interface
LVTTL
LVC MOS
HSTL
SSTL

LVTTL and LVC MOS do not specify any canonical termination method. Parallel termination at the end of the long t-line is an appropriate termination method.

Bidirectional Topography and Termination

The two basic subsets of bidirectional topographies are point-to-point and multi-point. A point-to-point topography has two transceivers (driver and receiver sharing one device pin), while a multi-point topography can have many transceivers. Whether or not a topography is point-to-point or multi-point defines important aspects of the interface that determine which termination strategies are appropriate and which are not.

Bidirectional Point-to-Point Topographies

The simplest bidirectional topography is point to point. That is, there are two transceivers connected by a transmission line. Because bidirectional interfaces need to operate equally well in both directions, symmetry of the topography is desirable. While asymmetrical topographies can be designed with reasonably good signal integrity, the easiest way to ensure good signal integrity is to keep the topography symmetrical. Thus any termination used on one side of the link should also be used on the other side of the link. Series termination ([Figure 205: Series Terminated Bidirectional Point-to-Point Topography: Not Recommended](#)) is rarely appropriate for bidirectional interfaces as incoming signals are attenuated by the series resistor of the receiving transceiver. Parallel termination (see the following figure) almost always achieves better

signal levels at both receivers. Controlled-impedance drivers, whether crudely controlled in the form of a weak LVCMOS driver or adaptively controlled in the form LVDCI or HSLVDCI, also can have good results as shown in [Figure 206: “Weak Driver” Bidirectional Point-to-Point Topography](#) (implemented with a low-drive strength LVCMOS driver). Always use IBIS simulation to determine the optimal termination resistor value, V_{TT} voltage level and VRP reference resistor values for these terminations.

Figure 204: Parallel Terminated Bidirectional Point-to-Point Topography




Figure 205: Series Terminated Bidirectional Point-to-Point Topography: Not Recommended





Figure 206: “Weak Driver” Bidirectional Point-to-Point Topography




In general, parallel resistive termination (R_p) has a value equal to the characteristic impedance Z_0 of the transmission line it is terminating. Some interfaces, such as DDR2 memory interfaces, use 75Ω termination resistors instead of 50Ω in an effort to open the data eye. In this case, the trade-off is eye height against a small amount of signal reflection from the impedance discontinuity. Controlled-impedance drivers are typically tuned such that the driver output impedance (R_O) is equal to the characteristic impedance (Z_0) of the transmission line it is terminating.

Assuming transmission lines with 50Ω characteristic impedance and a driver output impedance of 25Ω, 50Ω parallel terminations are appropriate (Figure 204: Parallel Terminated Bidirectional Point-to-Point Topography). Controlled-impedance drivers, whether implemented with DCI or with weak LVC MOS drivers, should be sized to have an output impedance (R_O) of 50Ω. An example of the use of a controlled-impedance driver would be the LVDCI_15 I/O standard. Weak LVC MOS drivers of 6 mA to 8 mA drive strength have an output impedance approximately equal to 50Ω (Figure 206: “Weak Driver” Bidirectional Point-to-Point Topography). Figure 206: “Weak Driver” Bidirectional Point-to-Point Topography

Parallel terminations have the best performance when V_{TT} (the voltage source connected to the parallel termination resistor) is equal to half of the signaling voltage, because this is typically the center voltage of the data eye. For 2.5V signals ($V_{CCO} = 2.5V$), V_{TT} is ideally 1.25V. In cases where this voltage is not available, it is advisable to use a Thevenin parallel termination. Thevenin parallel termination consists of a voltage divider with a parallel resistance equal to the characteristic impedance of the transmission line (50Ω in most cases). The divided voltage point is designed to be at V_{TT} . The following figure illustrates a Thevenin parallel termination powered from 2.5V V_{CCO} , made up of two 100Ω resistors, resulting in a V_{TT} of 1.25V and a parallel equivalent resistance (R_{PEQ}) of 50Ω.

Parallel termination can be less desirable than series termination or controlled-impedance drivers because it dissipates more power. This trade-off must be weighed against other trade-offs to determine the optimum termination topography for an interface.

Figure 207: Thevenin Parallel Termination (Bidirectional Point-to-Point Topography)



The following table lists example I/O interface types that can be used with the bidirectional point-to-point topography.

Table 149: Example I/O Interface Types for Bidirectional Point-to-Point I/O Topographies

I/O Interface
LVTTTL
LVC MOS
LVDCI
HSLVDCI
SSTL15
SSTL15 DCI
SSTL18 CLASS II
SSTL18 CLASS II DCI
HSTL CLASS II
HSTL CLASS II DCI

LVTTTL and LVC MOS do not specify any canonical termination method. Series termination is not recommended for bidirectional interfaces. Parallel termination and weak drivers, however, are both appropriate.

LVDCI and HSLVDCI both implicitly use controlled-impedance driver termination.

HSTL Class II specifies parallel termination at both transceivers. The termination voltage V_{TT} is defined as half of the supply voltage V_{CC0} . The designer can elect either not to use termination at all or to use a different termination. It is up to the designer to verify through simulation and measurement that the signal integrity at the receiver is adequate.

The JEDEC specifications for SSTL provide examples of both series termination and parallel termination. The termination voltage V_{TT} is defined as half of the supply voltage V_{CCO} . While the specification document provides examples depicting series termination at the drivers, it is important to note that the purpose of this is to attempt to match the impedance of the driver with that of the transmission line. Because the UltraScale architecture SSTL drivers target to have output impedances close to 40–50 Ω , better signal integrity can be achieved without any external source-series termination. When possible, it is a better starting point to consider the use of the 3-state DCI I/O standards (“T_DCI”), which provide internal parallel termination resistors that are only present when the output buffer is in 3-state. It is up to the designer to carefully choose the I/O standard(s) at the AMD device, drive strengths, and on-die termination (ODT) options at the other device(s) in the interface (usually DRAM ICs) and termination topography though careful simulation and measurement. See the *UltraScale Architecture SelectIO Resources User Guide* ([UG571](#)) for more details on the available I/O standards and options.

Bidirectional Multi-Point Topographies

In more complex topographies, any transceiver in a multi-point bus can transmit to all other transceivers. Usually these topographies can only run at very slow clock rates because they only support very slow signal rise times (10 ns to 50 ns). While useful in some situations, the drawbacks usually outweigh the benefits. The constraints involved in designing these topographies with good signal integrity are beyond the scope of this document.

PCB Technology Basics

Introduction

Printed circuit boards (PCBs) are electrical systems, with electrical properties as complicated as the discrete components and devices mounted to them. The PCB designer has complete control over many aspects of the PCB; however, current technology places constraints and limits on the geometries and resulting electrical properties. The following information is provided as a guide to the freedoms, limitations, and techniques for PCB designs using FPGAs.

PCB Structures

PCB technology has not changed significantly in the last few decades. An insulator substrate material (usually FR4, an epoxy/glass composite) with copper plating on both sides has portions of copper etched away to form conductive paths. Layers of plated and etched substrates are glued together in a stack with additional insulator substrates between the etched substrates. Holes are drilled through the stack. Conductive plating is applied to these holes, selectively forming conductive connections between the etched copper of different layers.

While there are advancements in PCB technology, such as material properties, the number of stacked layers used, geometries, and drilling techniques (allowing holes that penetrate only a portion of the stackup), the basic structures of PCBs have not changed. The structures formed through the PCB technology are abstracted to a set of physical/electrical structures: traces, planes (or planelets), vias, and pads.

Traces

A trace is a physical strip of metal (usually copper) making an electrical connection between two or more points on an X-Y coordinate of a PCB. The trace carries signals between these points.

Planes

A plane is an uninterrupted area of metal covering the entire PCB layer. A planelet, a variation of a plane, is an uninterrupted area of metal covering only a portion of a PCB layer. Typically, a number of planelets exist in one PCB layer. Planes and planelets distribute power to a number of points on a PCB. They are very important in the transmission of signals along traces because they are the return current transmission medium.

Vias

A via is a piece of metal making an electrical connection between two or more points in the Z space of a PCB. Vias carry signals or power between layers of a PCB. In current plated-through-hole (PTH) technology, a via is formed by plating the inner surface of a hole drilled through the PCB. In current microvias technology (also known as high-density interconnect or HDI), a via is formed with a laser by ablating the substrate material and deforming the conductive plating. These microvias cannot penetrate more than one or two layers, however, they can be stacked or stair-stepped to form vias traversing the full board thickness.

Pads and Antipads

Because PTH vias are conductive over the whole length of the via, a method is needed to selectively make electrical connections to traces, planes, and planelets of the various layers of a PCB. This is the function of pads and antipads.

Pads are small areas of copper in prescribed shapes. Antipads are small areas in prescribed shapes where copper is removed. Pads are used both with vias and as exposed outer-layer copper for mounting of surface-mount components. Antipads are used mainly with vias.

For traces, pads are used to make the electrical connection between the via and the trace or plane shape on a given layer. For a via to make a solid connection to a trace on a PCB layer, a pad must be present for mechanical stability. The size of the pad must meet drill tolerance/registration restrictions.

Antipads are used in planes. Because plane and planelet copper is otherwise uninterrupted, any via traveling through the copper makes an electrical connection to it. Where vias are not intended to make an electrical connection to the planes or planelets passed through, an antipad removes copper in the area of the layer where the via penetrates.

Lands

For the purposes of soldering surface mount components, pads on outer layers are typically referred to as lands or solder lands. Making electrical connections to these lands usually requires vias. Due to manufacturing constraints of PTH technology, it is rarely possible to place a via inside the area of the land. Instead, this technology uses a short section of trace connecting to a surface pad. The minimum length of the connecting trace is determined by minimum dimension specifications from the PCB manufacturer. Microvia technology is not constrained, and vias can be placed directly in the area of a solder land.

Dimensions

The major factors defining the dimensions of the PCB are PCB manufacturing limits, FPGA package geometries, and system compliance. Other factors such as Design For Manufacturing (DFM) and reliability impose further limits, but because these are application specific, they are not documented in this user guide.

The dimensions of the FPGA package, in combination with PCB manufacturing limits, define most of the geometric aspects of the PCB structures described in this section ([PCB Structures](#)), both directly and indirectly. This significantly constrains the PCB designer. The package ball pitch (1.0 mm for FF packages) defines the land pad layout. The minimum surface feature sizes of current PCB technology define the via arrangement in the area under the device. Minimum via diameters and *keep-out areas* around those vias are defined by the PCB manufacturer. These diameters limit the amount of space available in-between vias for routing of signals in and out of the via array underneath the device. These diameters define the maximum trace width in these *breakout* traces. PCB manufacturing limits constrain the minimum trace width and minimum spacing.

The total number of PCB layers necessary to accommodate an FPGA is defined by the number of signal layers and the number of plane layers.

- The number of signal layers is defined by the number of I/O signal traces routed in and out of an FPGA package (usually following the total User I/O count of the package).
- The number of plane layers is defined by the number of power and ground plane layers necessary to bring power to the FPGA and to provide references and isolation for signal layers.

Most PCBs for large FPGAs range from 12 to 22 layers.

System compliance often defines the total thickness of the board. Along with the number of board layers, this defines the maximum layer thickness, and therefore, the spacing in the Z direction of signal and plane layers to other signal and plane layers. Z-direction spacing of signal trace layers to other signal trace layers affects crosstalk. Z-direction spacing of signal trace layers to reference plane layers affects signal trace impedance. Z-direction spacing of plane layers to other plane layers affects power system parasitic inductance.

Z-direction spacing of signal trace layers to reference plane layers (defined by total board thickness and number of board layers) is a defining factor in trace impedance. Trace width (defined by FPGA package ball pitch and PCB via manufacturing constraints) is another factor in trace impedance. A designer often has little control over trace impedance in area of the via array beneath the FPGA. When traces escape the via array, their width can change to the width of the target impedance (usually 50Ω single-ended).

Decoupling capacitor placement and discrete termination resistor placement are other areas of trade-off optimization. DFM constraints often define a keep-out area around the perimeter of the FPGA (device footprint) where no discrete components can be placed. The purpose of the keep-out area is to allow room for assembly and rework where necessary. For this reason, the area just outside the keep-out area is one where components compete for placement. It is up to the PCB designer to determine the high priority components. Decoupling capacitor placement constraints are described in [Chapter 1: Power Distribution System in UltraScale Devices](#). Termination resistor placement constraints must be determined through signal integrity simulation, using IBIS or SPICE.

Basic PDS Principles

The purpose of the PDS and the properties of its components are discussed in this section. The important aspects of capacitor placement, capacitor mounting, PCB geometry, and PCB stackup recommendations are also described.

Noise Limits

In the same way that devices in a system have a requirement for the amount of current consumed by the power system, there is also a requirement for the cleanliness of the power. This cleanliness requirement specifies a maximum amount of noise present on the power supply, often referred to as ripple voltage (V_{RIPPLE}). Most digital devices, including all UltraScale architecture-based devices, require that VCC supplies not fluctuate more than the specifications documented in the device data sheet.

The power consumed by a digital device varies over time and this variance occurs on all frequency scales, creating a need for a wide-band PDS to maintain voltage stability.

- Low-frequency variance of power consumption is usually the result of devices or large portions of devices being enabled or disabled or event occurrences like data traffic or processing. This variance occurs in time frames from milliseconds to days.
- High-frequency variance of power consumption is the result of individual switching events inside a device. This occurs on the scale of the clock frequency and the first few harmonics of the clock frequency up to about 5 GHz.

Because the voltage level of VCC for a device is fixed, changing power demands are manifested as changing current demand. The PDS must accommodate these variances of current draw with as little change as possible in the power-supply voltage.

When the current draw in a device changes, the PDS cannot respond to that change instantaneously. As a consequence, the voltage at the device changes for a brief period before the PDS responds. Two main causes for this PDS lag correspond to the two major PDS components: the voltage regulator and decoupling capacitors.

The first major component of the PDS is the voltage regulator. The voltage regulator observes its output voltage and adjusts the amount of current it is supplying to keep the output voltage constant. Most common voltage regulators make this adjustment in milliseconds to microseconds. Voltage regulators effectively maintain the output voltage for events at all frequencies from DC to a few hundred kHz, depending on the regulator (some are effective at regulating in the low MHz). For transient events that occur at frequencies above this range, there is a time lag before the voltage regulator responds to the new current demand level.


For example, if the device's current demand increases in a few hundred picoseconds, the voltage at the device sags by some amount until the voltage regulator can adjust to the new, higher level of required current. This lag can last from microseconds to milliseconds. A second component is needed to substitute for the regulator during this time, preventing the voltage from sagging.

This second major PDS component is the decoupling capacitor (also known as a bypass capacitor). The decoupling capacitor works as the device's local energy storage. The capacitor cannot provide DC power because it stores only a small amount of energy (voltage regulator provides DC power). This local energy storage should respond very quickly to changing current demands. The capacitors effectively maintain power-supply voltage at frequencies from hundreds of kHz to hundreds of MHz (in the milliseconds to nanoseconds range). Discrete decoupling capacitors are not useful for events occurring above or below this range.

For example, if current demand in the device increases in a few picoseconds, the voltage at the device sags by some amount until the capacitors can supply extra charge to the device. If current demand in the device maintains this new level for many milliseconds, the voltage-regulator circuit, operating in parallel with the decoupling capacitors, replaces the capacitors by changing its output to supply this new level of current.


The following figure shows the major PDS components: the voltage regulator, the decoupling capacitors, and the active device being powered (FPGA).

Figure 208: Simplified PDS Circuit



The following figure shows a simplified PDS circuit with all reactive components represented by a frequency-dependent resistor.

Figure 209: Further Simplified PDS Circuit



Role of Inductance

Inductance is the electrical property of conductors by which a changing magnetic field creates an electromagnetic force or voltage. This field opposes the change of the current in the current path. Inductance is the reason why capacitors cannot respond instantaneously to transient currents or to changes that occur at frequencies higher than their effective range.

Inductance can be thought of as the momentum of charge. Charge moving through a conductor represents some amount of current. If the level of current changes, the charge moves at a different rate. Because momentum (stored magnetic-field energy) is associated with this charge, some amount of time and energy is required to slow down or speed up the charge flow. The greater the inductance, the greater the resistance to change, and the longer the time required for the current level to change. A voltage develops across the inductance as this change occurs.

The PDS, made up of a regulator and multiple stages of decoupling capacitors, accommodates the device current demand and responds to current transients as quickly as necessary to maintain the voltage within the specified limits. When these current demands are not met, the voltage across the device's power supply changes. This is observed as noise. Inductance in the current path of the capacitors should be minimized, because it retards the ability of decoupling capacitors to quickly respond to changing current demands.

Inductances occur between the FPGA and capacitors and between the capacitors and the voltage regulator (see [Figure 209: Further Simplified PDS Circuit](#)). These inductances occur as parasitics in the capacitors and in all PCB current paths. It is important that each of these parasitics be minimized.

Capacitor Parasitic Inductance

The capacitance value is often considered to be a capacitor's most important characteristic. In power system applications, the parasitic inductance (ESL) has the same or greater importance. Capacitor package dimensions (body size) determine the amount of parasitic inductance. Physically small capacitors usually have lower parasitic inductance than physically large capacitors.

Requirements for choosing decoupling capacitors:

- For a specific capacitance value, choose the smallest package available.
- or -
- For a specific package size (essentially a fixed inductance value), choose the highest capacitance value available in that package.


Surface-mount chip capacitors are the smallest capacitors available and are a good choice for discrete decoupling capacitors:

- For values from 100 μF to very small values such as 0.01 μF , ceramic X7R or X5R type capacitors are usually used. These capacitors have a low parasitic inductance and a low ESR, with an acceptable temperature characteristic.
- For larger values, such as 47 μF to 1000 μF , tantalum capacitors are usually used. These capacitors have a low parasitic inductance and a medium ESR, giving them a low Q factor and consequently a very wide range of effective frequencies.

If tantalum capacitors are not available or cannot be used, low-ESR, low-inductance electrolytic capacitors can be used, provided they have comparable ESR and ESL values. Other new technologies with similar characteristics are also available (Os-Con, POSCAP, and Polymer-Electrolytic SMT). Electrolytic or tantalum capacitors are generally part of the voltage regulator design. They are considered as part of the voltage regulator control loop and should be co-designed with the voltage regulator or control firmware setting. The ESR of these capacitors sets the limit of switching regulator ripple and switching noise along with regulator FET and output inductor design.


A *real* capacitor of any type then not only has capacitance characteristics but also inductance and resistance characteristics. The following figure shows the parasitic model of a real capacitor. A real capacitor should be treated as an RLC circuit (a circuit consisting of a resistor (R), an inductor (L), and a capacitor (C), connected in series).

Figure 210: Parasitics of a Real, Non-Ideal Capacitor



The following figure shows a *real* capacitor's impedance characteristic. Overlaid on this plot are curves corresponding to the capacitor's capacitance and parasitic inductance (ESL). These two curves combine to form the RLC circuit's total impedance characteristic, softened or sharpened by the capacitor's ESR.


Figure 211: Contribution of Parasitics to Total Impedance Characteristics



As capacitive value is increased, the capacitive curve moves down and left. As parasitic inductance is decreased, the inductive curve moves down and right. Because parasitic inductance for capacitors in a specific package is fixed, the inductance curve for capacitors in a specific package remains fixed.

As different capacitor values are selected in the same package, the capacitive curve moves up and down against the fixed inductance curve, as shown in the following figure.

Figure 212: Effective Frequency Example



The low-frequency capacitor impedance can be reduced by increasing the value of the capacitor; the high-frequency impedance can be reduced by decreasing the inductance of the capacitor. While it might be possible to specify a higher capacitance value in the fixed package, it is not possible to lower the inductance of the capacitor (in the fixed package) without putting more capacitors in parallel. Using multiple capacitors in parallel divides the parasitic inductance, and at the same time, multiplies the capacitance value. This lowers both the high and low frequency impedance at the same time.

PCB Current Path Inductance

The parasitic inductance of current paths in the PCB have three distinct sources:


- Capacitor mounting
- PCB power and ground planes
- FPGA mounting

Capacitor Mounting Inductance

Capacitor mounting refers to the capacitor's solder lands on the PCB, the trace (if any) between the land and via, and the via. The vias, traces, and capacitor mounting pads of a 2-terminal capacitor contribute inductance between 300 pH to 4 nH depending on the specific pad routing, via structure and PCB stackup.

Because the current path's inductance is proportional to the loop area the current traverses, it is important to minimize this loop size. The loop consists of the path through one power plane, up through one via, through the connecting trace to the land, through the capacitor, through the other land and connecting trace, down through the other via, and into the other plane, as shown in the following figure.

Figure 213: Example Cutaway View of PCB with Capacitor Mounting



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

Some PCB manufacturing processes allow via-in-pad geometries, an option for reducing parasitic inductance. Using multiple vias per land is important with ultra-low inductance capacitors, such as reverse aspect ratio capacitors that place wide terminals on the sides of the capacitor body instead of the ends.

PCB layout engineers often try to squeeze more parts into a small area by sharing vias among multiple capacitors. *This technique should not be used under any circumstances.* PDS improvement is very small when a second capacitor is connected to an existing capacitor's vias. For a larger improvement, optimize the total number of capacitors and improve the mounting via inductance path.

The capacitor mounting (lands, traces, and vias) typically contributes about the same amount or more inductance than the capacitor's own parasitic self-inductance. If the mounting via structure is not optimized, the capacitor might not be effective at all to PDS.

Plane Inductance

Some inductance is associated with the PCB power and ground planes. The geometry of these planes determines their inductance.

Current spreads out as it flows from one point to another (due to a property similar to *skin effect*) in the power and ground planes. Inductance in planes can be described as *spreading inductance* and is specified in units of henries per square. The square is dimensionless; the shape of a section of a plane, not the size, determines the amount of inductance.

Spreading inductance acts like any other inductance and resists changes to the amount of current in a power plane (the conductor). The inductance retards the capacitor's ability to respond to a device's transient currents and should be reduced as much as possible. Because the designer's control over the X-Y shape of the plane can be limited, the only controllable factor is the spreading inductance value. This is determined by the thickness of the dielectric separating a power plane from its associated ground plane.

