

Spartan-6 FPGA PCB Design Guide

UG393 (v1.0) September 21, 2009



Xilinx is disclosing this user guide, manual, release note, and/or specification (the "Documentation") to you solely for use in the development of designs to operate with Xilinx hardware devices. You may not reproduce, distribute, republish, download, display, post, or transmit the Documentation in any form or by any means including, but not limited to, electronic, mechanical, photocopying, recording, or otherwise, without the prior written consent of Xilinx. Xilinx expressly disclaims any liability arising out of your use of the Documentation. Xilinx reserves the right, at its sole discretion, to change the Documentation without notice at any time. Xilinx assumes no obligation to correct any errors contained in the Documentation, or to advise you of any corrections or updates. Xilinx expressly disclaims any liability in connection with technical support or assistance that may be provided to you in connection with the Information.

THE DOCUMENTATION IS DISCLOSED TO YOU "AS-IS" WITH NO WARRANTY OF ANY KIND. XILINX MAKES NO OTHER WARRANTIES, WHETHER EXPRESS, IMPLIED, OR STATUTORY, REGARDING THE DOCUMENTATION, INCLUDING ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT OF THIRD-PARTY RIGHTS. IN NO EVENT WILL XILINX BE LIABLE FOR ANY CONSEQUENTIAL, INDIRECT, EXEMPLARY, SPECIAL, OR INCIDENTAL DAMAGES, INCLUDING ANY LOSS OF DATA OR LOST PROFITS, ARISING FROM YOUR USE OF THE DOCUMENTATION.

© 2009 Xilinx, Inc. XILINX, the Xilinx logo, Virtex, Spartan, ISE, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. All other trademarks are the property of their respective owners.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
09/21/09	1.0	Initial Xilinx release.

Table of Contents

Revision History	2
Preface: About This Guide	
Guide Contents	5
Additional Documentation	5
Additional Support Resources	6
Chapter 1: PCB Technology Basics	
PCB Structures	7
Traces	7
Planes	7
Vias	8
Pads and Antipads	8
Lands	8
Dimensions	8
Transmission Lines	9
Return Currents	10
Chapter 2: Power Distribution System	
PCB Decoupling Capacitors	11
Recommended Capacitors per Device	11
Required PCB Capacitor Quantities	12
Capacitor Specifications	14
PCB Bulk Capacitors	14
PCB High-Frequency Capacitors	14
Capacitor Consolidation Rules	15
PCB Capacitor Placement and Mounting Techniques	15
PCB Bulk Capacitors	15
0805 Ceramic Capacitor	16
0402 Ceramic Capacitor	17
Basic PDS Principles	18
Noise Limits	18