For high-frequency power distribution systems, power and ground planes work in pairs, with their inductances coexisting dependently with each other. The spacing between the power and ground planes determines the pair's spreading inductance. The closer the spacing (the thinner the dielectric), the lower the spreading inductance. Approximate values of spreading inductance for different thicknesses of FR4 dielectric are shown in the following table.

Table 150: Capacitance and Spreading Inductance Values for Different Thicknesses of FR4 Power-Ground Plane Sandwiches

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

Decreased spreading inductance corresponds to closer spacing of VCC and GND planes. When possible, place the VCC planes directly adjacent to the GND planes in the PCB stackup. Facing VCC and GND planes are sometimes referred to as *sandwiches*. While the use of VCC – GND sandwiches was not necessary in the past for previous technologies (lead frames, wire bond packages), the speeds involved and the sheer amount of power required for fast, dense devices often demand it.

However, because of the presence of substrate decoupling capacitors in UltraScale architecture-based devices, there is a limit to the amount of fast transient current demanded from PCB decoupling capacitors. This means that there is little benefit from dielectric thicknesses below 50 μ (2 mil). Dielectric thickness of 50 μ or 75 μ between V_{CC} and GND layers is sufficient for UltraScale architecture-based devices.

Besides offering a low-inductance current path, power-ground sandwiches also offer some high-frequency decoupling capacitance. As the plane area increases and as the separation between power and ground planes decreases, the value of this capacitance increases. Capacitance per square inch is shown in the preceding table. However, the amount of capacitance arising from these PCB power-ground plane pairs is generally inconsequential, given the substrate decoupling capacitors present in UltraScale architecture-based devices.

FPGA Mounting Inductance

The PCB solder lands and vias that connect the FPGA power pins (V_{CC} and GND) contribute an amount of parasitic inductance to the overall power circuit. For existing PCB technology, the solder land geometry and the dogbone geometry are mostly fixed, and parasitic inductance of these geometries does not vary. Via parasitic inductance is a function of the via length and the proximity of the opposing current paths to one another.

The relevant via length is the portion of the via that carries transient current between the FPGA solder land and the associated V_{CC} or GND plane. Any remaining via (between the power plane and the PCB backside) does not affect the parasitic inductance of the via (the shorter the via between the solder lands and the power plane, the smaller the parasitic inductance). Parasitic via inductance in the FPGA mounting is reduced by keeping the relevant V_{CC} and GND planes as close to the FPGA as possible (close to the top of the PCB stackup).

Device pinout arrangement determines the proximity of opposing current paths to one another. Inductance is associated with any two opposing currents (for example, current flowing in a V_{CC} and GND via pair). A high degree of mutual inductive coupling between the two opposing paths reduces the loop's total inductance. Therefore, when given a choice, V_{CC} and GND vias should be as close together as possible.

The via field under an FPGA has many V_{CC} and GND vias, and the total inductance is a function of the proximity of one via to another:

- For core V_{CC} supplies (V_{CCINT} and V_{CCAUX}), opposing current is between the V_{CC} and GND pins.
- For I/O V_{CC} supplies (V_{CCO}), opposing current is between any I/O and its return current path, whether carried by a V_{CCO} or GND pin.

To reduce parasitic inductance:

- Core V_{CC} pins such as V_{CCINT} and V_{CCAUX} are placed in a checkerboard arrangement in the pinout.
- V_{CCO} and GND pins are distributed among the I/O pins.

Every I/O pin in the Kintex UltraScale and Virtex UltraScale FPGA pinouts is adjacent to a return-current pin.

FPGA pinout arrangement determines the PCB via arrangement. The PCB designer cannot control the proximity of opposing current paths but has control over the trade-offs between the capacitor's mounting inductance and FPGA's mounting inductance:

- Both mounting inductances are reduced by placing power planes close to the PCB stackup's top half and placing the capacitors on the top surface (reducing the capacitor's via length).
- If power planes are placed in the PCB stackup's bottom half, the capacitors are recommended to be mounted on the PCB backside. In this case, FPGA mounting vias are already long, and making the capacitor vias long (by coming down from the top surface) is a bad practice. A better practice is to take advantage of the short distance between the underside of the PCB and the power plane of interest, mounting capacitors on the underside.

PCB Stackup and Layer Order

V_{CC} and ground plane placement in the PCB stackup (the layer order) has a significant impact on the parasitic inductances of power current paths. Layer order must be considered early in the design process:

- High-priority supplies should be placed closer to the FPGA (in the PCB stackup's top half)
- Low-priority supplies should be placed farther from the FPGA (in the PCB stackup's bottom half)

Power supplies with high transient current should have the associated V_{CC} planes close to the top surface (FPGA side) of the PCB stackup. This decreases the vertical distance (V_{CC} and GND via length) that currents travel before reaching the associated V_{CC} and GND planes. To reduce spreading inductance, every V_{CC} plane should have an adjacent GND plane in the PCB stackup. The skin effect causes high-frequency currents to couple tightly, and the GND plane adjacent to a specific V_{CC} plane tends to carry the majority of the current complementary to that in the V_{CC} plane. Thus, adjacent V_{CC} and GND planes are treated as a pair.

Not all V_{CC} and GND plane pairs reside in the PCB stackup's top half because manufacturing constraints typically require a symmetrical PCB stackup around the center (with respect to dielectric thicknesses and etched copper areas). The PCB designer chooses the priority of the V_{CC} and GND plane pairs: high priority pairs carry high transient currents and are placed high in the stackup, while low priority pairs carry lower transient currents (or can tolerate more noise) and are placed in the lower part of the stackup.

Capacitor Effective Frequency

Every capacitor has a narrow frequency band where it is most effective as a decoupling capacitor. This band is centered at the capacitor's self-resonant frequency F_{RSELF} . The effective frequency bands of some capacitors are wider than others. A capacitor's ESR determines the capacitor's quality (Q) factor, and the Q factor can determine the width of the effective frequency band:

- Tantalum capacitors generally have a very *wide* effective band.

- Ceramic chip capacitors with a lower ESR, generally have a very *narrow* effective frequency band.

An ideal capacitor only has a capacitive characteristic, whereas *real* non-ideal capacitors also have a parasitic inductance (ESL) and a parasitic resistance (ESR). These parasitics work in series to form an RLC circuit ([Figure 211: Contribution of Parasitics to Total Impedance Characteristics](#)). The RLC circuit's resonant frequency is the capacitor's self-resonant frequency.

To determine the RLC circuit's resonant frequency, use the following equation:

Equation 4: Resonant Frequency

$$F = \frac{1}{2\pi\sqrt{LC}}$$

Another method of determining the self-resonant frequency is to find the minimum point in the impedance curve of the equivalent RLC circuit. The impedance curve can be computed or generated in SPICE using a frequency sweep. See [Simulation Methods](#) for other ways to compute an impedance curve.

It is important to distinguish between the capacitor's self-resonant frequency and the mounted capacitor's effective resonant frequency when the capacitor is part of the system, F_{RIS} . This corresponds to the resonant frequency of the capacitor with its parasitic inductance, plus the inductance of the vias, planes, and connecting traces between the capacitor and the FPGA.

The capacitor's self-resonant frequency, F_{RSELF} , (capacitor data sheet value) is much higher than its effective mounted resonant frequency in the system, F_{RIS} . Because the mounted capacitor's performance is most important, the mounted resonant frequency is used when evaluating a capacitor as part of the greater PDS.

Mounted parasitic inductance is a combination of the capacitor's own parasitic inductance and the inductance of: PCB lands, connecting traces, vias, and power planes. Vias traverse a full PCB stackup to the device when capacitors are mounted on the PCB backside. For a board with a finished thickness of 1.524 mm (60 mils), these vias contribute approximately 300 pH to 1,500 pH, (the capacitor's mounting parasitic inductance, L_{MOUNT}) depending on the spacing between vias. Wider-spaced vias and vias in thicker boards have higher inductance.

To determine the capacitor's total parasitic inductance in the system, L_{IS} , the capacitor's parasitic inductance, L_{SELF} , is added to the mounting's parasitic inductance, L_{MOUNT} :

$$L_{IS} = L_{SELF} + L_{MOUNT}$$

For example, using X7R Ceramic Chip capacitor in 0402 body size:

$$C = 0.01_{\mu} F \text{ (selected by user)}$$

$$L_{SELF} = 0.9 \text{ nH (capacitor data sheet parameter)}$$

$F_{RSELF} = 53 \text{ MHz}$ (capacitor data sheet parameter)

$L_{MOUNT} = 0.8 \text{ nH}$ (based on PCB mounting geometry)

To determine the effective in-system parasitic inductance (L_{IS}), add the via parasitics:

$L_{IS} = L_{SELF} + L_{MOUNT} = 0.9 \text{ nH} + 0.8 \text{ nH}$

$L_{IS} = 1.7 \text{ nH}$

The values from the example are used to determine the mounted capacitor resonant frequency (F_{RIS}). Using the following equation:

Equation 5: F_{RIS}

$$F_{RIS} = \frac{1}{2\pi\sqrt{L_{IS}C}}$$

$$F_{RIS} = \frac{1}{2\pi\sqrt{(1.7 \times 10^{-9} \text{ H}) \cdot (0.01 \times 10^{-6} \text{ F})}} = 38 \times 10^6 \text{ Hz}$$

F_{RSELF} is 53 MHz, but F_{RIS} is lower at 38 MHz. The addition of mounting inductances shifts the effective-frequency band down.

A decoupling capacitor is most effective at the narrow-frequency band around its resonant frequency, and thus, the resonant frequency must be reviewed when choosing a capacitor collection to build up a decoupling network. This being said, capacitors can be effective at frequencies considerably higher and lower than their resonant frequency. Recall that capacitors of differing values in the same package share the same inductance curve. As shown in [Figure 212: Effective Frequency Example](#), for any given frequency along the inductive portion of the curve, the capacitors are equally effective.

Capacitor Anti-Resonance

One problem associated with combinations of capacitors in a PDS of an FPGA is anti-resonant spikes in the PDS aggregate impedance. The cause for these spikes is a bad combination of energy storage elements in the PDS (intrinsic capacitances, discrete capacitors, parasitic inductances, and power and ground planes).

Anti-resonance can arise between any two consecutive stages of a power distribution system, such as between the high-frequency PCB capacitors and the PCB plane capacitance. The inter-plane capacitance of the power and ground planes generally has a high-Q factor. If the high-frequency PCB capacitors also are high-Q, the crossover point between the high-frequency discrete capacitors and the plane capacitance might exhibit a high-impedance anti-resonance peak. If the FPGA has a high transient current demand at this frequency (as a stimulus), a large noise voltage can occur.

To correct this type of problem, the characteristics of the high-frequency discrete capacitors or the characteristics of the VCC and ground planes must be changed, or FPGA activity shifted to a different frequency away from the resonance.

Capacitor Placement Background

To perform the decoupling function, capacitors should be close to the device being decoupled.

Increased spacing between the FPGA and decoupling capacitor increases the current flow distance in the power and ground planes, and it often increases the current path's inductance between the device and the capacitor.

The inductance of this current path (the loop followed by current as it travels from the V_{CC} side of the capacitor to the V_{CC} pin[s] of the FPGA, and from the GND pin[s] of the FPGA to the GND side of the capacitor[s]), is proportional to the loop area. Inductance is decreased by decreasing the loop area.

Shortening the distance between the device and the decoupling capacitor reduces the inductance, resulting in a less impeded transient current flow. Because of typical PCB dimensions, this lateral plane travel tends to be less important than the phase relationship between the FPGA noise source and the mounted capacitor.

The phase relationship between the FPGA's noise source and the mounted capacitor determines the capacitor's effectiveness. For a capacitor to be effective in providing transient current at a certain frequency (for example, the capacitor's resonant frequency), the phase relationship, based on the distance traveled by the current from the FPGA to the capacitor, must be within a fraction of the corresponding period.

The capacitor's placement determines the length of the transmission line interconnect (in this case, the power and ground plane pair) between the capacitor and FPGA. The propagation delay of this interconnect is the key factor.

FPGA noise falls into certain frequency bands, and different sizes of decoupling capacitors take care of different frequency bands. Thus, capacitor placement requirements are determined by each capacitor's effective frequency.

When the FPGA initiates a current demand change, it causes a small local disturbance in the PDS voltage (a point in the power and ground planes). Before it can counteract this, the decoupling capacitor must first sense a voltage difference.

A finite time delay occurs between the start of the disturbance at the FPGA power pins and the point when the capacitor senses the disturbance.

$$\text{Time Delay} = \frac{\text{Distance from the FPGA power pins to the capacitor}}{\text{Signal propagation speed through FR4 dielectric}}$$

The dielectric is the substrate of the PCB where the power planes are embedded.

Another delay of the same duration occurs when the compensation current from the capacitor flows to the FPGA. For any transient current demand in the FPGA, a round-trip delay occurs before any relief is seen at the FPGA.

- Negligible energy is transferred to the FPGA with placement distances greater than one quarter of a demand frequency's wavelength.
- Energy transferred to the FPGA increases from 0% at one-quarter of a wavelength to 100% at zero distance.
- Energy is transferred efficiently from the capacitor to the FPGA when capacitor placement is at a fraction of a quarter wavelength of the FPGA power pins. This fraction should be small because the capacitor is also effective at some frequencies (shorter wavelengths) above its resonant frequency.

One-tenth of a quarter wavelength is a good target for most practical applications and leads to placing a capacitor within one-fortieth of a wavelength of the power pins it is decoupling. The wavelength corresponds to the capacitor's mounted resonant frequency, F_{RIS} .

When using large numbers of external termination resistors or passive power filtering for transceivers, priority should be given to these over the decoupling capacitors. Moving away from the device in concentric rings, the termination resistors and transceiver supply filtering should be closest to the device, followed by the smallest-value decoupling capacitors, then the larger-value decoupling capacitors.

V_{REF} Stabilization Capacitors

In V_{REF} supply stabilization, one capacitor per pin is placed as close as possible to the V_{REF} pin. The capacitors used are in the 0.022 μF – 0.47 μF range. The V_{REF} capacitor's primary function is to reduce the V_{REF} node impedance, which in turn reduces crosstalk coupling. Because no low-frequency energy is needed, larger capacitors are not necessary.

This only applies when Internal V_{REF} is not used. Internal V_{REF} is an UltraScale architecture feature wherein the reference voltage rail is generated internally. See *UltraScale Architecture SelectIO Resources User Guide (UG571)* for more details on Internal V_{REF} .

Power Supply Consolidation

Powering 1.8V V_{CCO} and V_{CCAUX} from a common PCB plane is allowed in Kintex UltraScale and Virtex UltraScale FPGA designs. However, careful consideration must be given to power supply noise—in particular, any noise on the V_{CCO} rail should not violate the recommended operating condition range for the V_{CCAUX} supply. See the data sheet for these requirements.

Unconnected V_{CC0} Pins

In some cases, one or more I/O/PSIO banks in an FPGA are not used (for example, when an FPGA has far more I/O pins than the design requires). In these cases, it might be desirable to leave the bank's associated V_{CC0} pins unconnected, as it can free up some PCB layout constraints (less voiding of power and ground planes from via antipads, less obstacles to signals entering and exiting the pinout array, more copper area available for other planelets in the otherwise used plane layer).

Leaving the V_{CC0} pins of unused I/O/PSIO banks floating reduces the level of ESD protection on these pins and the I/O pins in the bank. For maximum ESD protection in an unused bank, all V_{CC0} pins in that bank should be connected together to the same potential, whether that be a valid V_{CC0} voltage, or a floating plane. I/O/PSIO pins are also recommended to be connected to the same potential as V_{CC0} , or they can be left floating. If using SYSMON VAUX inputs to monitor voltages on the board, it is not recommended to ground V_{CC0} pins of any I/O/PSIO banks because this can interfere with ADC accuracy. In this case, tie the V_{CC0} of all unused I/O banks to an existing I/O/PSIO voltage plane that is applicable to the bank type of the unused bank.

Transmission Lines

The combination of a signal trace and a reference plane forms a transmission line. All I/O signals in a PCB system travel through transmission lines.

For single-ended I/O interfaces, both the signal trace and the reference plane are necessary to transmit a signal from one place to another on the PCB. For differential I/O interfaces, the transmission line is formed by the combination of two traces and a reference plane. While the presence of a reference plane is not strictly necessary in the case of differential signals, it is necessary for practical implementation of differential traces in PCBs.

Good signal integrity in a PCB system is dependent on having transmission lines with controlled impedance. Impedance is determined by the geometry of the traces and the dielectric constant of the material in the space around the signal trace and between the signal trace and the reference plane.

The dielectric constant of the material in the vicinity of the trace and reference plane is a property of the PCB laminate materials, and in the case of surface traces, a property of the air or fluid surrounding the board. PCB laminate is typically a variant of FR4, though it can also be an exotic material.

While the dielectric constant of the laminate varies from board to board, it is fairly constant within one board. Therefore, the relative impedance of transmission lines in a PCB is defined most strongly by the trace geometries and tolerances. Impedance variance can occur based on the presence or absence of glass in a local portion of the laminate weave, but this rarely poses issues except in high-speed (>6 Gb/s) interfaces.

Return Currents

An often neglected aspect of transmission lines and their signal integrity is return current. It is incorrect to assume that a signal trace by itself forms a transmission line. Currents flowing in a signal trace have an equal and opposite complementary current flowing in the reference plane beneath them. The relationship of the trace voltage and trace current to reference plane voltage and reference plane current defines the characteristic impedance of the transmission line formed by the trace and reference plane. While interruption of reference plane continuity beneath a trace is not as dramatic in effect as severing the signal trace, the performance of the transmission line and any devices sharing the reference plane is affected.

It is important to pay attention to reference plane continuity and return current paths. Interruptions of reference plane continuity, such as holes, slots, or isolation splits, cause significant impedance discontinuities in the signal traces. They can also be a significant source of crosstalk and contributor to Power Distribution System (PDS) noise. The importance of return current paths cannot be underestimated.

PCB Materials and Traces

The choice of transmission media, whether PCB materials or cable type, can have a large impact on system performance. Although any transmission medium is lossy at gigahertz frequencies, this chapter provides some guidelines on managing signal attenuation so as to obtain optimal performance for a given application.

How Fast is Fast?

Signal edges contain frequency components called harmonics. Each harmonic is a multiple of the signal frequency and has significant amplitude up to a frequency determined by the following equation:

$$f \approx \frac{0.35}{T}$$

Where:

f = Frequency in GHz

T = The smaller of signal rise (T_r) or fall (T_f) time in ns

Because dielectric losses in a PCB are frequency dependent, a bandwidth of concern must be determined to find the total loss of the PCB. Frequencies must start at the operation frequency and extend to the frequency in the preceding equation. For example, a 10 Gb/s signal with a 10 ps rise time has a bandwidth from 10 GHz to 35 GHz.

Dielectric Losses

The amount of signal energy lost into the dielectric is a function of the material's characteristics. Some parameters used to describe the material include relative permittivity ϵ_r (also known as the dielectric constant) and loss tangent. Skin effect is also a contributor to energy loss at line speeds in the gigahertz range.

Relative Permittivity

Relative permittivity is a measure of the effect of the dielectric on the capacitance of a conductor. The higher the relative permittivity, the slower a signal travels on a trace and the lower the impedance of a given trace geometry. A lower ϵ_r is almost always preferred.

Although the relative permittivity varies with frequency in all materials, FR4 exhibits wide variations in ϵ_r with frequency. Because ϵ_r affects impedance directly, FR4 traces can have a spread of impedance values with increasing frequency. While this spread can be insignificant at 1.125 Gb/s, it can be a concern at 10 Gb/s operation.

Loss Tangent

Loss tangent is a measure of how much electromagnetic energy is lost to the dielectric as it propagates down a transmission line. A lower loss tangent allows more energy to reach its destination with less signal attenuation.

As frequency increases, the magnitude of energy loss increases as well, causing the highest frequency harmonics in the signal edge to suffer the most attenuation. This appears as a degradation in the rise and fall times.

Skin Effect and Resistive Losses

The skin effect is the tendency for current to flow preferentially near the outer surface of a conductor. This is mainly due to the magnetic fields in higher frequency signals pushing current flow in the perpendicular direction towards the perimeter of the conductor.

As current density near the surface increases, the effective cross-sectional area through which current flows decreases. Resistance increases because the effective cross-sectional area of the conductor is now smaller. Because this skin effect is more pronounced as frequency increases, resistive losses increase with signaling rates.

Resistive losses have a similar effect on the signal as loss tangent. Rise and fall times increase due to the decreased amplitude of the higher harmonics, with the highest frequency harmonics being most affected. In the case of 10 Gb/s signals, even the fundamental frequency can be attenuated to some degree when using FR4.

For example, an 8 mil wide trace at 1 MHz has a resistance on the order of 0.06Ω /inch, while the same trace at 10 Gb/s has a resistance of just over 1Ω /inch. Given a 10 inch trace and 1.6V voltage swing, a voltage drop of 160 mV occurs from resistive losses of the fundamental frequency, not including the losses in the harmonics and dielectric loss.

Choosing the Substrate Material

The goal in material selection is to optimize both performance and cost for a particular application.

FR4, the most common PCB substrate material, provides good performance with careful system design. For long trace lengths or high signaling rates, a more expensive substrate material with lower dielectric loss must be used.

Substrates, such as Nelco, have lower dielectric loss and exhibit significantly less attenuation in the gigahertz range, thus increasing the maximum bandwidth of PCBs. At 3.125 Gb/s, the advantages of Nelco over FR4 are added voltage swing margin and longer trace lengths. At 10 Gb/s, a low-loss dielectric like Nelco is necessary unless high-speed traces are kept very short.

The choice of substrate material depends on the total length of the high-speed trace and also the signaling rate.

What-if analysis can be done in HSPICE simulation to evaluate various substrate materials. By varying the dielectric constant, loss tangent, and other parameters of the PCB substrate material. The impact on eye quality can be simulated to justify the use of higher cost materials. The impact of other parameters such as copper thickness can also be explored.

Traces

Trace Geometry

For any trace, its characteristic impedance is dependent on its stackup geometry as well as the trace geometry. In the case of differential traces, the inductive and capacitive coupling between the tightly coupled pair also determines the characteristic impedance of the traces.

The impedance of a trace is determined by its inductive and capacitive coupling to nearby conductors. For example, these conductors can be planes, vias, pads, connectors, and other traces, including the other closely coupled trace in a differential pair. The substrate properties, conductor properties, flux linkage area, and distance to a nearby conductor determine the amount of coupling and hence, the contribution to the final impedance.

2D field solvers are necessary in resolving these complex interactions and contribute to the calculation of the final impedance of the trace. They are also a useful tool to verify existing trace geometries.

Wider traces create a larger cross-sectional area for current to flow and reduce resistive losses in high-speed interfaces. Use the widest traces that space constraints allow. Because trace width tolerances are expressed in absolute terms, a wider trace also minimizes the percentage variation of the manufactured trace, resulting in tighter impedance control along the length of the transmission line.

Sometimes, striplines are preferred over microstrips because the reference planes on both sides of the trace provide radiation shielding. Microstrips are shielded on only one side (by the reference plane) because they run on the top-most or bottom-most layers, leaving the other side exposed to the environment.

For best results, the use of a 2D or 3D field solver is recommended for verification.

Trace Characteristic Impedance Design for High-Speed Transceivers

Because the transceivers use differential signaling, the most useful trace configurations are differential edge-coupled stripline and differential microstrip. While some backplanes use the differential broadside-coupled stripline configuration, it is not recommended for 10 Gb/s operation, because the P and N vias are asymmetrical and introduce common-mode non-idealities.

With few exceptions, 50Ω characteristic impedance (Z_0) is used for transmission lines in the channel. In general, when the width/spacing (W/S) ratio is greater than 0.4 (8 mil wide traces with 20 mil separation), coupling between the P and N signals affects the trace impedance. In this case, the differential traces must be designed to have an odd mode impedance (Z_{0O}) of 50Ω , resulting in a differential impedance (Z_{DIFF}) of 100Ω , because $Z_{DIFF} = 2 \times Z_{0O}$.

The same W/S ratio also must be less than 0.8, otherwise strong coupling between the traces requires narrower, lossier traces for a Z_{0O} of 50Ω . To clarify, with Z_{0O} at 50Ω , an even mode impedance (Z_{0E}) of 60Ω or below is desired.

The following figures show example cross sections of differential structures.

Figure 214: Differential Edge-Coupled Centered Stripline




Figure 215: Differential Edge-Coupled Offset Stripline




Figure 216: Centered Broadside-Coupled Stripline





Figure 217: Differential Microstrip



A good PCB manufacturer understands controlled impedance and allows fine adjustments for line widths to produce a Z_{00} of 50Ω . The PCB manufacturer also provides the parameters necessary for the specific PCB layout. Some parameters can be calculated or simulated from the guideline outlined in the example. Although $\pm 10\%$ tolerance on Z_{00} is typical and can provide adequate performance, the additional cost of a tighter tolerance results in better channel performance.

Trace Routing

High-speed serial differential traces are routed with the highest priority to ensure that the optimal path is available to these critical traces. This reduces the need for bends and vias and minimizes the potential for impedance transitions. Traces must be kept straight, short, and with as few layer changes as possible. The impact of vias is discussed in [Differential Vias](#).

Routing of high-speed traces must be avoided near other traces or other potential sources of noise. Traces on neighboring signal planes should run perpendicular to minimize crosstalk.

Striplines are to be used whenever possible, as are the uppermost and lowermost stripline layers to minimize via stubs. When the stackup is being planned, these layers must be placed as close to the top and bottom layers whenever possible.

Design constraints might require microstrips for the BGA exit path or from via to connector launch or SMT pads. In such cases, the microstrip trace must be kept as short as possible.

Mitered 45-degree bends are recommended (as opposed to 90-degree bends). At a 90-degree bend, the effective width of the trace changes, causing an impedance discontinuity due to the capacitive coupling of the additional conductor area to the reference plane.

The two traces of a differential pair must be length-matched to eliminate skew. Skew creates mismatches in the common mode and reduces the differential voltage swing as a result.

Plane Splits

Ground planes should be used as reference planes for signals, as opposed to noisier power planes. Each reference plane should be contiguous for the length of the trace, because routing over plane splits creates an impedance discontinuity. In this case, the impedance of the trace changes because its coupling to the reference plane is changed abruptly at the plane split.

Return Currents

Routing over plane splits also creates issues with the return current. High-speed signals travel near the surface of the trace due to the skin effect mentioned in [Dielectric Losses](#). Meanwhile, the return current also travels near the surface of the tightly coupled reference plane.

Because of the tight coupling, the return current has the tendency to travel close to the original signal-carrying trace. At the plane split, the return current can no longer follow the same path parallel to the trace, but must instead find an alternative route.

A plane split causes a suboptimal current return path and increases the current loop area, thereby increasing the inductance of the trace at the plane split, changing the impedance of the trace.

Simulating Lossy Transmission Lines

Due to the different modeling implementations used by various circuit simulators (frequency-domain versus time-domain techniques), it is important to check that the models accurately reflect actual losses. One method is to compare the models against known published configurations.

Cable

Cables are controlled-impedance transmission lines due to the constant physical dimensions of conductor and dielectric along the length of the cable. The highest quality cable shows little variation in these dimensions and also has a wide bandwidth with low loss at high frequencies.

Connectors

The connectors attached to cables should exhibit low parasitic inductance, low-parasitic capacitance, and low crosstalk for high bandwidth operation.

Skew Between Conductors

When selecting a cable, look for a specification of the skew between the conductors in a cable. If the conductors are not length matched, the skew appears in the common mode and directly reduces the eye height.

Simulation Methods

Simulation methods, ranging from very simple to very complex, exist to predict the PDS characteristics. An accurate simulation result is difficult to achieve without using a fairly sophisticated simulator and taking a significant amount of time.

Basic lumped RLC simulation is one of the simplest simulation methods. Though it does not account for the distributed behavior of a PDS, it is a useful tool for selecting and verifying that combinations of decoupling capacitor values will not lead to large anti-resonances. Lumped RLC simulation is a good method for establishing equivalence of decoupling networks.

Lumped RLC simulation is performed either in a version of SPICE or other circuit simulator, or by using a mathematical tool like Mathcad or Microsoft Excel. Istvan Novak publishes a free Excel spreadsheet for lumped RLC simulation (among other useful tools for PDS simulation) on his website under Tool Download:

<http://www.electrical-integrity.com>

The following table also lists a few EDA tool vendors for PDS design and simulation. These tools span a wide range of sophistication levels.

Table 151: EDA Tools for PDS Design and Simulation

Tool	Vendor	Website URL
ADS	Keysight Technologies	http://www.keysight.com
SIwave, HFSS	Ansys	http://www.ansys.com
Sigrity	Cadence	http://www.cadence.com
Hyperlynx PI	Mentor	http://www.mentor.com

PDS Measurements

Measurements can be used to determine whether a PDS is adequate. PDS noise measurements are a unique task, and many specialized techniques have been developed. This section describes the noise magnitude and noise spectrum measurements.

Noise Magnitude Measurement

Noise measurement must be performed with a high-bandwidth oscilloscope (minimum 3 GHz oscilloscope and 1.5 GHz probe or direct coaxial connection) on a design running realistic test patterns. The measurement is taken at the device's power pins (referred to as a *spyhole measurement*) or at an unused I/O driven High or Low.

V_{CCINT} and V_{CCAUX} can only be measured at the PCB backside vias. V_{CCO} can also be measured this way, but more accurate results are obtained by measuring static (fixed logic level) signals at unused I/Os in the bank of interest.

When making the noise measurement on the PCB backside, the via parasitics in the path between the measuring point and FPGA must be considered. Any voltage drop occurring in this path is not accounted for in the oscilloscope measurement.

PCB backside via measurements also have a potential problem: decoupling capacitors are often mounted directly underneath the device, meaning the capacitor lands connect directly to the V_{CC} and GND vias with surface traces. These capacitors confuse the measurement by acting like a short circuit for the high-frequency AC current. To make sure the measurements are not shorted by the capacitors, remove the capacitor at the measurement site (keep all others to reflect the real system behavior).

When measuring V_{CCO} noise, the measurement can be taken at an I/O pin configured as a driver to logic 1 or logic 0. In most cases, the same I/O standard should be used for this measurement as for the other signals in the bank. Measuring a static logic 0 shows the crosstalk (via field, PCB routing, package routing) induced on the victim. Measuring a static logic 1 shows all the same crosstalk components as well as the noise present on the V_{CCO} net for the I/O bank. By subtracting (coherently in time) the noise measured on static logic 0 from the noise measured on static logic 1, the noise on V_{CCO} at the die can be viewed. For an accurate result, the static logic 0 and static logic 1 noise must be measured at the same I/O location. This means storing the time-domain waveform information from both logic states and performing the subtraction operation on the two waveforms in a post-process math computation tool such as MATLAB® or Excel.

Oscilloscope Measurement Methods


There are two basic ways of using the oscilloscope to view power system noise, each for a different purpose. The first surveys all possible noise events, while the second is useful for focusing on individual noise sources.

- Place the oscilloscope in infinite persistence mode to acquire all noise over a long time period (many seconds or minutes). If the design operates in many different modes, using different resources in different amounts, these various conditions and modes should be in operation while the oscilloscope is acquiring the noise measurement.
- Place the oscilloscope in averaging mode and trigger on a known aggressor event. This can show the amount of noise correlated with the aggressor event (any events asynchronous to the aggressor are removed through averaging).

Power system noise measurements should be made at a few different FPGA locations to ensure that any local noise phenomena are captured.

The following figure shows an averaged noise measurement taken at the V_{CCO} pins of a sample design. In this case, the trigger was the clock for an I/O bus interface sending a 1-0-1-0 pattern at 250 Mb/s.


Figure 218: Averaged Measurement of V_{CC0} Supply with Multiple I/O Sending Patterns at 250 Mb/s



UG583_c3_11_112113

The following figure shows an infinite persistence noise measurement of the same design with a wider variety of I/O activity. Because the infinite persistence measurement catches *all* noise events over a long period, both correlated and non-correlated with the primary aggressor, all power system excursions are shown.

Figure 219: Infinite Persistence Measurement of Same Supply



UG583_c3_12_112113

The measurement shown in the preceding figures represents the peak-to-peak noise. If the peak-to-peak noise is outside the specified acceptable voltage range (data sheet value, $V_{CC} \pm 5\%$), the decoupling network is inadequate or a problem exists in the PCB layout.

Noise Spectrum Measurements

Having the necessary information to improve the decoupling network requires additional measurements. To determine the frequencies where the noise resides, noise power spectrum measurement is necessary. A spectrum analyzer or a high-bandwidth oscilloscope coupled with FFT math functionality can accomplish this.

The FFT math function can be built into the oscilloscope, however, many of these functions do not have resolution sufficient to give a clear picture of the noise spectrum. Alternatively, a long sequence of time-domain data can be captured from an oscilloscope and converted to frequency domain using MATLAB or other post-processing software supporting FFT. This method has the advantage of showing as much resolution as you are willing to process. If neither math capacity is available, the noise frequency content can be approximated by visually examining the time-domain waveform and estimating the individual periodicities present in the noise.

A spectrum analyzer is a frequency-domain instrument, showing the frequency content of a voltage signal at its inputs. Using a spectrum analyzer, you see the exact frequencies where the PDS is inadequate.

Excessive noise at a certain frequency indicates a frequency where the PDS impedance is too high for the device's transient current demands. Using this information, the designer can modify the PDS to accommodate the transient current at the specific frequency. This is accomplished by either adding capacitors with effective frequencies close to the noise frequency or otherwise lowering the PDS impedance at the critical frequency.


The noise spectrum measurement should be taken in the same manner as the peak-to-peak noise measurement, directly underneath the device, or at a static I/O driven High or Low. A spectrum analyzer takes its measurements using a 50Ω cable instead of an active probe.

- A good method attaches the measurement cable through a coaxial connector tapped into the power and ground planes close to the device. This is *not* available in most cases.
- Another method attaches the measurement cable at the lands of a decoupling capacitor in the vicinity of the device that has been removed. The cable's center conductor and shield are soldered directly to the capacitor lands. Alternatively, a probe station with 50Ω RF probes can be used to touch the decoupling capacitor lands.

To protect the spectrum analyzer's sensitive front-end circuitry, add a DC blocking capacitor or attenuator in line. This isolates the spectrum analyzer from the device supply voltage.

The following figure is an example of a noise spectrum measurement of the V_{CC0} power-supply noise, with multiple I/O sending patterns at 100 MHz.

Figure 220: Screenshot of Spectrum Analyzer Measurement of V_{CC0}



UG583_03_13_112113

Optimum Decoupling Network Design

If a highly optimized PDS is needed, measurements and simulations of a prototype system can inform the PDS design. Using knowledge of the noise spectrum generated by the prototype system along with knowledge of the system’s power system impedance, the unique transient current of the design can be determined and accommodated.

To measure the noise spectrum of the design under operating conditions, use either a spectrum analyzer or an oscilloscope with FFT. The power system impedance can be determined either through direct measurement or simulation, or a combination of these two as there are often many variables and unknowns.

Both the noise spectrum and the impedance are functions of frequency. By examining the quotient of these per frequency point, transient current as a function of frequency is computed (see the following equation):

$$I(f) = \frac{V(f) \text{ From Spectrum Analyzer}}{Z(f) \text{ From Network Analyzer}}$$

Using the data sheet's maximum voltage ripple value, the impedance value needed at all frequencies can be determined. This yields a target impedance as a function of frequency. A specially designed capacitor network can accommodate the specific design's transient current.

Troubleshooting

In some cases the proper design work is done up-front, but noise problems still exist. This next section describes possible issues and suggested resolution methods.

Possibility 1: Excessive Noise from Other Devices on the PCB

Sometimes ground and/or power planes are shared among many devices, and noise from an inadequately decoupled device affects the PDS at other devices. Common causes of this noise are:

- RAM interfaces with inherently high-transient current demands resulting either from temporary periodic contention or high-current drivers
- Large ASICs

When unacceptable amounts of noise are measured locally at these devices, the local PDS and the component decoupling networks should be analyzed.

Possibility 2: Parasitic Inductance of Planes, Vias, or Connecting Traces

Sometimes the decoupling network capacitance is adequate, but there is too much inductance in the path from the capacitors to the FPGA.

Possible causes are:

- Wrong decoupling capacitor connecting-trace geometry or solder-land geometry
- The path from the capacitors to the FPGA is too long
 - and/or -
- A current path in the power vias traverses an exceptionally thick PCB stackup

For inadequate connecting trace geometry and capacitor land geometry, review the loop inductance of the current path. If the vias for a decoupling capacitor are spaced a few millimeters from the capacitor solder lands on the board, the current loop area is greater than necessary.

To reduce the current loop area, vias should be placed directly against capacitor solder lands. *Never* connect vias to the lands with a section of trace.

Other improvements of geometry are via-in-pad (via under the solder land), not shown, and via-beside-pad (vias straddle the lands instead of being placed at the ends of the lands). Double vias also improve connecting trace geometry and capacitor land geometry.

Exceptionally thick boards (> 3.2 mm or 127 mils) have vias with higher parasitic inductance.

To reduce the parasitic inductance, move critical V_{CC}/GND plane sandwiches close to the top surface where the FPGA is located, and place the capacitors on the top surface where the FPGA is located.

Possibility 3: I/O Signals in PCB are Stronger Than Necessary

If noise in the V_{CCO} PDS is still too high after refining the PDS, the I/O interface slew rate and/or drive strength can be reduced. This applies to both outputs from the FPGA and inputs to the FPGA. In severe cases, excessive overshoot on inputs to the FPGA can reverse-bias the IOB clamp diodes, injecting current into the V_{CCO} PDS.

If large amounts of noise are present on V_{CCO} , the drive strength of these interfaces should be decreased, or different termination should be used (on input or output paths).

Possibility 4: I/O Signal Return Current Traveling in Sub-Optimal Paths

I/O signal return currents can also cause excessive noise in the PDS. For every signal transmitted by a device into the PCB (and eventually into another device), there is an equal and opposite current flowing from the PCB into the device's power/ground system. If a low-impedance return current path is not available, a less optimal, higher impedance path is used. When I/O signal return currents flow over a less optimal path, voltage changes are induced in the PDS, and the signal can be corrupted by crosstalk. This can be improved by ensuring every signal has a closely spaced and fully intact return path.

Methods to correct a sub-optimal return current path:

- Restrict signals to fewer routing layers with verified continuous return current paths.
- Provide low-impedance paths for AC currents to travel between reference planes (high-frequency decoupling capacitors at PCB locations where layer transitions occur).

Design of Transitions for High-Speed Signals

Each transition in the channel must be designed to minimize any negative impact on the link performance. This chapter addresses the interface at either end of a transmission line.

Transmission lines have defined and controlled characteristic impedance along their length. However, the three-dimensional structures that they interface do not have easily defined or constant impedance along the signal path. Software tools such as 3D field solvers are necessary for computing the impedance that a 10 Gb/s signal sees as it passes through these structures, while 2D field solvers are sufficient for computing transmission line characteristic impedance.

PCB designers can use the analyses and examples in this chapter to assist the design of such a channel. Cases not covered in this chapter might need further simulation and analysis.

Excess Capacitance and Inductance

Most differential transitions are overly capacitive. The P and N paths couple to each other, increasing capacitance. Many transitions have a frequency response identical to that of a lumped capacitor over a wide frequency band.

By design, adding inductance cancels this excess capacitance in many cases except when impacted by density concerns and physical limitations. While techniques such as blind vias, solder balls on a larger pitch, and very small via pads reduce capacitance, they are not always feasible in a design.

Time domain reflectometry (TDR) techniques, either through simulation or measurement, allow the designer to identify excess capacitance or excess inductance in a transition.

Time Domain Reflectometry

To make TDR measurements, a step input is applied to the interconnect. The location and magnitude of the excess capacitance or inductance that the voltage step experiences as it traverses the interconnect can be determined through observing the reflected signal.

A shunt capacitance (see the following figure) causes a momentary dip in the impedance, while a series inductance (see [Figure 222: TDR Signature of Series Inductance](#)) causes an impedance discontinuity in the opposite direction. T_d is the propagation delay through the first transmission line segment on the left. The reflected wave due to the impedance discontinuity takes $2 * T_d$ to return to the TDR port. If the signal propagation speed through the transmission line is known, the location of the excess capacitance or inductance along the channel can be calculated.

Figure 221: TDR Signature of Shunt Capacitance





Figure 222: TDR Signature of Series Inductance



The magnitude of this excess capacitance (C) or inductance (L) can also be extracted from the TDR waveform by integrating the normalized area of the transition's TDR response. The respective equations for capacitance and inductance are:

Equation 6: Capacitance


$$C = -\frac{2}{Z_0} \int_{t_1}^{t_2} \frac{V_{tdr}(t) - v_{step}}{V_{step}} dt$$

Equation 7: Inductance

$$L = 2Z_0 \int_{t_1}^{t_2} \frac{V_{tdr}(t) - v_{step}}{V_{step}} dt$$

The following figure shows the integration of the normalized TDR area.

Figure 223: Integration of Normalized TDR Area



The results using these equations are not sensitive to rise time variation and are valid for simulated TDR measurements provided that the leading and trailing transmission lines are very close to 50Ω . However, for actual measurements, accuracy is very dependent on Z_0 .


BGA Package

Each signal path within the BGA package is carefully designed to optimize signal integrity. Traces supporting single-ended I/O are nominally designed for 50Ω trace impedance. Traces supporting high-speed SerDes I/O are designed for nominally 100Ω differential impedance. Special care is taken in the design of signal paths to optimize discontinuities such as solder balls and substrate vias to minimize their effect on signal integrity. A 3D full-wave electromagnetic solver and a vector network analyzer are used to model and measure package performance.

SMT Pads

For applications that require AC coupling between transmitter and receiver, SMT pads are introduced in the channel to allow coupling capacitors to be mounted. Standard SMT pads have excess capacitance due to plate capacitance to a nearby reference plane. In the following figure example, a 5 mil trace with a Z_0 of 50Ω transitions to an 0402 SMT pad that is 28 mils wide, all over 3 mils of FR4.

Figure 224: 2D Field-Solver Analysis of 5 Mil Trace and 28 Mil Pad



Using a 2D field solver on these dimensions yields a Z_0 of 50Ω for the 5 mil trace. The Z_0 for the 0402 pad is 16Ω because the pad has too much capacitance and too little inductance, resulting in an impedance of less than 50Ω . Performance of this transition can be optimized in one of two ways.

The first method makes the trace the same width as the pad and moves the ground plane deeper into the stackup to maintain the Z_0 of the transition at 50Ω . This method does not require any special analysis, but there might be some error due to the fringing capacitance of the SMT capacitor body. Trace density is limited because traces are now 28 mils wide.

The second method, shown in the following figure, clears the ground plane underneath the pad, which removes much of the excess capacitance caused by the plate capacitance between the pad and the ground plane. This technique allows for greater trace density than the first method, but requires 3D field-solver analysis or measurement along with several board iterations to get the desired performance.


Figure 225: Transition Optimization



The 2D field-solver example shows that close to 50Ω can be achieved if the ground plane under the pad footprint is cleared out. A 3D field solver is then used to verify this result to a greater degree of accuracy.


The following figure shows the ground plane cleared away exactly as it was for the 2D simulation. Using frequency domain analysis within HFSS, there is a 20 dB (10x) improvement in return loss using this technique.

Figure 226: Ansoft HFSS Model of Pad Clear-Out




The following figure shows the return loss comparison between 0402 pad structures with linear scale.

Figure 227: Return Loss Comparison Between 0402 Pad Structures



The approximately -40 dB/decade slope in the following figure shows good fit to the frequency response of a lumped capacitor.

Figure 228: Return Loss Comparison Between 0402 Pad Structures on Log (Frequency) Scale



Next, using simulated measurements on the same transition modeled in HFSS, the time-domain performance of this transition can be measured by doing a TDR on the S-parameter results from the earlier frequency domain analysis.

In the following figures, the red curve with the large capacitive dip corresponds to the SMT pad without the ground plane cleared from underneath. The blue curve shows that clearing out the ground plane removes much of the excess capacitance. This improvement can be quantified using [Equation 6: Capacitance](#) and [Equation 7: Inductance](#).

Figure 229: TDR Results Comparing 0402 Pad Structures





Figure 230: TDR Results Comparing 0402 Pad Structures



As shown from the following figures, clearing the ground plane under SMT pads yields a significant improvement in the performance of an SMT pad transition. Excess capacitance is reduced by 15x, and return loss is improved by 20 dB.

Figure 231: 840 fF Excess Capacitance with Ground Plane Intact





Figure 232: 57 fF Excess Capacitance with Ground Plane Intact




Differential Vias

The most common transition is the differential via where the signal pair must transition from an upper stripline layer or top microstrip to a lower stripline layer or bottom microstrip.

The following figure shows a Ground-Signal-Signal-Ground (GSSG) type differential via. Ground vias are connected to each ground plane in the stackup, while signal layers only contain pads for the entry and exit layers.

Figure 233: Differential Via Design Example

Via Diameter = 12 mils (0.012 inches)
Pad Diameter = 22 mils
Annular Ring = 5 mils
GSSG Via Pitch = 40 mils
Oblong Antipads = ~55 mils x 95 mils,
aligned with ground pads



UG583_c4_13_112113


A key advantage of a GSSG via is that it allows for the signal's return current to flow in the ground via near the corresponding signal via, reducing excess inductance. The signal path is also symmetrical between the P and N halves of the differential signal, which is critical in controlling common-mode artifacts due to P/N imbalance.

The larger oblong antipads reduce excess fringing capacitance between the via body and the surrounding planes edges. Unused pads are also removed.

A good starting point is to use the dimensions shown in the preceding figure as an example differential via design for an 80 mil board. To accommodate density constraints or the lack thereof, the dimensions can be scaled accordingly to preserve the ratios of each dimension relative to the others. Such scaling preserves the impedance performance of the differential via while allowing variation in overall size to better suit specific applications. These final dimensions are limited by manufacturability and density constraints.


While the via length can be varied by a small amount to suit boards that are thicker or thinner than the 80 mil example, changing the ratio of the via length relative to other dimensions affects the via's impedance. For this and other configurations of differential vias, it is best to simulate a model using 3D field-solver tools to ensure that performance targets are met.

Figure 234: Differential GSSG Via in 16-Layer PCB from Pins L11 and L6



As a general rule, the P and N paths need to be kept at equal lengths through a transition. Where possible, via stub length should be kept to a minimum by traversing the signal through the entire length of the vias. The analysis shown in the following figure compares the S-parameter return loss for common-mode (SCC11) and differential (SDD11) responses.

Figure 235: Simulated Return Loss Comparing Differential and Common-Mode Losses for L11 and L6 GSSG Vias



From the graph in the preceding figure, the common-mode response is 20 dB worse in terms of return loss. The much worse common-mode response relative to the differential response is the reason why it is a good idea to reduce P/N skew as much as possible before entering a transition. The 60/40 rule of thumb is 40 dB of return loss at 1 GHz, which implies 60 fF of excess capacitance. Because excess capacitance is a single pole response, simple extrapolation rules can be used. For example, a shift to 34 dB return loss doubles the excess capacitance. Due to the excellent performance characteristics of GSSG vias, even long via stubs only double the differential via's capacitance at the most.

P/N Crossover Vias

Some transceivers offer the ability to independently switch the polarity of the transmit and receive signal pairs. This functionality eliminates the need to cross over the P/N signals at the board level, which in turn significantly enhances signal integrity. If possible, P/N crossover vias are to be avoided and the polarity switch of the transceiver should be used.

SMA Connectors

Well-designed SMA connectors can reduce debugging time and allow a high-performance channel to be designed correctly on the first pass. SMA connectors that perform well at 10 Gb/s need to be simulated, designed, and manufactured to meet this performance target. Vendors can also offer design services that ensure that the connector works well on a specific board. Assembly guidelines are crucial in ensuring that the process of mating the connector to the board is well-controlled to give the specified performance.

AMD uses precision SMA connectors from Rosenberger and other precision connector manufacturers because of their excellent performance and because of the points listed in the previous paragraph.

Backplane Connectors

There are numerous signal integrity issues associated with backplane connectors including:

- P/N signal skew
- Crosstalk
- Stubs due to connector pins


Some connector manufacturers offer not only S parameters, models, and layout guidelines for their connectors but also design support, seminars, and tutorials.

Microstrip/Stripline Bends

A bend in a PCB trace is a transition. When routing differential traces through a 90° corner, the outer trace is longer than the inner trace, which introduces P/N imbalance. Even within a single trace, signal current has the tendency to hug the inside track of a corner, further reducing the actual delay through a bend.

To minimize skew between the P and N paths, 90° turns in microstrips or striplines are routed as two 45° bends to give mitered corners. The addition of a jog-out also allows the trace lengths to be matched. The following figure shows example bends in traces.

Figure 236: Example Design for 90 Degree Bends in Traces



UG583_c4_16_112113

Turns add capacitance because the trace at a 90° corner is 41% wider. That difference is reduced to 8% with a 45° turn. The addition of plane cutouts to a depth of 30 mils act to reduce this amount of excess capacitance. The trace was not widened to maintain 50Ω with the plane cutouts in place.

When this mitered bend is simulated with the jog-out and plane cutouts, excess capacitance is reduced and P/N length and phase matching is improved. Without jog-outs, the P/N length mismatch is 16 mils. Given FR4 material, the 16 mil difference translates to a phase mismatch of 4.8° at 5 GHz, or 2.68 ps (0.0268 UI) at 10 Gb/s.

The following figures show that phase mismatch is reduced to 0.75° with jog-outs and 0.3° with jog-outs and plane cutouts. The combination of jog-outs and plane cutouts yields simulation results that show the excess capacitance of the structure is reduced to 65 fF.

Designers are tempted to widen lines to compensate for the characteristic impedance increase as the lines are separated and couple less strongly. However, even without widening the lines, the combined capacitance of the corners and jog-outs is still overly capacitive, and therefore the uncoupled section of the jog-out must not be widened.

Figure 237: Simulated TDR of 45 Degree Bends with Jog-Outs




Figure 238: Simulated Return Loss of 45 Degree Bends with Jog-Outs






Figure 239: Simulated Phase Response of 45 Degree Bends with Jog-Outs



For wide traces, curved routing can also be helpful as shown in the following figure.

Figure 240: Measured TDR of 45 Degree Bends with and without Jog-Outs



UG583_c4_20_112113

Maximum Current Draw for V_{CCINT} in UltraScale+ Devices

AMD UltraScale™ and AMD UltraScale+™ devices are designed to work over a wide temperature range, and the temperature of the device and package is directly related to the amount of current draw. This is especially relevant on the V_{CCINT} power rail, which is the rail that would typically require the most power and thus current draw of all the power rails in the device. For large AMD Virtex™ UltraScale+™ FPGA designs in very high-speed applications, it might be possible to exceed the maximum current that the package is rated for due to the effect of electromigration.

The following tables are a guide to the maximum possible sustained current draw on the V_{CCINT} rail for UltraScale+ devices to guarantee reliable operation for ten years at a constant 100°C junction temperature for commercial (XC) devices, and for three years at a constant 125°C junction temperature for automotive (XA) and defense (XQ) devices.

Table 152: Maximum V_{CCINT} Current for Artix UltraScale+ FPGAs

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
XCAU7P-FCVA289	5.48
XCAU10P-UBVA368	4.4
XCAU10P-SBVB484	15
XAAU10P-SBVB484	10
XCAU10P-FFVB676	33
XAAU10P-FFVB676	23
XCAU15P-UBVA368	6.4
XCAU15P-SBVB484	15
XAAU15P-SBVB484	10
XCAU15P-FFVB676	33
XAAU15P-FFVB676	23
XCAU20P-FFVB676	52
XCAU20P-SFVB784	57
XCAU25P-FFVB676	52
XCAU25P-SFVB784	57

Table 153: Maximum V_{CCINT} Current for Kintex UltraScale+ FPGAs

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
XCKU3P-FFVA676	43
XCKU3P-FFVB676	52
XCKU3P-SFVB784	57
XCKU3P-FFVD900	107
XCKU5P-FFVA676	43
XCKU5P-FFVB676	52
XQKU5P-FFRB676	43
XCKU5P-SFVB784	57
XQKU5P-SFRB784	57
XCKU5P-FFVD900	107
XCKU9P-FFVE900	114
XCKU11P-FFVD900	64
XCKU11P-FFVA1156	41
XCKU11P-FFVE1517	60
XCKU13P-FFVE900	115
XCKU15P-FFVA1156	50
XQKU15P-FFRA1156	50
XCKU15P-FFVE1517	73
XQKU15P-FFRE1517	73
XCKU15P-FFVA1760	77
XCKU15P-FFVE1760	76
XCKU19P-FSVJ1760	250
XCKU19P-FFVJ1760	250
XCKU19P-FFVB2104	86

Table 154: Maximum V_{CCINT} Current for AMD Virtex™ UltraScale+™ FPGAs

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
XCVU3P-FFVC1517	40
XQVU3P-FFRC1517	40
XCVU5P-FLVA2104	120
XCVU5P-FLVB2104	93
XCVU5P-FLVC2104	126
XCVU7P-FLVA2104	120
XQVU7P-FLRA2104	120
XCVU7P-FLVB2104	93

Table 154: Maximum V_{CCINT} Current for AMD Virtex™ UltraScale+™ FPGAs (cont'd)

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
XQVU7P-FLRB2104	93
XCVU7P-FLVC2104	126
XCVU9P-FLGA2104	99
XCVU9P-FLGB2104	93
XQVU9P-FLQA2104	75
XCVU9P-FLGC2104	137
XCVU9P-FSGD2104	217
XQVU9P-FLQB2104	54
XQVU9P-FSQD2104	124
XCVU9P-FLGA2577	174
XCVU11P-FLGF1924	102
XCVU11P-FLGB2104	118
XCVU11P-FLGC2104	141
XQVU11P-FLRC2104	141
XCVU11P-FSGD2104	200
XCVU11P-FLGA2577	163
XCVU13P-FHGA2104	94
XCVU13P-FHGB2104	114
XCVU13P-FHGC2104	129
XQVU13P-FHQA2104	53
XQVU13P-FHQB2104	81
XCVU13P-FIGD2104	200
XQVU13P-FIQD2104	110
XCVU13P-FLGA2577	162
XCVU13P-FSGA2577	158
XCVU15P-FSVA3824	275
XCVU15P-FSVB3824	314
XCVU19P-FSVA3824	275
XCVU19P-FSVB3824	314
58G PAM4 Devices	
XCVU23P-VSVA1365	130
XCVU23P-FSVJ1760	250
XCVU27P-FIGD2104	192
XCVU27P-FSGA2577	158
XCVU29P-FIGD2104	192
XCVU29P-FSGA2577	158

Table 154: Maximum V_{CCINT} Current for AMD Virtex™ UltraScale+™ FPGAs (cont'd)

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
HBM (4 GB) Devices	
XCVU31P-FSVH1924	68
HBM (8 GB) Devices	
XCVU33P-FSVH2104	68
XCVU35P-FSVH2104	135
XCVU35P-FSVH2892	135
XCVU37P-FSVH2892	200
XQVU37P-FSQH2892	116
HBM (16 GB) Devices	
XCVU45P-FSVH2104	135
XCVU45P-FSVH2892	135
XCVU47P-FSVH2892	200
XCVU57P-FSVK2892	218

 Table 155: Maximum V_{CCINT} Current for Zynq UltraScale+ MPSoCs

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
CG Devices	
XCZU1CG-SBVA484	7
XCZU1CG-UBVA494	3.16
XCZU1CG-SFVA625	11
XCZU1CG-SFVC784	7
XCZU2CG-SBVA484	9
XCZU2CG-UBVA530	6
XCZU2CG-SFVA625	10
XCZU2CG-SFVC784	7
XCZU3CG-SBVA484	9
XCZU3CG-UBVA530	6
XCZU3CG-SFVA625	10
XCZU3CG-SFVC784	7
XCZU3TCG-SFVC784	14
XCZU3TCG-SFVD784	23
XCZU4CG-SFVC784	11
XCZU4CG-FBVB900	52
XCZU5CG-SFVC784	14
XCZU5CG-FBVB900	52

Table 155: Maximum V_{CCINT} Current for Zynq UltraScale+ MPSoCs (cont'd)

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
XCZU6CG-FFVC900	49
XCZU6CG-FFVB1156	62
XCZU7CG-FBVB900	43
XCZU7CG-FFVC1156	36
XCZU7CG-FFVF1517	73
XCZU9CG-FFVC900	49
XCZU9CG-FBVB1156	62
EG Devices	
XCZU1EG-SBVA484	7
XAZU1EG-SBVA484	7
XCZU1EG-UBVA494	3.16
XCZU1EG-SFVA625	11
XAZU1EG-SFVA625	11
XCZU1EG-SFVC784	7
XAZU1EG-SFVC784	7
XCZU2EG-SBVA484	9
XAZU2EG-SBVA484	9
XCZU2EG-UBVA530	6
XCZU2EG-SFVA625	10
XAZU2EG-SFVA625	10
XCZU2EG-SFVC784	7
XAZU2EG-SFVC784	7
XCZU3EG-SBVA484	9
XAZU3EG-SBVA484	9
XQZU3EG-SFRA484	9
XCZU3EG-UBVA530	6
XCZU3EG-SFVA625	10
XAZU3EG-SFVA625	10
XCZU3EG-SFVC784	7
XAZU3EG-SFVC784	7
XQZU3EG-SFRC784	7
XCZU3TEG-SFVC784	14
XAZU3TEG-SFVC784	10
XCZU3TEG-SFVD784	23
XAZU3TEG-SFVD784	16
XCZU4EG-SFVC784	11
XCZU4EG-FBVB900	52

Table 155: Maximum V_{CCINT} Current for Zynq UltraScale+ MPSoCs (cont'd)

Device	Maximum Possible V _{CCINT} Current (A) (100°C, 10 Years)
XQZU4EG-SFRC784	11
XCZU5EG-SFVC784	14
XCZU5EG-FBVB900	52
XCZU6EG-FFVC900	49
XCZU6EG-FFVB1156	62
XCZU7EG-FBVB900	43
XCZU7EG-FFVC1156	36
XCZU7EG-FFVF1517	73
XCZU9EG-FFVC900	49
XQZU9EG-FFRC900	49
XCZU9EG-FFVB1156	62
XQZU9EG-FFRB1156	62
XCZU11EG-FFVC1156	39
XQZU11EG-FFRC1156	39
XCZU11EG-FFVB1517	50
XCZU11EG-FFVF1517	55
XAZU11EG-FFVF1517	55
XCZU11EG-FFVC1760	70
XQZU11EG-FFRC1760	70
XCZU15EG-FFVC900	51
XQZU15EG-FFRC900	51
XCZU15EG-FFVB1156	63
XQZU15EG-FFRB1156	63
XCZU17EG-FFVB1517	62
XCZU17EG-FFVC1760	78
XCZU17EG-FFVD1760	79
XCZU17EG-FFVE1924	76
XCZU19EG-FFVB1517	62
XQZU19EG-FFRB1517	62
XCZU19EG-FFVC1760	78
XQZU19EG-FFRC1760	78
XCZU19EG-FFVD1760	79
XCZU19EG-FFVE1924	76
EV Devices	
XCZU4EV-SFVC784	11
XAZU4EV-SFVC784	11
XCZU4EV-FBVB900	52

Table 155: Maximum V_{CCINT} Current for Zynq UltraScale+ MPSoCs (cont'd)

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
XCZU5EV-SFVC784	14
XAZU5EV-SFVC784	14
XQZU5EV-SFRC784	14
XQZU5EV-FFRB900	52
XCZU5EV-FBVB900	52
XCZU7EV-FBVB900	43
XAZU7EV-FBVB900	43
XQZU7EV-FFRB900	43
XCZU7EV-FFVC1156	36
XQZU7EV-FFRC1156	36
XCZU7EV-FFVF1517	73

 Table 156: Maximum V_{CCINT} Current for Zynq UltraScale+ RFSocS

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
Gen 1 Devices	
XCZU21DR-FFVD1156	50
XQZU21DR-FFRD1156	50
XCZU25DR-FFVE1156	50
XCZU25DR-FSVE1156	50
XCZU25DR-FFVG1517	50
XCZU25DR-FSVG1517	50
XCZU27DR-FFVE1156	50
XCZU27DR-FSVE1156	50
XCZU27DR-FFVG1517	50
XCZU27DR-FSVG1517	50
XCZU28DR-FFVE1156	50
XQZU28DR-FFRE1156	50
XCZU28DR-FSVE1156	50
XCZU28DR-FFVG1517	50
XQZU28DR-FFRG1517	50
XCZU28DR-FSVG1517	50
XCZU29DR-FFVF1760	50
XQZU29DR-FFRF1760	50
XCZU29DR-FSVF1760	50

Table 156: Maximum V_{CCINT} Current for Zynq UltraScale+ RFSocS (cont'd)

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
Gen 2 Devices	
XCZU39DR-FFVF1760	50
XCZU39DR-FSVF1760	50
Gen 3 Devices	
XCZU42DR-FSVE1156	64
XCZU42DR-FFVE1156	59
XCZU43DR-FSVE1156	50
XCZU43DR-FFVG1517	50
XCZU43DR-FSVG1517	50
XCZU43DR-FFVE1156	50
XCZU46DR-FFVH1760	50
XCZU46DR-FSVH1760	50
XCZU47DR-FFVE1156	50
XCZU47DR-FSVE1156	50
XCZU47DR-FFVG1517	50
XCZU47DR-FSVG1517	50
XCZU48DR-FFVE1156	50
XCZU48DR-FSVE1156	50
XQZU48DR-FFRE1156	50
XQZU48DR-FFQE1156	50
XQZU48DR-FSRE1156	50
XQZU48DR-FSQE1156	50
XCZU48DR-FFVG1517	50
XCZU48DR-FSVG1517	50
XQZU48DR-FFQG1517	50
XQZU48DR-FSRG1517	50
XQZU48DR-FSQG1517	50
XCZU49DR-FFVF1760	50
XCZU49DR-FSVF1760	50
XQZU49DR-FFQF1760	50
XQZU49DR-FSQF1760	50
XQZU49DR-FSRF1760	50
DFE Devices	
XCZU63DR-FFVE1156	59
XCZU63DR-FSVE1156	59
XCZU64DR-FFVE1156	59
XCZU64DR-FSVE1156	59

Table 156: Maximum V_{CCINT} Current for Zynq UltraScale+ RFSOCs (cont'd)

Device	Maximum Possible V_{CCINT} Current (A) (100°C, 10 Years)
XCZU65DR-FSVE1156	64
XCZU65DR-FFVE1156	59
XCZU67DR-FSVE1156	64
XCZU67DR-FFVE1156	59

Memory Derating Tables

The skew numbers presented in [Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices](#) were calculated to operate the memory interfaces at their maximum data rates. Some of those skew limits can be relaxed if the respective memory interface is not intended to be operated at the maximum data rate. The tables in this appendix represent the extent to which the respective skew numbers can be relaxed based on the speed rating of the FPGA, the memory component rating, and the actual speed at which the system is operated. For example, for DDR3 DQ to DQS (see the following table) using an FPGA rated at 2133 Mb/s and a memory component rated at 1333 Mb/s while operating at 1333 Mb/s, the skew can be relaxed from 10 ps to 87 ps. The derating tables are valid for both PL and PS memory interfaces.

Note: The shaded numbers in the following tables represent the maximum skew (plus or minus) allowed for the given FPGA speed rating and memory component rating when operated at that speed.

Table 157: DDR3 Data to DQS Skew Limit

FPGA Rating (Mb/s)		Memory Component Rating (Mb/s)					
Rated	Actual	2133	1866	1600	1333	1066	800
2133	2133	10	N/A	N/A	N/A	N/A	N/A
	1866	39	26	N/A	N/A	N/A	N/A
	1600	83	70	52	N/A	N/A	N/A
	1333	146	133	114	87	N/A	N/A
	1066	150	150	150	150	146	N/A
	800	150	150	150	150	150	150
1866	1866	18	10	N/A	N/A	N/A	N/A
	1600	63	50	31	N/A	N/A	N/A
	1333	125	112	94	66	N/A	N/A
	1066	150	150	150	150	125	N/A
	800	150	150	150	150	150	150
1600	1600	36	23	10	N/A	N/A	N/A
	1333	99	86	68	40	N/A	N/A
	1066	150	150	150	134	99	N/A
	800	150	150	150	150	150	150
1333	1333	64	51	33	10	N/A	N/A
	1066	150	145	126	99	64	N/A
	800	150	150	150	150	150	150

Table 157: DDR3 Data to DQS Skew Limit (cont'd)

FPGA Rating (Mb/s)		Memory Component Rating (Mb/s)					
Rated	Actual	2133	1866	1600	1333	1066	800
1066	1066	99	86	68	40	10	N/A
	800	150	150	150	150	150	99
800	800	150	148	130	102	67	10

Notes:

- See [Table 33: DDR3 SDRAM Data Group Skew Constraints](#) and [Table 40: DDR3 DIMM Data Group Skew Constraints](#) for the original specifications associated with this table.

Table 158: DDR3 Address/Command/Control to CK Skew Limit

FPGA Rating (Mb/s)		Memory Component Rating (Mb/s)					
Rated	Actual	2133	1866	1600	1333	1066	800
2133	2133	8	N/A	N/A	N/A	N/A	N/A
	1866	8	8	N/A	N/A	N/A	N/A
	1600	32	32	32	N/A	N/A	N/A
	1333	77	77	77	77	N/A	N/A
	1066	119	119	119	119	119	N/A
	800	150	150	150	150	150	150
1866	1866	8	8	N/A	N/A	N/A	N/A
	1600	32	32	32	N/A	N/A	N/A
	1333	77	77	77	77	N/A	N/A
	1066	119	119	119	119	119	N/A
	800	150	150	150	150	150	150
1600	1600	32	28	8	N/A	N/A	N/A
	1333	77	77	77	77	N/A	N/A
	1066	119	119	119	119	119	N/A
	800	150	150	150	150	150	150
1333	1333	58	48	28	8	N/A	N/A
	1066	119	119	119	119	119	N/A
	800	150	150	150	150	150	150
1066	1066	119	119	101	81	8	N/A
	800	150	150	150	150	150	150
800	800	150	150	150	150	150	150

Notes:

- See [Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints](#) and [Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints](#) for the original specifications associated with this table.

Table 159: DDR4 Data to DQS Skew Limit

FPGA Rating (Mb/s)	Running At (Mb/s)	Memory Component Rating (Mb/s)						
		3200	2933	2666	2400	2133	1866	1600
2666	2666	18	16	10	N/A	N/A	N/A	N/A
	2400	38	37	31	24	N/A	N/A	N/A
	2133	64	63	57	50	46	N/A	N/A
	1866	98	96	90	83	79	66	N/A
	1600	142	141	135	128	124	110	92
2400	2400	25	23	17	10	N/A	N/A	N/A
	2133	51	49	43	36	32	N/A	N/A
	1866	84	82	77	70	65	52	N/A
	1600	129	127	121	114	110	97	79
2133	2133	29	27	21	14	10	N/A	N/A
	1866	62	61	55	48	44	30	N/A
	1600	107	105	99	92	88	75	57
1866	1866	42	40	35	28	23	10	N/A
	1600	87	85	79	72	68	55	37
1600	1600	60	58	52	45	41	28	10

Notes:

- See [Table 26: DDR4 SDRAM Data Group Skew Constraints](#) and [Table 43: DDR4 DIMM Skew Constraints](#) for the original specifications associated with this table.

Table 160: DDR4 Address/Command/Control to CK Skew Limit

FPGA Rating (Mb/s)	Running At (Mb/s)	Memory Component Rating (Mb/s)						
		3200	2933	2666	2400	2133	1866	1600
2666	2666	23	15	8	N/A	N/A	N/A	N/A
	2400	65	57	50	43	N/A	N/A	N/A
	2133	117	109	102	95	77	N/A	N/A
	1866	150	150	150	150	144	124	N/A
	1600	150	150	150	150	150	150	150
2400	2400	30	22	15	8	N/A	N/A	N/A
	2133	82	74	67	60	42	N/A	N/A
	1866	149	141	134	127	109	89	N/A
	1600	150	150	150	150	150	150	150
2133	2133	48	40	33	26	8	N/A	N/A
	1866	115	107	100	93	75	55	N/A
	1600	150	150	150	150	150	144	129
1866	1866	68	60	53	46	28	8	N/A
	1600	150	149	142	135	117	97	82

Table 160: DDR4 Address/Command/Control to CK Skew Limit (cont'd)

FPGA Rating (Mb/s)	Running At (Mb/s)	Memory Component Rating (Mb/s)						
		3200	2933	2666	2400	2133	1866	1600
1600	1600	83	75	68	61	43	23	8

Notes:

- See [Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints](#) and [Table 44: DDR4 DIMM Address, Command, and Control Skew Constraints](#) for the original specifications associated with this table.

Table 161: RLD RAM 3 DQ/DM to DK_P/N Skew Limit

Rated	Running at	1200	1066	933	800
FPGA Rating (MHz)		Memory Component Rating (MHz)			
1200	1200	5	N/A	N/A	N/A
	1066	31	26	N/A	N/A
	933	65	60	45	N/A
	800	109	104	89	64
1066	1066	10	5	N/A	N/A
	933	43	38	23	N/A
	800	88	83	68	43
933	933	25	20	5	N/A
	800	70	65	50	25
800	800	50	45	30	5

Notes:

- See [Table 88: RLD RAM 3 Memory Skew Constraints](#) for the original specifications associated with this table.

Table 162: RLD RAM 3 DQ to QK_P/N Skew Limit

Rated	Running at	1200	1066	933	800
FPGA Rating (MHz)		Memory Component Rating (MHz)			
1200	1200	5	N/A	N/A	N/A
	1066	31	6	N/A	N/A
	933	65	40	9	N/A
	800	109	84	54	13
1066	1066	30	5	N/A	N/A
	933	63	38	8	N/A
	800	108	83	53	11
933	933	60	35	5	N/A
	800	105	80	50	8
800	800	105	77	46	5

Notes:

- See [Table 88: RLD RAM 3 Memory Skew Constraints](#) for the original specifications associated with this table.

Table 163: RLD RAM 3 DK_P/N to CK_P/N Skew Limit

Rated	Running at	1200	1066	933	800
FPGA Rating (MHz)		Memory Component Rating (MHz)			
1200	1200	5	N/A	N/A	N/A
	1066	57	29	N/A	N/A
	933	124	96	60	N/A
	800	150	150	149	101
1066	1066	33	5	N/A	N/A
	933	100	72	36	N/A
	800	150	150	125	77
933	933	69	41	5	N/A
	800	150	130	94	46
800	800	118	89	53	5

Notes:

- See [Table 88: RLD RAM 3 Memory Skew Constraints](#) for the original specifications associated with this table.

Table 164: RLD RAM 3 ADDR/CMD to CK_P/N Skew Limit

Rated	Running at	1200	1066	933	800
FPGA Rating (MHz)		Memory Component Rating (MHz)			
1200	1200	5	N/A	N/A	N/A
	1066	57	42	N/A	N/A
	933	124	109	74	N/A
	800	150	150	150	128
1066	1066	20	5	N/A	N/A
	933	87	72	37	N/A
	800	150	150	126	91
933	933	55	40	5	N/A
	800	144	129	94	59
800	800	90	75	40	5

Notes:

- See [Table 95: QDR II+ SRAM Skew Constraints](#) for the original specifications associated with this table.

Table 165: QDR II+ D to K_P/N Skew Limit

FPGA Rating (MHz)		Memory Component Rating (MHz)		
Rated	Actual	633	600	550
633	633	6	N/A	N/A
	600	27.72	27.72	N/A
	550	65.6	65.6	65.6

Table 165: QDR II+ D to K_P/N Skew Limit (cont'd)

FPGA Rating (MHz)		Memory Component Rating (MHz)		
Rated	Actual	633	600	550
600	633	N/A	N/A	N/A
	600	6	6	N/A
	550	43.88	43.88	43.88
550	633	N/A	N/A	N/A
	600	N/A	N/A	N/A
	550	6	6	6

Notes:

- See [Table 95: QDR II+ SRAM Skew Constraints](#) for the original specifications associated with this table.

Table 166: QDR II+ Q to CQ_P/N Skew Limit

Rated	Actual	633	600	550
FPGA Rating (MHz)		Memory Component Rating (MHz)		
633	633	6	N/A	N/A
	600	27.72	27.72	N/A
	550	65.6	65.6	6
600	633	N/A	N/A	N/A
	600	6	6	N/A
	550	43.88	43.88	6
550	633	N/A	N/A	N/A
	600	N/A	N/A	N/A
	550	66	66	6

Notes:

- See [Table 95: QDR II+ SRAM Skew Constraints](#) for the original specifications associated with this table.

Table 167: QDR II+ Data Skew Limit

Rated	Actual	633	600	550
FPGA Rating (MHz)		Memory Component Rating (MHz)		
633	633	5	N/A	N/A
	600	42.88	42.88	N/A
	550	64.6	64.6	64.6
600	633	N/A	N/A	N/A
	600	5	5	N/A
	550	42.88	42.88	42.88

Table 167: QDR II+ Data Skew Limit (cont'd)

Rated	Actual	633	600	550
FPGA Rating (MHz)		Memory Component Rating (MHz)		
550	633	N/A	N/A	N/A
	600	N/A	N/A	N/A
	550	5	5	5

Notes:

- See [Table 95: QDR II+ SRAM Skew Constraints](#) for the original specifications associated with this table.

Table 168: QDR II+ ADDR/CMD Skew Limit

Rated	Actual	633	600	550
FPGA Rating (MHz)		Memory Component Rating (MHz)		
633	633	6	N/A	N/A
	600	49.44	49.44	N/A
	550	125.2	125.2	125.2
600	633	N/A	N/A	N/A
	600	6	6	N/A
	550	81.76	81.76	81.76
550	633	N/A	N/A	N/A
	600	N/A	N/A	N/A
	550	6	6	6

Notes:

- See [Table 95: QDR II+ SRAM Skew Constraints](#) for the original specifications associated with this table.

Table 169: QDR II+ ADDR/CMD to CLK Skew Limit for Clock Point-to-Point and CA Fly-by

FPGA Rating (MHz)		Memory Component Rating (MHz)		
Rated	Actual	633	600	550
633	633	6	N/A	N/A
	600	49.44	49.44	N/A
	550	125.2	125.2	125.2
600	633	N/A	N/A	N/A
	600	6	6	N/A
	550	81.76	81.76	81.76
550	633	N/A	N/A	N/A
	600	N/A	N/A	N/A
	550	6	6	6

Notes:

- See [Table 95: QDR II+ SRAM Skew Constraints](#) for the original specifications associated with this table.

Table 170: QDR II+ ADDR/CMD to CLK Skew Limit for Clock T-Branch and CA Fly-by

Rated	Actual	633	600	550
FPGA Rating (MHz)		Memory Component Rating (MHz)		
633	633	34	N/A	N/A
	600	77.44	77.44	N/A
	550	153.2	153.2	153.2
600	633	N/A	N/A	N/A
	600	34	34	N/A
	550	109.76	109.76	109.76
550	633	N/A	N/A	N/A
	600	N/A	N/A	N/A
	550	34	34	34

Notes:

- See [Table 95: QDR II+ SRAM Skew Constraints](#) for the original specifications associated with this table.

Table 171: QDR II+ ADDR/CMD to CLK Skew Limit for Clock T-Branch and CA T-Branch

Rated	Actual	633	600	550
FPGA Rating (MHz)		Memory Component Rating (MHz)		
633	633	6	N/A	N/A
	600	49.44	49.44	N/A
	550	125.2	125.2	125.2
600	633	N/A	N/A	N/A
	600	6	6	N/A
	550	81.76	81.76	81.76
550	633	N/A	N/A	N/A
	600	N/A	N/A	N/A
	550	6	6	6

Notes:

- See [Table 95: QDR II+ SRAM Skew Constraints](#) for the original specifications associated with this table.

Material Properties and Insertion Losses

This appendix lists various material types and their associated characteristics along with relative cost multipliers to assist with choosing the best printed circuit board material.

Table 172: Material Parameters and Cost Multipliers

	Rogers Corporation (Rogers_4350)	Panasonic (MEGTRON 6)	Isola Group (I-Speed)	Isola Group (FR408HRIS)	Isola Group (Standard FR4_LP)	Isola Group (Standard FR4_LP)
DK	3.48	3.49	3.25	3.65	4.4	4.4
DF	0.0057	0.006	0.0074	0.0115	0.017	0.017
Surface Roughness	0.5 μm	2 μm	2 μm	2 μm	5 μm	8 μm
Copper Conductivity	4×10^7 S/m	4×10^7 S/m	4×10^7 S/m	4×10^7 S/m	4×10^7 S/m	4×10^7 S/m
Cost	3.8X	3.5X	2.3X	1.7X	1.2X	1X

Table 173: Insertion Loss @ 2.4 GHz

IL @ 2.4GHz (dB/in)	Rogers Corporation (Rogers_4350)	Panasonic (MEGTRON 6)	Isola Group (I-Speed)	Isola Group (FR408HRIS)	Isola Group (Standard FR4_LP)	Isola Group (Standard FR4_LP)
W = 4 mil	-0.318	-0.469	-0.474	-0.528	-0.658	-0.692
W = 6 mil	-0.246	-0.362	-0.367	-0.422	-0.539	-0.565
W = 8 mil	-0.211	-0.302	-0.308	-0.359	-0.467	-0.488
W = 10 mil	-0.184	-0.263	-0.272	-0.323	-0.427	-0.444
W = 12 mil	-0.167	-0.233	-0.246	-0.293	-0.399	-0.414

Notes:

- 1 mil = 0.0254 mm

Table 174: Insertion Loss @ 4.0 GHz

IL @ 4 GHz (dB/in)	Rogers Corporation (Rogers_4350)	Panasonic (MEGTRON 6)	Isola Group (I-Speed)	Isola Group (FR408HRIS)	Isola Group (Standard FR4_LP)	Isola Group (Standard FR4_LP)
W = 4 mil	-0.452	-0.667	-0.679	-0.765	-0.954	-1.002

Table 174: Insertion Loss @ 4.0 GHz (cont'd)

IL @ 4 GHz (dB/in)	Rogers Corporation (Rogers_4350)	Panasonic (MEGTRON 6)	Isola Group (I-Speed)	Isola Group (FR408HRIS)	Isola Group (Standard FR4_LP)	Isola Group (Standard FR4_LP)
W = 6 mil	-0.354	-0.518	-0.529	-0.616	-0.794	-0.829
W = 8 mil	-0.306	-0.435	-0.446	-0.531	-0.696	-0.724
W = 10 mil	-0.270	-0.381	-0.397	-0.481	-0.641	-0.665
W = 12 mil	-0.246	-0.340	-0.362	-0.440	-0.603	-0.624

Notes:

- 1 mil = 0.0254 mm

Table 175: Insertion Loss @ 5.0 GHz

IL @ 5 GHz (dB/in)	Rogers Corporation (Rogers_4350)	Panasonic (MEGTRON 6)	Isola Group (I-Speed)	Isola Group (FR408HRIS)	Isola Group (Standard FR4_LP)	Isola Group (Standard FR4_LP)
W = 4 mil	-0.531	-0.774	-0.788	-0.896	-1.126	-1.178
W = 6 mil	-0.417	-0.604	-0.618	-0.728	-0.942	-0.981
W = 8 mil	-0.362	-0.509	-0.525	-0.628	-0.831	-0.861
W = 10 mil	-0.321	-0.447	-0.467	-0.572	-0.768	-0.794
W = 12 mil	-0.294	-0.400	-0.427	-0.524	-0.724	-0.747

Notes:

- 1 mil = 0.0254 mm

Table 176: Insertion Loss @ 6.0 GHz

IL @ 6 GHz (dB/in)	Rogers Corporation (Rogers_4350)	Panasonic (MEGTRON 6)	Isola Group (I-Speed)	Isola Group (FR408HRIS)	Isola Group (Standard FR4_LP)	Isola Group (Standard FR4_LP)
W = 4 mil	-0.607	-0.872	-0.890	-1.018	-1.290	-1.356
W = 6 mil	-0.479	-0.683	-0.699	-0.832	-1.085	-1.134
W = 8 mil	-0.417	-0.577	-0.596	-0.720	-0.961	-0.999
W = 10 mil	-0.370	-0.508	-0.532	-0.658	-0.891	-0.924
W = 12 mil	-0.339	-0.455	-0.487	-0.604	-0.842	-0.871

Notes:

- 1 mil = 0.0254 mm

Table 177: Insertion Loss @ 12.9 GHz

IL @ 6 GHz (dB/in)	Rogers Corporation (Rogers_4350)	Panasonic (MEGTRON 6)	Isola Group (I-Speed)	Isola Group (FR408HRIS)	Isola Group (Standard FR4_LP)	Isola Group (Standard FR4_LP)
W = 4 mil	-1.099	-1.440	-1.485	-1.754	-2.307	-2.694
W = 6 mil	-0.882	-1.147	-1.191	-1.466	-1.996	-2.275

Table 177: Insertion Loss @ 12.9 GHz (cont'd)

IL @ 6 GHz (dB/in)	Rogers Corporation (Rogers_4350)	Panasonic (MEGTRON 6)	Isola Group (I-Speed)	Isola Group (FR408HRIS)	Isola Group (Standard FR4_LP)	Isola Group (Standard FR4_LP)
W = 8 mil	-0.775	-0.983	-1.032	-1.294	-1.792	-2.020
W = 10 mil	-0.695	-0.878	-0.932	-1.198	-1.683	-1.877
W = 12 mil	-0.642	-0.795	-0.863	-1.115	-1.607	-1.776

Notes:

- 1 mil = 0.0254 mm

Additional Resources and Legal Notices

Finding Additional Documentation

Technical Information Portal

The AMD Technical Information Portal is an online tool that provides robust search and navigation for documentation using your web browser. To access the Technical Information Portal, go to <https://docs.amd.com>.

Documentation Navigator

Documentation Navigator (DocNav) is an installed tool that provides access to AMD Adaptive Computing documents, videos, and support resources, which you can filter and search to find information. To open DocNav, do the following:

- From the AMD Vivado™ IDE, select **Help** → **Documentation and Tutorials**.
- On Windows, click the **Start** button and select **AMDDesignTools** → **DocNav**.
- At the Linux command prompt, enter `docnav`.

Note: For more information on DocNav, refer to the *Documentation Navigator User Guide* ([UG968](#)).

Design Hubs

AMD Design Hubs provide links to documentation organized by design tasks and other topics, which you can use to learn key concepts and address frequently asked questions. To access the Design Hubs, do the following:

- In DocNav, click the **Design Hubs View** tab.
- Go to the [Design Hubs](#) web page.

Support Resources

For support resources such as Answers, Documentation, Downloads, and Forums, see [Support](#).

References

These documents provide supplemental material useful with this guide:

1. *Kintex UltraScale FPGAs Data Sheet: DC and AC Switching Characteristics* ([DS892](#))
2. *Virtex UltraScale FPGAs Data Sheet: DC and AC Switching Characteristics* ([DS893](#))
3. *Kintex UltraScale+ FPGAs Data Sheet: DC and AC Switching Characteristics* ([DS922](#))
4. *Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics* ([DS923](#))
5. *Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics* ([DS925](#))
6. *Zynq UltraScale+ RFSoc Data Sheet: DC and AC Switching Characteristics* ([DS926](#))
7. *Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics* ([DS930](#))
8. *UltraScale Architecture GTH Transceivers User Guide* ([UG576](#))
9. *UltraScale Architecture GTY Transceivers User Guide* ([UG578](#))
10. *Kintex UltraScale and Virtex UltraScale FPGAs Schematic Review Checklist* ([XTP344](#))
11. *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification* ([UG575](#))
12. *UltraScale Architecture SelectIO Resources User Guide* ([UG571](#))
13. *Spartan UltraScale+ FPGAs SelectIO Resources User Guide* ([UG861](#))
14. *Interfacing 7 Series FPGAs High-Performance I/O Banks with 2.5V and 3.3V I/O Standards* ([XAPP520](#))
15. *UltraScale Architecture Configuration User Guide* ([UG570](#))
16. *UltraScale Architecture-Based FPGAs Memory IP LogiCORE IP Product Guide* ([PG150](#))
17. *UltraScale Devices Gen3 Integrated Block for PCI Express LogiCORE IP Product Guide* ([PG156](#))
18. *UltraScale Devices Integrated 100G Ethernet LogiCORE IP Product Guide* ([PG165](#))
19. *Integrated Interlaken up to 150G LogiCORE IP Product Guide* ([PG169](#))
20. *Zynq UltraScale+ RFSoc RF Data Converter LogiCORE IP Product Guide* ([PG269](#))
21. *UltraScale+ FPGAs and Zynq UltraScale+ Devices Schematic Review Checklist* ([XTP427](#))
22. *UltraScale Architecture and Product Data Sheet: Overview* ([DS890](#))
23. [VESA DisplayPort Standard Version 1](#)

- 24. [JESD84-B451, Embedded Multi-media Card \(eMMC\), Electrical Standard \(5.1\)](#)
- 25. [Zynq UltraScale+ Device Technical Reference Manual \(UG1085\)](#)
- 26. [Managing Power and Performance with the Zynq UltraScale+ MPSoC \(WP482\)](#)
- 27. [UltraScale Architecture System Monitor User Guide \(UG580\)](#)
- 28. [Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide \(UG1075\)](#)
- 29. [UltraScale+ Devices Integrated Block for PCI Express LogiCORE IP Product Guide \(PG213\)](#)
- 30. [UltraScale+ Devices Integrated 100G Ethernet Subsystem Product Guide \(PG203\)](#)
- 31. [Interlaken 150G LogiCORE IP Product Guide \(PG212\)](#)
- 32. [Recommended Design Rules and Strategies for BGA Devices User Guide \(UG1099\)](#)

Revision History

The following table shows the revision history for this document.

Section	Revision Summary
12/23/2025 Version 1.29	
General updates	<ul style="list-style-type: none"> • Removed description of PCB decoupling capacitors from Chapter 1: Power Distribution System in UltraScale Devices. • Removed recommended decoupling capacitor tables for XC, XA, and XQ devices (Appendix A, Appendix B, and Appendix C). • Replaced XPE with PDM throughout.
Overview	Rewrote section to apply to Spartan UltraScale+ FPGAs.
Fly-by and Clamshell Topologies	Removed important note.
Routing Rule Changes for Thicker Printed Circuit Boards	<ul style="list-style-type: none"> • Replaced links to Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923) with Spartan UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS930). • In Table 107: DDR4 Maximum Data Rates per Board Thickness, removed 1-Rank, 2-Rank, and 4-Rank DIMM rows.
alert_n	<ul style="list-style-type: none"> • Updated first paragraph. • Updated paragraph before Figure 105: Address/Command/Control Topology for DDR4 SDRAM Clamshell.
Table 116: LPDDR5 Signal Definitions	Removed CS1_A, CS1_B and ZQ_A_C, ZQ_B_D.
Figure 109: LPDDR5 1x16 Single Rank	<ul style="list-style-type: none"> • Updated figure title. • Updated ZQ to pull up to VDDQ. • Removed CS_0[1] to CS[1] connection.
Figure: LPDDR5 2x16 Single/Dual Rank	Removed figure.

Section	Revision Summary
Figure 110: LPDDR5 1x32 Single Rank	<ul style="list-style-type: none"> Updated figure title. Added pull-down resistors on CS[0]_A and CS[1]_A. Updated ZQ_A to pull up to VDDQ. Removed CS_0[1] to CS[1]_A connection. Removed CS_0[3] to CS[1]_B connection. Renamed CS_0[2] to CS_0[1].
Supply Voltages between FPGA and LPDDR5 Device	Added section.
Table 118: Physical Design Rules for LPDDR5 Signals	Added note 2.
Table 119: Skew Constraint Rules for LPDDR5 Signals	Removed notes about individual constraint sets and Quad-die/Eight-die devices.
Table 120: LPDDR4x Signal Definitions	<ul style="list-style-type: none"> Removed CS0_A, CS1_A. Updated CKE0_A, CKE0_B. Removed ZQ1. Updated DMI[1:0]_A, DMI[1:0]_B signals. Updated LPDDR4x Device and FPGA subheadings. Removed table note.
Supply Voltage between FPGA and LPDDR4x	<ul style="list-style-type: none"> Updated section title. Updated introductory text before figures. Removed the following figures: <ul style="list-style-type: none"> Connections for a 2x32 LPDDR4x Interface Connections for a 1x32 (or x48) with ECC LPDDR4x Interface Connections for a 2x16 LPDDR4x Interface Connections for a Pin Efficient 2x32 LPDDR4x Interface Connections for a Pin Efficient 1x32 LPDDR4x Interface
Table 123: Skew Constraint Rules for LPDDR4x Signals1	<ul style="list-style-type: none"> Updated group for Command to Clock A and Command to Clock B. Removed note about signals only being present on dual rank devices.
Unused ADC and DAC Power Pins	Updated second bullet.
SD/SDIO	Added skew guidance, including for when using a level shifter.
Table 131: PCB Design Checklist for PS-GTR	Added notes to PS_MGTRAVCC and PS_MGTRAVTT saying that PS MGT should be powered to allow zeroization to occur.
Table 132: MIPI D-PHY PCB Trace and Skew Guidelines	Removed sentence about reference stackup being used for longer trace lengths from note 2.
12/19/2024 Version 1.28	
Chapter 3: PCB Guidelines for Memory Interfaces in Spartan UltraScale+ Devices	Added new chapter.
AC/DC Coupling Guidelines	Changed DC common mode point from 2.1V to 2.07V.
Recommended Decoupling Capacitor Quantities for Artix UltraScale+, Kintex UltraScale+, Spartan UltraScale+, and Virtex UltraScale+ XC Devices	Added Spartan UltraScale+ FPGAs.

Section	Revision Summary
Table 1	Added XCZU4CG-SFVE784, XCZU5CG-SFVE784, XCZU4EG-SFVE784, XCZU5EG-SFVE784, XCZU4EV-SFVE784, and XCZU5EV-SFVE784.
Table 1	Added XAAU7P-FCVA289.
Table 2	Added XAZU3TEG-SFVC784.
Table 3	Changed capacitances for $V_{CC_PSINTFP}$, $V_{CC_PSINTLP}$, V_{CC_PSAUX} , V_{CC_PSPLL} , $V_{CC_PSINTFP_DDR}$, and V_{CCO_PSIOX} from 10 μ F to 2.2 μ F.
Table 5	Change capacitance for V_{CCINT_AMS} from 1.0 μ F to 2.2 μ F.
11/14/2023 Version 1.27	
Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices	Removed XSR from Figure 15: System Clock .
Decoupling for Commercial Grade XC Devices	Added XCAU7P-FCVA289 to Table 1.
Decoupling for Commercial Grade XA Devices	Added XAZU4EV-SFVC784 to Table 4.
Appendix A: Maximum Current Draw for V_{CCINT} in UltraScale+ Devices	Added new appendix.
05/19/2023 Version 1.26	
Chapter 1: Power Distribution System in UltraScale Devices : Updated notes after Figure 4: Always On: Cost-Optimized Power Rail Consolidation , Figure 6: Always On: Power and/or Efficiency-Optimized Power Rail Consolidation for Low-Power Devices , Figure 7: Always On: PL Performance-Optimized Rail Consolidation and Figure 8: Full Power Domain Flexibility Consolidation .	
Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices : Replaced <i>address</i> with <i>clock</i> in item 4 of General Memory Routing Guidelines .	
Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS : Added XCZU63DR-FFVE1156, XCZU63DR-FSVE1156, XCZU64DR-FFVE1156, XCZU64DR-FSVE1156, XCZU65DR-FSVE1156, XCZU67DR-FSVE1156, XQZU65DR-FFRE1156, and XQZU67DR-FFRE1156 device-packages to Table 156: Maximum VCCINT Current for Zynq UltraScale+ RFSocS .	
Decoupling for Commercial Grade XC Devices: Updated pin count for XCAU7P-SBVC484 in Table 1. Removed XQKU15P-FFRA1156 and XQKU15P-FFRE1517 from Table 2. Added notes to Table 4 and Table 5. Added XCZU63DR-FFVE1156, XCZU63DR-FSVE1156, XCZU64DR-FFVE1156, XCZU64DR-FSVE1156, XCZU65DR-FSVE1156, and XCZU67DR-FSVE1156 to Table 6, Table 8, and Table 9.	
Decoupling for Automotive Grade XA Devices: Corrected part and package numbers in Table 4.	
03/06/2023 Version 1.25	
Chapter 1: Power Distribution System in UltraScale Devices : Moved all capacitor tables to Decoupling for Commercial Grade XC Devices. Added XCKU19P to Table 1-11. Added XCVU57P to Table 1-12. Added XCZU1CG and XCZU1EG to Table 1-13. Added Table 1-14.	
Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices : Removed ECC Connection Rules for DDR3 SDRAM section. In LPDDR4 x32 with ECC Memory Interface Signals and Connections , replaced SDP and DDP with "single-rank" and "dual-rank."	
Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS : Moved decoupling guidelines to Decoupling for Commercial Grade XC Devices.	
Decoupling for Commercial Grade XC Devices: Added new appendix.	
Decoupling for Automotive Grade XA Devices: Added new appendix.	
Decoupling for Defense Grade XQ Devices: Added new appendix.	
07/27/2022 Version 1.24	
Chapter 1: Power Distribution System in UltraScale Devices : Updated notes 1 and 2 in Table 1-2, Table 1-3, and Table 1-8. In Table 1-4, added 1.0 μ F sub-column to V_{CCINT}/V_{CCINT_IO} , V_{CCAUX}/V_{CCAUX_IO} , and HDIO/HPIO columns, added XCAU10P-UBVA368 and XCAU15P-UBVA368, updated capacitor recommendations for V_{CCINT}/V_{CCINT_IO} at 100 μ F, 47 μ F, and 10 μ F, and updated notes 1, 2, and 3. Updated notes 1, 2, and 3 in Table 1-5, Table 1-6, Table 1-7, and Table 1-10. Added Recommended Decoupling Capacitor Quantities for Zynq UltraScale+ Automotive Devices.	

Section	Revision Summary
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Added ECC Connection Rules for DDR4 SDRAM and ECC Connection Rules for DDR3 SDRAM. Added L2 breakouts before and after L3 in Figure 30: Address, Command, and Control Fly-by Termination for DDR4 SDRAM, Figure 33: Clock Fly-by Termination for DDR4 SDRAM, Figure 36: Address Fly-by Termination for DDR3 SDRAM and Figure 38: Clock Fly-by Termination for DDR3 SDRAM. Changed DRAM to SRAM in Figure 80: Address and Command Point-to-Point Routing and Termination for QDR II+ SRAM.</p>
	<p>Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Updated notes 1 and 2 in Table 3-1.</p>
	<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added bullet about ensuring SD signals have a total delay of less than 1.3 ns to SD/SDIO. Added bullet about PS-GTR reference clock to list of recommendations in Table 131: PCB Design Checklist for PS-GTR.</p>
	<p>Chapter 8: Migration between UltraScale and UltraScale+ FPGAs: Updated VCCINT_IO connection and R1 for 3, -2, -1 in Table 137: VCCINT/VCCINT_IO/VCCBRAM Connection Matrix.</p>
	<p>Chapter 9: Migration between Zynq UltraScale+ MPSoCs and Packages: In Table 143: MPSoC PS Voltage Matrix by Speed/Temperature Grade, changed VPS_MGTRAVCC for -3E from 0.90V to 0.85V.</p>
04/15/2022 Version 1.23	
	<p>Chapter 1: Power Distribution System in UltraScale Devices: Added XCZU1CG and XCZU1EG devices to Table 1-10. Added UBVA494 to Recommended Decoupling Capacitor Quantities for Zynq UltraScale+ Devices in UBVA494 and UBVA530 Packages.</p>
	<p>Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Added XQZU65DR-FFRE1156 and XQZU67DR-FFRE1156 to Table 3-1.</p>
01/20/2022 Version 1.22.1	
General updates	Updated description of reset_n termination throughout.
01/06/2022 Version 1.22	
	<p>Chapter 1: Power Distribution System in UltraScale Devices: Added $V_{CC_PSTINFP}/V_{CC_PSTINTLP}$ to Table 1. Added Table 1-4. Removed $V_{CCINT}/V_{CCBRAM}/V_{CCINT_IO}$ from column heading in Table 1-8. Added 1.0 μF capacitor recommendation for $V_{CC_PSTINFP_DDR}$ in Table 1-13.</p>
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: In Table 9: Reference Stackup updated L1 and L16 thickness from 2.5 to 0.6 mil. Replaced V_{DD2} with V_{DDQ} in Table 45: Signal Connection Matrix for x32 Dual Channel LPDDR4 SDP without ECC, Table 46: Signal Connection Matrix for x32 Dual Channel LPDDR4 DDP without ECC, Table 55: Signal Connection Matrix for x32 Dual Channel LPDDR4 Single-Rank with ECC and Table 56: Signal Connection Matrix for x32 Dual Channel LPDDR4 Dual-Rank with ECC. In Table 45: Signal Connection Matrix for x32 Dual Channel LPDDR4 SDP without ECC and Table 55: Signal Connection Matrix for x32 Dual Channel LPDDR4 Single-Rank with ECC, updated PCB termination at far end for CKE0. In Table 46: Signal Connection Matrix for x32 Dual Channel LPDDR4 DDP without ECC and Table 56: Signal Connection Matrix for x32 Dual Channel LPDDR4 Dual-Rank with ECC, updated PCB termination at far end for CKE0 and CKE1. Removed option 1 (80Ω termination to V_{TT}) from Figure 48: Clock Enable Routing for LPDDR4 without ECC and Figure 54: CKE0/CKE1 Routing for LPDDR4 with ECC. In Table 68: Routing Guidelines for HDI LPDDR4 Signals, replaced UBVA530 with FGA530 package, added row for Top-L5-Top to blind/buried via routing, and added rows for Minimum spacing: CAC Signals, Minimum spacing: CAC Signals to CK, and Minimum spacing: CAC/CK to DQ/DQS.</p>
	<p>Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Added XQZU48DR-FFRE1156, XQZU48DR-FSRG1517, XQZU49DR-FSRF1760, XCZU65DR-FFVE1156, and XCZU67DR-FFVE1156 to Table 3-1. Added XCZU65DR-FFVE1156 and XCZU67DR-FFVE1156 to Table 156: Maximum VCCINT Current for Zynq UltraScale+ RFSocS. Removed Table 3-5: Data Converter Bandwidth. Updated first sentence in paragraph after Table 124: Balun Specification Recommendations. Added paragraphs about P to N skew requirements and derating to Recommended Clocking Options. In Table 125: Signal Integrity Specifications for DAC/ADC Pairs and Clocks, changed column heading to Maximum Frequency, added row for CMRR under RFDC clocking, and added notes. In AC/DC Coupling Guidelines, added paragraph about AC coupling, updated second paragraph, and added notes. In Table 126: Trace Length Matching, added comments for DAC/ADC data converters, updated comments for DAC data converters, ADC data converters and ADC clocks, and added DAC clock. Removed Figure 3-18: Analog_SYSREF Levels. Added important note to Pi Network for Improved Return Loss in Gen 1 Devices (XCZU25DR/ XCZU27DR/XCZU28DR/XCZU29DR). In SYSREF, removed bullet about Analog_SYSREF and PL_SYSREF, and added bullet about multi-chip synchronization. Added two bullets to Unused ADC & DAC Pins. Updated first two paragraphs in Power Regulation and Decoupling for ADC and DAC Supplies. Updated ADC_AVCC nominal voltage in Table 129: ADC and DAC Voltage Supply Specifications for Gen 3 Devices¹. Added sentence about glitches in PDN to Power Delivery Network Design for Time Division Duplex.</p>
	<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added bullet about 10 kΩ pull-up resistor to eMMC. Added bullet about interconnected JTAG chains to JTAG.</p>
	<p>Appendix B: Memory Derating Tables: Added sentence about derating tables to first paragraph.</p>

Section	Revision Summary
06/03/2021 Version 1.21	
<p>Chapter 1: Power Distribution System in UltraScale Devices: Added Recommended Decoupling Capacitor Quantities for Zynq UltraScale+ Devices in UBVA494 and UBVA530 Packages. In Table 1-14 added row for 1.0 μF.</p>	
<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Added PCB Routing Guidelines for LPDDR4 Memories in High-Density Interconnect Boards.</p>	
02/12/2021 Version 1.20	
<p>Chapter 1: Power Distribution System in UltraScale Devices: Added XCKU19P to Table 1-5. Added XCVU23P-FSVJ1760 to Table 1-6. Added XCVU57P-FSVK2892 to Table 1-8. Added VU57P to Table 1-9. Updated first sentence in V_{CCINT_VCU} Plane Design and Power Delivery.</p>	
<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Updated item 13 in General Memory Routing Guidelines. Updated first paragraph in PCB Guidelines for DDR4 SDRAM (PL and PS). Added Routing Rule Changes for Thicker Printed Circuit Boards.</p>	
<p>Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Added XCZU42DR to Table 3-1. Added paragraph about clock forwarding capability in Gen 3 RFSoc devices to Recommended Clocking Options. Added Table 129: ADC and DAC Voltage Supply Specifications for Gen 3 Devices¹. Updated Powering RFSocS with Switch Regulators. Added Power Delivery Network Design for Time Division Duplex.</p>	
<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added bullet about device without DQS pin to DDR Mode (100 MHz). In SD/SDIO, added note about external pull-up resistor after fifth bullet, and added two bullets about level shifters.</p>	
<p>Chapter 12: PCB Technology Basics: Replaced I/O with I/O/PSIO in Unconnected V_{CCO} Pins.</p>	
09/02/2020 Version 1.19	
<p>Chapter 1: Power Distribution System in UltraScale Devices: In Table 1-5, updated packages for XQKU5P and XCVU7P, added row for XCVU23P-VSVA1365, and updated note 3. In Table 1-10, updated packages for XCZU3CG, XCZU6CG, XCZU9CG, XCZU3EG, XCZU6EG, XCZU9EG, and XCZU15EG. Replaced Table 1-21: Maximum VCCINT Current for Virtex UltraScale+ Devices with Table 1-11, Table 1-12, and Table 1-13.</p>	
<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Updated items 2 and 13 in General Memory Routing Guidelines. In Table 17: DDR4 SDRAM I/O Signal Description, updated address, write enable, row address strobe, column address strobe, chip select, and alert_n signals. Updated alert_n.</p>	
<p>Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Updated V_{CCSDFEC} and Migration. Added XCZU43DR, XCZU46DR, XCZU47DR, XCZU48DR, and XCZU49DR to Table 3-1 and Table 156: Maximum VCCINT Current for Zynq UltraScale+ RFSocS. In Table 3-1, removed note about connecting V_{CCINT}, V_{CCBRAM}, and V_{CCINT_IO} together for non-L speed grades. Added Maximum Current Draw for Zynq UltraScale+ RFSocS. Added notes to Table 126: Trace Length Matching. Removed Table 3-8: Analog_SYSREF Requirements and Table 3-9: DC Coupling Input Parameters. Updated VCCINT_AMS nominal voltage in Table 3-11.</p>	
<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added note to JTAG and PCIe. Added bullets about termination and signal integrity simulation to PS Reference Clock. Added bullet about SDIO 0/1 power control signal to SD/SDIO. Added recommended capacitor values for PCIe in Table 131: PCB Design Checklist for PS-GTR.</p>	
<p>Chapter 10: Migration between Virtex UltraScale+ VU13P Devices and VU27P/VU29P Devices: Updated 4.7 μF to 10 μF in Table 145: Decoupling for Non-2LE Devices and Table 146: Decoupling for -2LE Devices.</p>	
<p>Appendix B: Memory Derating Tables: Updated Table 166: QDR II+ Q to CQ_P/N Skew Limit.</p>	
11/26/2019 Version 1.18	
<p>Chapter 1: Power Distribution System in UltraScale Devices: In Table 1-6, updated existing XCVU19P-FSVA3824 and XCVU19P-FSVB3824 recommendations, and added XCVU19P-FSVA3824 and XCVU19P-FSVB3824 recommendations for thicker boards.</p>	
<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Updated CK to DQS skew constraints in Table 54: LPDDR4 without ECC Skew Constraints and Table 67: LPDDR4 with ECC Skew Constraints.</p>	
<p>Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Updated all bandwidths in Table 3-5. Updated Anaren part number and bandwidth in Choosing the Appropriate Balun. Updated second paragraph in AC/DC Coupling Guidelines. Updated note in Isolation Recommendations. Removed Table 3-7: Available S-Parameter Models for ADC/DAC Channels. Updated Pi Network for Improved Return Loss in Gen 1 Devices (XCZU25DR/ XCZU27DR/XCZU28DR/XCZU29DR) heading. In Table 128: ADC and DAC Voltage Supply Specifications for Gen 1 and Gen 2 Devices¹, replaced LDO with VRM and updated note 4.</p>	

Section	Revision Summary
Chapter 12: PCB Technology Basics : Updated second paragraph in Unconnected V_{CC0} Pins .	
08/29/2019 Version 1.17	
Chapter 1: Power Distribution System in UltraScale Devices : Updated note 4 in Table 1-5. In Table 1-6, added XCVU19P, updated note 4, and added note 5. Updated note 4 in Table 1-7. In Table 1-8, added XCVU45P and XCVU47P, and updated note 4. Added VU45P and VU47P to Table 1-9. Removed heading for V _{CC_PSBATT} at 100 µF from Table 1-11.	
Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices : Added note about multiple DIMMs in DDR3/DDR4 UDIMM/RDIMM/SODIMM/LRDIMM Routing Guidelines (PL and PS) . Updated DIMM labels in Figure 42: CMD/ADDR Fly-by Routing for DDR3/DDR4 DIMM Topology , Figure 44: DQS Fly-by Routing for DDR3/DDR4 DIMM Topology , and Figure 45: DQ Fly-by Routing for DDR3/DDR4 DIMM Topology . Updated PCB termination at far end for CS0 in Table 55: Signal Connection Matrix for x32 Dual Channel LPDDR4 Single-Rank with ECC .	
Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC : Updated trace delay to 500 ps in fifth bullet of QSPI .	
06/26/2019 Version 1.16	
Chapter 1: Power Distribution System in UltraScale Devices : Updated second paragraph in Recommended PCB Capacitors per Device. Added note at end of Step Load Assumptions. Updated Table 1-2, Table 1-3, Table 1-5, Table 1-6, Table 1-9, Table 1-10, and Table 1-11. Added Table 1-7 and Table 1-8. Removed sections <i>PCB Decoupling Capacitors for Virtex UltraScale+ 58G-Enabled Devices</i> and <i>PCB Decoupling Capacitors for Virtex UltraScale+ High Bandwidth Memory Devices</i> . Added Capacitor Specifications for XC Devices. Updated capacitor values and part numbers in V_{CC_PSDDR_PLL} Supply , including Figure 1: Connecting V_{CC_PSDDR_PLL} . Updated 10 µF 0402 and 47 µF 0603 part numbers in Table 1-17. Removed sections <i>Capacitor Specifications</i> and <i>Capacitor Consolidation Rules</i> .	
Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices : Updated Table 11: Example Signal Group Skew Constraint and Figure 13: Skew Example for Clock Signals . Added VRP (PL) and ZQ (PS) to Table 17: DDR4 SDRAM I/O Signal Description and Table 28: DDR3 SDRAM Interface Signal Description . Added vias to Figure 40: CK Point-to-Point Routing for DDR3/DDR4 DIMM Topology to Figure 45: DQ Fly-by Routing for DDR3/DDR4 DIMM Topology . Added reset_n . Removed note about pin ending in termination from Table 56: Signal Connection Matrix for x32 Dual Channel LPDDR4 Dual-Rank with ECC . Added VRP to Table 80: RDRAM 3 Memory Interface Signal Description , Table 89: QDR II+ SRAM Interface Signal Description , and Table 96: QDR-IV SRAM Interface Signal Description .	
Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS : Updated capacitance values in Figure 114: Recommended Plane Connection for VCCSDFEC and VCCBRAM . Updated Table 3-1, Table 3-2, and Table 3-3. Removed Table 3-2: Programmable Logic Rail Decoupling Guidelines for L Devices.	
Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC : In Table 131: PCB Design Checklist for PS-GTR , changed filter capacitor recommendation for PS_MGTRAVCC and PS_MGTRAVTT from 4.7 µF to 10 µF.	
Chapter 6: PCB Routing Guidelines for MIPI D-PHY : Added new chapter.	
05/09/2019 Version 1.15	
Chapter 1: Power Distribution System in UltraScale Devices : Updated recommended note in Step Load Assumptions. Added note 4 to Table 1-11.	
Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices : Added LRDIMM to DDR3/DDR4 UDIMM/RDIMM/SODIMM/LRDIMM Routing Guidelines (PL and PS) , DDR3 UDIMM/RDIMM/SODIMM/LRDIMM Routing Constraints , and DDR4 UDIMM/RDIMM/SODIMM/LRDIMM Routing Constraints headings. Added CK (A/B) to DQS0/1 (A/B) to Table 54: LPDDR4 without ECC Skew Constraints and Table 67: LPDDR4 with ECC Skew Constraints .	
Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS : Added XCZU39DR-FFVF1760 and XCZU39DR-FSVF1760 to Table 3-1 and Table 3-2. Updated ground stitching bullet in Analog Ground to Digital Ground Connection . Updated ADC input, DAC output, and DAC clock input bandwidths in Table 3-5. In Choosing the Appropriate Balun , updated Anaren bullets and added bullet for Mini Circuits TCM2-33X+. Updated Table 125: Signal Integrity Specifications for DAC/ADC Pairs and Clocks . Updated Trace Routing Impedance Recommendation . Removed Figure 3-10: Ground Stitching around RFSoc Pins . Removed Table 3-8: Signal Integrity Specifications for DAC/ADC Clocks and Reference Clocks and Table 3-9: Isolation Recommendations for ADC & DAC Pairs . Added routing guidelines to Trace Routing Impedance Recommendation . Updated Figure 123: Cross-Section of Guard Traces . Added Table 126: Trace Length Matching . Added paragraph about electrical length calculation after Figure 126: Via Configuration for Signal Layer Transition . Removed <i>Inter-Pair Skew</i> section. Added device numbers to Pi Network for Improved Return Loss in Gen 1 Devices (XCZU25DR/ XCZU27DR/XCZU28DR/ XCZU29DR) heading. Added paragraph about AC and DC coupling after Table 3-9. Removed <i>Ground Plane Cutout Under Passive Components</i> section. Updated Sample Stackup . Replaced LDO with VRM in Power Regulation and Decoupling for ADC and DAC Supplies and Figure 144: Power Regulation for Analog Supplies . Added Powering RFSocS with Switch Regulators .	

Section	Revision Summary
	<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added note about PS_INIT_B to PS_INIT_B, PS_PROG_B, and PS_DONE. Updated first two bullets in PS Reset (External System Reset and POR Reset). Added note to Table 131: PCB Design Checklist for PS-GTR.</p>
	<p>Appendix B: Memory Derating Tables: Updated table title and removed note 1 in Table 158: DDR3 Address/Command/Control to CK Skew Limit and Table 160: DDR4 Address/Command/Control to CK Skew Limit.</p>
01/04/2019 Version 1.14	
	<p>Chapter 1: Power Distribution System in UltraScale Devices: Added XQKU5P-SFRB784, XQKU5P-FFRB676, XQKU15P-FFRA1156, and XQKU15P-FFRE1517 to Table 1-5. Added XQVU3P-FFRC1517, XQVU7P-FLRA2104, XQVU7P-FLRB2104, and XQVU11P-FLRC2104 to Table 1-6. Added PCB Decoupling Capacitors for Virtex UltraScale+ 58G-Enabled Devices. In Table 1-10, added note 5, and updated Example Part Number column heading, part number for 680 µF capacitor, and note 1. Added XQZU3EG-SFRA484, XQZU3EG-SFRC784, XQZU9EG-FFRC900, XQZU9EG-FFRB1156, XQZU11EG-FFRC1156, XQZU11EG-FFRC1760, XQZU15EG-FFRC900, XQZU15EG-FFRB1156, XQZU19EG-FFRB1517, XQZU19EG-FFRC1760, XQZU5EV-SFRC784, XQZU5EV-FFRB900, XQZU7EV-FFRB900, and XQZU7EV-FFRC1156 to Table 1-10. In Table 1-14, updated suggested part number for 680 µF capacitor.</p>
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Updated guidelines 12 and 13 in General Memory Routing Guidelines. Replaced V_{TT} with V_{DDQ} in PCB termination for CKE0 and CKE1 in Table 45: Signal Connection Matrix for x32 Dual Channel LPDDR4 SDP without ECC, Table 46: Signal Connection Matrix for x32 Dual Channel LPDDR4 DDP without ECC, Table 55: Signal Connection Matrix for x32 Dual Channel LPDDR4 Single-Rank with ECC, and Table 56: Signal Connection Matrix for x32 Dual Channel LPDDR4 Dual-Rank with ECC. Added note 2 to Table 54: LPDDR4 without ECC Skew Constraints, and note 3 to Table 67: LPDDR4 with ECC Skew Constraints.</p>
	<p>Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Added XQZU28DR-FFRE1156 and XQZU28DR-FFRG1517 to Table 3-1. Replaced “PL GPIO pins” with “dedicated clock inputs” in first paragraph of SYSREF. Added note 1 to Table 3-8. Added DAC_AVTT to second paragraph in AC/DC Coupling Guidelines. Replaced “rectangular” with “circular” in title of Figure 128: Circular Antipad. Updated coupling value in Table 3-8. Added Figure 130: Analog SYSREF Receiver Equivalent Circuit and Table 3-9. In Table 128: ADC and DAC Voltage Supply Specifications for Gen 1 and Gen 2 Devices, added note 1 and removed Maximum Current column. In Unused ADC & DAC Pins, removed second bullet, updated third bullet, and removed Table 3-19: Leakage Current Values for ADC and DAC Tiles.</p>
	<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added note about PS_PROG_B and PS_POR_B to PS_INIT_B, PS_PROG_B, and PS_DONE.</p>
	<p>Chapter 10: Migration between Virtex UltraScale+ VU13P Devices and VU27P/VU29P Devices: Added new chapter.</p>
	<p>Appendix B: Memory Derating Tables: Updated skew values in second paragraph.</p>
08/15/2018 Version 1.13	
	<p>Chapter 1: Power Distribution System in UltraScale Devices: Added description of slew rates after Table 1. Added XCVU13P-FSGA2577 to Table 1-6. In Table 1-10, updated type, ESL maximum, upper ESR range, and suggested part number columns, and added a row for 0.47 µF. Updated capacitor part numbers after Figure 2: Filtering Circuit Layout. Added FSVH1924, FSVH2104, and FSVH2892 packages to Table 1-21.</p>
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Removed LPDDR3 and LPDDR4 routing guidelines from General Memory Routing Guidelines. Added note to Overview. Updated Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints, Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints, Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints, Table 43: DDR4 DIMM Skew Constraints, Table 54: LPDDR4 without ECC Skew Constraints, Table 67: LPDDR4 with ECC Skew Constraints, and Table 79: LPDDR3 SDRAM Memory Skew Constraints. Added dm[9:0] to Table 28: DDR3 SDRAM Interface Signal Description and remove note 1. Added note about ODT pins and removed ODT0/1 from FPGA Pins column in Table 45: Signal Connection Matrix for x32 Dual Channel LPDDR4 SDP without ECC, Table 46: Signal Connection Matrix for x32 Dual Channel LPDDR4 DDP without ECC, Table 55: Signal Connection Matrix for x32 Dual Channel LPDDR4 Single-Rank with ECC, and Table 56: Signal Connection Matrix for x32 Dual Channel LPDDR4 Dual-Rank with ECC.</p>
	<p>Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Added -2LI to first sentence in Separate V_{CCINT} and V_{CCBRAM}/V_{CCINT_IO} (-1LI, -2LI, -2LE). Added FSVE1156, FSVG1517, FSVE1156, FSVG1517, FSVE1156, FSVG1517, and FSVF1760 packages to Table 3-1 and Table 3-2. Updated description of RF signal chain after Table 124: Balun Specification Recommendations. Updated guard traces bullet in Isolation Recommendations. Added Figure 116: Balun Isolation with GND Stitching and GND Plane, Figure 117: RF Shield Solder Mask, Figure 125: P/N Trace Pair with 180-Degree Bend, Figure 126: Via Configuration for Signal Layer Transition, and Figure 127: Signal Trace Routing. Updated SYSREF. Added sentence about separate power supplies for ADC and DAC supplies to Power Regulation and Decoupling for ADC and DAC Supplies.</p>
	<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Replaced VCCO_MIO0 with VCCO_PSIO[0]. Removed bullet about 2.00 kΩ from SPI.</p>

Section	Revision Summary
04/10/2018 Version 1.12.1	
Appendix B: Memory Derating Tables: Removed ± sign from all table entries.	
04/09/2018 Version 1.12	
Chapter 1: Power Distribution System in UltraScale Devices: Added PCB Decoupling Capacitors for Virtex UltraScale+ High Bandwidth Memory Devices.	
Chapter 4: PCB Guidelines for Zynq UltraScale+ RFSocS: Added new chapter.	
Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: In last bullet of SD/SDIO, changed CMD3 to DAT3. Updated second paragraph in Unconnected V _{CCO} Pins.	
Appendix C: Material Properties and Insertion Losses: Added new appendix.	
Appendix D: Additional Resources and Legal Notices: Added DS926 and PG269 to References.	
02/22/2018 Version 1.11	
General updates	Reorganized content between Chapter 1: Power Distribution System in UltraScale Devices, Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC, Chapter 7: Migration between UltraScale Devices and Packages, and Chapter 8: Migration between UltraScale and UltraScale+ FPGAs.
Chapter 1: Power Distribution System in UltraScale Devices: Updated first paragraph in Recommended PCB Capacitors per Device and added bullet item to resource usage list. Added ideal value of 0.47 μF to Table 1-14. Added Recommended Decoupling Capacitor Quantities for UltraScale and UltraScale+ Devices, Recommended Decoupling Capacitor Quantities for Artix UltraScale+, Kintex UltraScale+, and Virtex UltraScale+ Devices, and Recommended Decoupling Capacitor Quantities for Zynq UltraScale+ Devices. Moved and updated Power Supply Consolidation Solutions for Zynq UltraScale+ MPSoCs from Chapter 8: Migration between UltraScale and UltraScale+ FPGAs. Updated Table 8: Number of Power Rails for Typical Configurations for Application Requiring Full Power Domain Flexibility.	

Section	Revision Summary
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Rewrote Overview. In Table 9: Reference Stackup, updated descriptions of L2, L4, L13, and L15 layers. In General Memory Routing Guidelines, removed description of maximum routing length in guideline 2, and updated guidelines 5, 6, 14, and 23. Added Figure 21: Simulations With and Without Ground Stitching Vias. Reversed order of PCB Guidelines for DDR3/3L SDRAM (PL and PS) and PCB Guidelines for DDR4 SDRAM (PL and PS) sections. Removed Table 2-11: PCB Guidelines for DDR4 SDRAM and Table 2-22: PCB Guidelines for DDR3 SDRAM. In Table 29: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals, Table 31: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals, Table 20: DDR4 SDRAM Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals, Table 24: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals, Table 36: Impedance, Length, and Spacing Guidelines for One-Slot DIMM Address, Command, and Control Signals to Table 38: Impedance, Length, and Spacing Guidelines for DIMM Data Signals, updated L0 and L1 trace lengths, and added note 2. In Table 30: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals, Table 22: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals, and Table 35: Impedance, Length, and Spacing Guidelines for DIMM Clock Signals, updated L0 and L1 trace lengths, and note 2. Updated skew constraints for signal group data to DQS in Table 33: DDR3 SDRAM Data Group Skew Constraints, Table 26: DDR4 SDRAM Data Group Skew Constraints, Table 40: DDR3 DIMM Data Group Skew Constraints, and Table 43: DDR4 DIMM Skew Constraints. Updated note after Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints. Added alert_n to Table 17: DDR4 SDRAM I/O Signal Description. Added note to Fly-by and Clamshell Topologies. Added note to Table 18: Allowable Mirror Pins for DDR4 SDRAM. Added reset_n. Updated Figure 32: Address/Command/Control Topology for DDR4 SDRAM Clamshell. In Table 21: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Address/Command/Control Signals and Table 23: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Clock Signals, updated L0 and L1 trace lengths, and added note 1. Updated note after Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints, Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints, Table 44: DDR4 DIMM Address, Command, and Control Skew Constraints, Table 88: RLD RAM 3 Memory Skew Constraints, Table 95: QDR II+ SRAM Skew Constraints, and Table 100: QDR-IV SRAM Skew Constraints. Added Table 55: Signal Connection Matrix for x32 Dual Channel LPDDR4 Single-Rank with ECC. Added Signal Segment column to Table 79: LPDDR3 SDRAM Memory Skew Constraints. Updated L1 trace length in Table 81: Clamshell Routing Impedance, Length, and Spacing Guidelines for Address and Command Signals, Table 82: RLD RAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals, Table 84: RLD RAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Clock Signals to Table 86: RLD RAM 3 Memory Impedance, Length, and Spacing Guidelines for Data Signals, Table 93: QDR II+ SRAM Impedance, Length, and Spacing Guidelines for Clock and Data Signals, Table 97: QDR-IV SRAM Impedance, Length, and Spacing Guidelines for DQ/Address/Command Signals, and Table 98: QDR-IV SRAM Impedance, Length, and Spacing Guidelines for CK/DK/QK Signals. Added VREFOA and VREFDQ. Removed instances of <i>length constraints</i> from introductory text in RLDRAM 3 Memory Routing Constraints, including Table 87: RLD RAM 3 Memory Maximum Delay Constraints. Removed skew constraints (mil) column from Table 88: RLD RAM 3 Memory Skew Constraints. Added PCB Guidelines for LPDDR4 Memories without ECC (PS), PCB Guidelines for LPDDR4 Memories with ECC (PS), and PCB Guidelines for LPDDR3 SDRAM (PL and PS) from Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC.</p>
	<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Renamed chapter title. Removed section <i>Design Example for x32 LPDDR4 with ECC</i>. In CAN, changed PCB and package skew to ± 100 ps. Added bullet about series resistor to eMMC. In Standard and High-Speed SDR Interfaces, updated heading and changed PCB and package skew to ± 100 ps. Updated heading in HS200 (200 MHz) and High-Speed DDR Interfaces. Removed TDO from JTAG. In SDR Mode, changed PCB and package skew to ± 100 ps. In DDR Mode (100 MHz), updated second bullet and added third bullet. In PS Reset (External System Reset and POR Reset), replaced VCCO_PSI0[3] with VCCO_MIO0 in first bullet, and added two new bullets. Updated first and second bullets in PS_INIT_B, PS_PROG_B, and PS_DONE. Updated fifth bullet in QSPI. Updated first bullet in Real-Time Clock. In SD/SDIO, updated first bullet, removed bullet about PCB and package delay skew, and added bullet about 10 kΩ pull-up resistor. In SPI, Triple Time Counter, and Watchdog Timer, changed PCB and package skew to ± 100 ps. Removed Trace B section. In UART, changed MIO trace delay to 1.30 ns and removed bullet about matching TX line. In USB 2.0, updated PCB and package delays, and added bullet about 30Ω series resistor. Added recommendation to tie pins to ground in Table 131: PCB Design Checklist for PS-GTR.</p>
	<p>Chapter 8: Migration between UltraScale and UltraScale+ FPGAs: Renamed chapter title. Updated 7. Memory Interface PCB Routing, including Table 140: DDR4 SDRAM Performance in UltraScale+ FPGAs in the Presence of Migration for Single-Rank Component</p>
	<p>Chapter 9: Migration between Zynq UltraScale+ MPSoCs and Packages: Added -2E speed grade to Table 142: MPSoC PL Voltage Matrix by Speed/Temperature Grade. Added XCZU2EG/XCZU3EG to first sentence in 4. -3 Speed Migration. Updated note in 5. VCU Migration. Updated Figure 184: VCCINT_VCU Connection Options for CG/EG/EV Migration (Schematic). Removed Figure 6-8: VCCINT_VCU Connection Options for CG/EG/EV Migration (PCB). Updated first two sentences in Example: Schematic/PCB Options to Enable VCU Migration from EV to CG/EG Devices. Removed section 7. Application Processor and Real-Time Processor Migration.</p>
	<p>Chapter 12: PCB Technology Basics: Replaced “Kintex UltraScale and Virtex UltraScale FPGAs” with “UltraScale architecture-based devices.” In Table 151: EDA Tools for PDS Design and Simulation, updated ADS and Siwave vendors, replaced Specctraquest Power Integrity with Sigrity, and removed Speed 2000, PowerSI, PowerDC.</p>

Section	Revision Summary
<p>Appendix B: Memory Derating Tables: Added ± to all entries in Table 157: DDR3 Data to DQS Skew Limit, Table 165: QDR II+ D to K_P/N Skew Limit to Table 169: QDR II+ ADDR/CMD to CLK Skew Limit for Clock Point-to-Point and CA Fly-by, and Table 171: QDR II+ ADDR/CMD to CLK Skew Limit for Clock T-Branch and CA T-Branch. Updated entries in Table 159: DDR4 Data to DQS Skew Limit to Table 164: RLDRAM 3 ADDR/CMD to CK_P/N Skew Limit.</p>	
01/30/2017 Version 1.10	
<p>General updates</p>	<p>Replaced “midpoint” with “midrange” throughout.</p>
<p>Chapter 1: Power Distribution System in UltraScale Devices: Added Figure 1. Added RLD1517 and RLF1924 packages for XQKU115 devices to Table 1-2. Updated V_{CCINT} values in Table 1-21.</p>	
<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Updated guideline 12 in General Memory Routing Guidelines. Added note after Figure 42: CMD/ADDR Fly-by Routing for DDR3/DDR4 DIMM Topology, Figure 44: DQS Fly-by Routing for DDR3/DDR4 DIMM Topology and Figure 45: DQ Fly-by Routing for DDR3/DDR4 DIMM Topology.</p>	
<p>Chapter 8: Migration between UltraScale and UltraScale+ FPGAs: Replaced HR I/O with HD I/O in Table 4-1. Added <i>Migration between UltraScale and UltraScale+ FPGAs</i>.</p>	
<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Replaced HR I/O with HD I/O in Table 5-1. Added V_{CCINT_VCU} Decoupling Capacitor Recommendations. Added V_{CCO_PSDDR} to Table 5-3. Added description of V_{CCINT_VCU} to Video Codec Unit (EV Devices Only). Updated PCB Guidelines for DDR3 and DDR4 SDRAM. Added DQ bus width to Figure 5-8. Updated third bullet about PCB and package delay skew in SD/SDIO. Updated list of filters for PS_MGTRAVCC and PS_MGTRAVTT in Table 131: PCB Design Checklist for PS-GTR.</p>	
<p>Chapter 9: Migration between Zynq UltraScale+ MPSoCs and Packages: Added new chapter.</p>	
<p>Appendix B: Memory Derating Tables: Replaced Mb/s with MHz as units for memory component rating in Table 165: QDR II+ D to K_P/N Skew Limit to Table 171: QDR II+ ADDR/CMD to CLK Skew Limit for Clock T-Branch and CA T-Branch.</p>	
<p>Appendix D: Additional Resources and Legal Notices: Added <i>Zynq UltraScale+ Device Technical Reference Manual (UG1085)</i>, <i>Managing Power and Performance with the Zynq UltraScale+ MPSoC (WP482)</i>, <i>UltraScale Architecture System Monitor User Guide (UG580)</i>, <i>Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)</i>, <i>UltraScale+ Devices Integrated Block for PCI Express LogiCORE IP Product Guide (PG213)</i>, <i>UltraScale+ Devices Integrated 100G Ethernet Subsystem Product Guide (PG203)</i>, and <i>Interlaken 150G LogiCORE IP Product Guide (PG212)</i> to References.</p>	
11/18/2016 Version 1.9	
<p>Chapter 1: Power Distribution System in UltraScale Devices: Updated chapter title. Updated resource usage bulleted list in Recommended PCB Capacitors per Device. Added Step Load Assumptions, including Table 1, and <i>Maximum Current Draw for V_{CCINT} in UltraScale+ Devices</i>. Added RBA676 and RFA1156 packages to Table 1-2.</p>	
<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Updated guideline 14 in General Memory Routing Guidelines. Added Figure 20: Example of Ground Stitching (Clamshell) Red: Power, Green: Ground and Figure 23: Schematic Example of VTT Resistor and Capacitor Connections. Distinguished between fly-by and clamshell topology in Overview. Added CS0_n, CS1_n to Table 17: DDR4 SDRAM I/O Signal Description. Added Fly-by and Clamshell Topologies and Using Address Mirroring to Ease Clamshell Routing. Updated first sentence in DDR4 SDRAM Address, Command, and Control Fly-by and Clamshell Topologies. Added “fly-by” to title of Table 20: DDR4 SDRAM Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals. Added Figure 32: Address/Command/Control Topology for DDR4 SDRAM Clamshell, Figure 34: Clock Topology for DDR4 SDRAM Clamshell, Table 21: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Address/Command/Control Signals, and Table 23: DDR4 SDRAM Clamshell Impedance, Length, Width, and Spacing Guidelines for Clock Signals. Updated DDR4 SDRAM Data Signals Point-to-Point for Fly-by and Clamshell Configurations.</p>	
<p>Chapter 7: Migration between UltraScale Devices and Packages: Updated recommendations in 15. Power Supplies and Thermal Considerations.</p>	
<p>Chapter 8: Migration between UltraScale and UltraScale+ FPGAs: Updated chapter title. Added notes 4 through 7 to Table 4-1. Added XCVU9P-FSGD2104, XCVU11P-FLGD2104, and XCVU13P-FIGD2104 devices, and notes 4 through 7 to Table 4-2. Added Table 138: Corresponding Half-Banks and Full-Banks. Removed sentence about UltraScale+ FPGAs having stricter ESD handling requirements than UltraScale FPGAs from 10. ESD Requirements.</p>	
<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Updated chapter title. Updated Table 5-1. Added LPDDR4 Address Copy, page 142, LPDDR4 Address Copy, page 149, and LPDDR3 Address Copy, page 161. Removed bullet about PCB and package delay skew from SD/SDIO.</p>	
07/18/2016 Version 1.8	
<p>Chapter 1: Power Distribution System in UltraScale Devices: Added note 2 to Table 1-2 and Table 1-3.</p>	

Section	Revision Summary
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Clarified guideline 5 in General Memory Routing Guidelines. Replaced CD with CK in Figure 14: Interpretation of CK to DQS Skew Specification title. Added important note before Figure 19: Example of Ground Stitching (Fly-by). Replaced reset_b with reset_n throughout. Updated reset_n, and reset_n, including removal of Figure 2-17 and Figure 2-22. Removed note about package delay (P0) from Table 33: DDR3 SDRAM Data Group Skew Constraints, Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints, Table 26: DDR4 SDRAM Data Group Skew Constraints, Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints, Table 40: DDR3 DIMM Data Group Skew Constraints, Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints, Table 43: DDR4 DIMM Skew Constraints, Table 44: DDR4 DIMM Address, Command, and Control Skew Constraints, Table 88: RDRAM 3 Memory Skew Constraints, Table 95: QDR II+ SRAM Skew Constraints, and Table 100: QDR-IV SRAM Skew Constraints. Added we, ras, cas, and TEN to Table 17: DDR4 SDRAM I/O Signal Description. Removed L2 column from Table 35: Impedance, Length, and Spacing Guidelines for DIMM Clock Signals. In Table 36: Impedance, Length, and Spacing Guidelines for One-Slot DIMM Address, Command, and Control Signals, removed L2 column, and updated L1 parameters. Added Table 37: Impedance, Length, and Spacing Guidelines for Two-Slot DIMM Address, Command, and Control Signals. Replaced dm with dk/qk in Table 85: RDRAM 3 Memory Impedance, Length, and Spacing Guidelines for dk and qk Signals.</p>
	<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added Zynq UltraScale+ MPSoCs and V_{CCINT}/V_{CCINT_IO} at 0.47 μF to Table 5-1. Added note 2 to Table 5-3. Added <i>Video Codec Unit (EV Devices Only)</i>. Added reset_n to Table 5-5 and Table 5-14. Updated Figure 5-3. In Table 5-11 and Table 5-13, removed note about package delay (P0), updated skew constraints (mil), and removed CK to DQS (A)/(B) row. Updated L2 spacing to other group signals in Table 5-4 to Table 5-7, and Table 5-11. Updated Figure 5-4. Added <i>PCB Guidelines for LPDDR3 SDRAM</i>. Updated recommendations for PS_MGTRREF and PS_MGTRAVTT in Table 131: PCB Design Checklist for PS-GTR.</p>
	<p>Chapter 7: Migration between UltraScale Devices and Packages: Added recommendation to use Delphi thermal model to 15. Power Supplies and Thermal Considerations.</p>
	<p>Chapter 8: Migration between UltraScale and UltraScale+ FPGAs: Removed PCI Express from UltraScale+ FPGA Migration Checklist.</p>
	<p>Appendix B: Memory Derating Tables: Updated second paragraph at start of appendix. Updated cells in Table 157: DDR3 Data to DQS Skew Limit to Table 160: DDR4 Address/Command/Control to CK Skew Limit, added notes to all tables, and updated titles of Table 158: DDR3 Address/Command/Control to CK Skew Limit and Table 160: DDR4 Address/Command/Control to CK Skew Limit.</p>
<p>04/26/2016 Version 1.7</p>	
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: In General Memory Routing Guidelines, updated guidelines 10 and 21, and added guidelines 11 and 23. Updated paragraph after Table 12: Example CK to DQS Skew Constraint. Removed DIMM from Figure 14: Interpretation of CK to DQS Skew Specification. Added introductory sentences for Figure 2-17 and Figure 37: Termination for reset_n in DDR3 DRAM. In Table 28: DDR3 SDRAM Interface Signal Description, updated descriptions of cke[1:0] and reset_b. Added heading to and updated reset_n. Updated introductory sentence before Figure 39: Data Signals Point-to-Point for DDR3 SDRAM and Figure 35: Data Signals Point-to-Point for DDR4 SDRAM. Added note 3 to Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints, Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints, and Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints. In Table 17: DDR4 SDRAM I/O Signal Description, updated descriptions of cke and reset_b. Added notes about pin rules and swapping after Table 28: DDR3 SDRAM Interface Signal Description and Table 17: DDR4 SDRAM I/O Signal Description. Added $\pm 10\%$ to all entries for Z_{DIFF} in Table 22: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals and Table 24: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals.</p>
	<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added $V_{CC_PSDDR_PLL}$ Supply. Updated <i>PCB Guidelines for DDR3 and DDR4 SDRAM</i>. Updated L2 column in Table 5-5, Table 5-6, Table 5-7, Table 5-8, Table 5-9, Table 5-8, and Table 5-11. Updated signal names and connections in Figure 5-3, Figure 5-4, and Figure 5-8. Made CK routing pull up to V_{TT} in Figure 5-3. Updated DQ width in Figure 5-4. Updated Table 5-11 and Table 5-14. Added channel B to ECC unit in Figure 5-3 and Figure 5-4. Made ending via pull up to V_{TT} in Figure 5-4. Added channel A label to Figure 5-5. Updated ECC unit in Figure 5-6. Updated trace lengths in Table 5-5 through Table 5-11. Updated L3 Z_{diff} in Table 5-9. Replaced DRAM with memory device in Figure 5-7 and Figure 5-8. Updated L1 spacing in Table 5-10. Updated Table 5-13. Added <i>PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC</i>.</p>
	<p>Chapter 7: Migration between UltraScale Devices and Packages: Updated 15. Power Supplies and Thermal Considerations.</p>
	<p>Appendix B: Memory Derating Tables: Added notes to Table 159: DDR4 Data to DQS Skew Limit and Table 160: DDR4 Address/Command/Control to CK Skew Limit.</p>
	<p>Appendix D: Additional Resources and Legal Notices: Added <i>Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics (DS925)</i> to References.</p>

Section	Revision Summary
02/04/2016 Version 1.6	
<p>Chapter 1: Power Distribution System in UltraScale Devices: Updated paragraph after resource usage bullets in Recommended PCB Capacitors per Device. In Table 1-2, added SFVA784 package for XCKU035 and XCKU040, and FFVA1156 package for XCKU095. Added footnote with definition of electrical inch to <i>0805 Ceramic Capacitor</i>.</p>	
<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Updated Overview. Updated note in Table 9: Reference Stackup. Significant revision to the General Memory Routing Guidelines section including adding Table 10: Example Signal to Signal Skew Constraint through Table 12: Example CK to DQS Skew Constraint and Figure 12: Skew Example for DQ to DQS to Figure 14: Interpretation of CK to DQS Skew Specification. Added the Adjusting for Different Stack-Ups section. Updated cke[1:0] and reset_b descriptions in Table 28: DDR3 SDRAM Interface Signal Description and Table 17: DDR4 SDRAM I/O Signal Description. Added Figure 2-17, Figure 37: Termination for reset_n in DDR3 DRAM, Figure 2-21, and Figure 31: Termination for reset_n in DDR4 DRAM. Replaced strobe with dqs in Table 31: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals, Table 24: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Data Signals, and Table 38: Impedance, Length, and Spacing Guidelines for DIMM Data Signals. In Table 33: DDR3 SDRAM Data Group Skew Constraints, Table 26: DDR4 SDRAM Data Group Skew Constraints, Table 40: DDR3 DIMM Data Group Skew Constraints, and Table 43: DDR4 DIMM Skew Constraints replaced strobe with DQS and clock with CK and updated note 1. Added new rows and updated note 1 and 2 in Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints, Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints, Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints, and Table 44: DDR4 DIMM Address, Command, and Control Skew Constraints. In Table 40: DDR3 DIMM Data Group Skew Constraints, updated CK to DQS constraints and added notes 2 and 5. In Table 43: DDR4 DIMM Skew Constraints, updated CK to DQS constraints and added note 2. Updated note 1 in Table 88: RLD RAM 3 Memory Skew Constraints, Table 95: QDR II+ SRAM Skew Constraints, and Table 100: QDR-IV SRAM Skew Constraints. Removed PCB Guidelines for LPDDR3 Memories section.</p>	
<p>Chapter 8: Migration between UltraScale and UltraScale+ FPGAs: In Table 4-1, removed XCKU7P in all packages, and added XCKU3P-FFVB676, XCKU3P-FFVD900, XCKU5P-SFVB784, XCKU5P-FFVA676, and XCKU5P-FFVD900. Updated packages for XCVU9P, XCVU11P, and XCVU13P in Table 4-2. Added A676 and A1156 packages to Table 135: UltraScale to UltraScale+ FPGA Migration Paths by Package.</p>	
<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added XCZU15EG to Table 5-1. In Table 5-3, replaced VCCPSGTA with PS_MGTRAVCC, VCCPSDDR with VCC_PSINTFP_DDR, and added underscores to other column headings. Moved PCB Guidelines for LPDDR4 Memories without ECC from Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices and expanded section. Added PCB Guidelines for LPDDR4 Memories with ECC.</p>	
11/24/2015 Version 1.5	
<p>Chapter 5: PCB Guidelines for the PS Interface in the Zynq UltraScale+ MPSoC: Added Zynq UltraScale+ MPSoC information.</p>	
<p>Chapter 8: Migration between UltraScale and UltraScale+ FPGAs: Added UltraScale+ FPGA information.</p>	
09/16/2015 Version 1.4	
<p>Chapter 1: Power Distribution System in UltraScale Devices: Added V_{CCO} (Bank 0) column to Table 1-2 and Table 1-3.</p>	

Section	Revision Summary
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Added QDR-IV SRAM to chapter title and Overview. In General Memory Routing Guidelines, updated second, fifth, seventh, and eighth guidelines, and added third, fourth, ninth and tenth guidelines. Updated Figure 18: Breakout Region Routing. Added <code>reset_b</code> to control signals in Table 28: DDR3 SDRAM Interface Signal Description and Table 17: DDR4 SDRAM I/O Signal Description. Updated L3 trace length in Table 29: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals, Table 30: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals, Table 20: DDR4 SDRAM Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals, Table 22: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals, Table 82: RLD RAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals, Table 84: RLD RAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Clock Signals, and Table 90: QDR II+ SRAM Impedance, Length, and Spacing Guidelines for Address and Command Signals. Removed selected L2 breakouts from Figure 36: Address Fly-by Termination for DDR3 SDRAM, Figure 38: Clock Fly-by Termination for DDR3 SDRAM, Figure 30: Address, Command, and Control Fly-by Termination for DDR4 SDRAM, Figure 33: Clock Fly-by Termination for DDR4 SDRAM, Figure 71: Address, Command, and Control Fly-by Termination for RLD RAM 3 Memories, and Figure 79: Address and Command Fly-by Routing and Termination for QDR II+ SRAM. Updated table title and maximum length/delay constraints column in Table 32: DDR3 SDRAM Total Length/Delay Constraints, Table 25: DDR4 SDRAM Total Length/Delay Constraints, Table 39: DDR3 DIMM Total Length/Delay Constraints, Table 42: DDR4 DIMM Total Length/Delay Constraints, Table 87: RLD RAM 3 Memory Maximum Delay Constraints, Table 94: QDR II+ SRAM Maximum Length/Delay Constraints, Table 99: QDR-IV SRAM Maximum Length/Delay Constraints, and Table 2-58. Added skew constraints (mil) column and added table note to Table 33: DDR3 SDRAM Data Group Skew Constraints, Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints, Table 26: DDR4 SDRAM Data Group Skew Constraints, Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints, Table 40: DDR3 DIMM Data Group Skew Constraints, Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints, Table 43: DDR4 DIMM Skew Constraints, Table 44: DDR4 DIMM Address, Command, and Control Skew Constraints, Table 88: RLD RAM 3 Memory Skew Constraints, Table 95: QDR II+ SRAM Skew Constraints, Table 100: QDR-IV SRAM Skew Constraints, and Table 2-59. Replaced device with FPGA in DDR3 SDRAM Address, Command, and Control Fly-by Termination. Updated Table 26: DDR4 SDRAM Data Group Skew Constraints title. Updated Table 41: DDR3 DIMM Address, Command, and Control Skew Constraints title and data to strobe skew constraints. Added Table 44: DDR4 DIMM Address, Command, and Control Skew Constraints. Updated Figure 68: Clamshell RLD RAM 3 Memory with Width Expansion. Removed parenthetical comments from descriptions in Table 80: RLD RAM 3 Memory Interface Signal Description and Table 89: QDR II+ SRAM Interface Signal Description. Added "qvd to qk_p/n" row to Table 88: RLD RAM 3 Memory Skew Constraints. Updated Figure 80: Address and Command Point-to-Point Routing and Termination for QDR II+ SRAM. Replaced <code>K_P/N</code> with <code>K</code> and <code>K_B</code> in QDR II+ SRAM Clock (K, K_B, and BWS) T-Branch Routing and Termination. Replaced <code>k_p/k_n</code> with <code>k/k_b</code> in QDR II+ SRAM Clock and Data Signals (d/k/k_b) Point-to-Point Routing. Replaced <code>cq_p/cq_n</code> with <code>cq/cq_b</code> in QDR II+ SRAM Clock and Data Signals (q/cq/cq_b) Point-to-Point Routing. Added signal groups "k to k_b" and "q to q_b" to Table 95: QDR II+ SRAM Skew Constraints. Added signal group "ck_p to ck_n" to Table 2-59. Added PCB Guidelines for QDR-IV SRAM.</p>
	<p>Chapter 7: Migration between UltraScale Devices and Packages: Replaced <i>pin</i> with <i>footprint</i> throughout. Updated 1. Footprint Compatibility between Packages. Replaced XCVU160 with XCVU190 in Example. Added MGTRREF and MGTAVTTRCAL.</p>
	<p>Appendix B: Memory Derating Tables: Added new appendix.</p>
03/09/2015 Version 1.3	
	<p>Chapter 1: Power Distribution System in UltraScale Devices: In Table 1-2, added XCKU035 and XCKU100 devices and removed note about 0402 2.2 μF capacitors. Added Table 1-3.</p>

Section	Revision Summary
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Removed LPDDR2 throughout. Updated fourth and fifth guidelines in General Memory Routing Guidelines. Updated L0 trace length to 0.0–0.8/1.2 in Impedance, Length, and Spacing Guidelines tables. Updated L1 trace length to Address L1 + 0.25 in Table 30: DDR3 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals, Table 22: DDR4 SDRAM Impedance, Length, and Spacing Guidelines for Clock Signals, and Table 35: Impedance, Length, and Spacing Guidelines for DIMM Clock Signals. Updated address/command/control signal group in Table 32: DDR3 SDRAM Total Length/Delay Constraints, Table 25: DDR4 SDRAM Total Length/Delay Constraints, Table 39: DDR3 DIMM Total Length/Delay Constraints, and Table 42: DDR4 DIMM Total Length/Delay Constraints. Updated skew constraints in Table 33: DDR3 SDRAM Data Group Skew Constraints, Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints, Table 26: DDR4 SDRAM Data Group Skew Constraints, Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints, Table 40: DDR3 DIMM Data Group Skew Constraints, Table 43: DDR4 DIMM Skew Constraints, Table 88: RLD RAM 3 Memory Skew Constraints, and Table 95: QDR II+ SRAM Skew Constraints. Added note about skew specifications to Table 33: DDR3 SDRAM Data Group Skew Constraints, Table 34: DDR3 SDRAM Address, Command, and Control Skew Constraints, Table 26: DDR4 SDRAM Data Group Skew Constraints, Table 27: DDR4 SDRAM Address, Command, and Control Skew Constraints, Table 43: DDR4 DIMM Skew Constraints, Table 88: RLD RAM 3 Memory Skew Constraints, Table 95: QDR II+ SRAM Skew Constraints, and Table 2-59. Updated table note referring to item 5 of General Memory Routing Guidelines. Updated skew constraints for “data to strobe” signal group in Table 26: DDR4 SDRAM Data Group Skew Constraints and Table 43: DDR4 DIMM Skew Constraints. Added row for differential signals to Table 88: RLD RAM 3 Memory Skew Constraints. Updated all trace lengths in Table 90: QDR II+ SRAM Impedance, Length, and Spacing Guidelines for Address and Command Signals. Updated L1 trace length to 1.0–4.0 in Table 93: QDR II+ SRAM Impedance, Length, and Spacing Guidelines for Clock and Data Signals. Updated Figure 68: Clamshell RLD RAM 3 Memory with Width Expansion. Updated title of Figure 71: Address, Command, and Control Fly-by Termination for RLD RAM 3 Memories. Removed “R = 50Ω” label from Figure 77: Design Example for QDR II+ SRAM Devices with Width Expansion. Added pull-up resistor to CA line in Figure 78: Design Example for Single QDR II+ SRAM Device. In Figure 2-52 and Figure 2-53, replaced V_{TT} with V_{DDQ} and changed pull-up resistance from 120Ω to 60Ω. Removed “spacing in addr/cmd/ctrl” row from Table 2-56.</p>
	<p>Chapter 7: Migration between UltraScale Devices and Packages: Added new chapter.</p>
	<p>Chapter 12: PCB Technology Basics: Updated second paragraph of Input Thresholds. Updated second paragraph of V_{REF} Stabilization Capacitors.</p>
	<p>Appendix D: Additional Resources and Legal Notices: Updated list of documents in References.</p>
12/01/2014 Version 1.2	
	<p>Chapter 1: Power Distribution System in UltraScale Devices: Removed recommendation asking customers to run their own power integrity simulations. In Table 1-2, replaced “should” with “must” in notes 1 and 2 and added note 4. Removed Table 1-2: Virtex UltraScale Devices Power Supply Decoupling Capacitors. In Table 1-14, added row for 470 μF, and updated ESL maximum and suggested part numbers for 100 μF, 47 μF, and 4.7 μF capacitors. In PCB Bulk Capacitors, added 470 μF to first paragraph and replaced niobium oxide with polymer aluminum in second paragraph.</p>
	<p>Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices: Added sentence about inner signal layer propagation time to first paragraph of Reference Stackup. Updated Thickness column in Table 9: Reference Stackup. Added second, fifth, and sixth guidelines and Figure 11: UltraScale Device Breakout Single-Ended Routing Options to General Memory Routing Guidelines. Replaced “length matching” with “skew” throughout. Changed units in skew constraints column from mils to picoseconds and updated notes in all skew constraints tables. Replaced “routing length” with “flight times” in Important notes after skew constraints tables. Removed DDR3 SDRAM ODT Settings and DDR4 SDRAM ODT Settings sections. Updated L0 trace length to 0–0.6/1.2 and added table note in all impedance, length, and spacing guidelines tables. Added DDR3/DDR4 UDIMM/RDIMM/SODIMM/LRDIMM Routing Guidelines (PL and PS). Replaced “depth” with “width” in RLDRAM 3 Memory Topology and Routing Guidelines for Clamshell and Fly-by Configurations. Updated maximum PCB via count to 6 in Table 81: Clamshell Routing Impedance, Length, and Spacing Guidelines for Address and Command Signals and Table 82: RLD RAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals. Updated L2 trace length in Table 82: RLD RAM 3 Memory Fly-by Impedance, Length, and Spacing Guidelines for Address, Command, and Control Signals to Table 86: RLD RAM 3 Memory Impedance, Length, and Spacing Guidelines for Data Signals. Updated title of Figure 74: dk_p/n Point-to-Point Routing. Added Figure 75: qk_p/n Point-to-Point Routing. In Table 86: RLD RAM 3 Memory Impedance, Length, and Spacing Guidelines for Data Signals, added Z_0 parameter and removed Z_{diff} and differential trace row. Updated DATA IN and DATA OUT ports in Figure 77: Design Example for QDR II+ SRAM Devices with Width Expansion. Added Figure 78: Design Example for Single QDR II+ SRAM Device. Added QDR II+ SRAM Topology and Routing Guidelines for T-Branch Configuration. Removed QDR II+ SRAM Clock (k and cq) Differential Point-to-Point Routing section. Added QDR II+ SRAM Clock (K, K_B, and BWS) T-Branch Routing and Termination. Added clock to QDR II+ SRAM Clock and Data Signals (d/k/k_b) Point-to-Point Routing and QDR II+ SRAM Clock and Data Signals (q/cq/cq_b) Point-to-Point Routing. Removed resistor from Figure 84: Clock and Data Signals (q/cq/cq_b) Point-to-Point Routing. Updated Table 95: QDR II+ SRAM Skew Constraints. Removed “clock” from LPDDR3 Memory Address, Command, and Control Point-to-Point Routing. Added Important note after Table 2-59.</p>
	<p>Chapter 11: SelectIO Signaling: Removed VRN from Unidirectional Topographies and Termination.</p>
	<p>Chapter 12: PCB Technology Basics: Replaced “inductance” with “impedance” in Figure 212: Effective Frequency Example.</p>

Section	Revision Summary
08/28/2014 Version 1.1	
Chapter 1: Power Distribution System in UltraScale Devices : Replaced 0603 capacitor with 0805 capacitor throughout. In Recommended PCB Capacitors per Device, added alternate network example, Recommended icon, and Table 1-2. Updated Table 1-2. In Table 1-2, updated body size and voltage rating of 4.7 μ F capacitor and removed note 3 about ESR. Updated first paragraph in <i>PCB Bulk Capacitors</i> .	
Chapter 2: PCB Guidelines for Memory Interfaces in Kintex UltraScale+, Virtex UltraScale+, and Zynq UltraScale+ Devices : Added new chapter.	
Chapter 12: PCB Technology Basics : In Noise Limits , updated first paragraph and updated low-frequency variance bullet. Updated first paragraph in Role of Inductance . Updated description of tantalum capacitors in Capacitor Parasitic Inductance . Updated Capacitor Mounting Inductance . Updated second trade-off bullet in FPGA Mounting Inductance . Updated terminology in Noise Magnitude Measurement .	
Appendix D: Additional Resources and Legal Notices : Added <i>UltraScale Architecture GTY Transceivers User Guide (UG578)</i> to References .	
12/10/2013 Version 1.0	
General updates	Initial Xilinx release.

Please Read: Important Legal Notices

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes. THIS INFORMATION IS PROVIDED "AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AUTOMOTIVE APPLICATIONS DISCLAIMER

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

Copyright

© Copyright 2013–2025 Advanced Micro Devices, Inc. AMD, the AMD Arrow logo, Artix, Kintex, Spartan, UltraScale, UltraScale+, Virtex, Vivado, Zynq, and combinations thereof are trademarks of Advanced Micro Devices, Inc. AMBA, AMBA Designer, Arm, ARM1176JZ-S, CoreSight, Cortex, PrimeCell, Mali, and MPCore are trademarks of Arm Limited in the US and/or elsewhere. The DisplayPort Icon is a trademark of the Video Electronics Standards Association, registered in the U.S. and other countries. MATLAB and Simulink are registered trademarks of The MathWorks, Inc. PCI, PCIe, and PCI Express are trademarks of PCI-SIG and used under license. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.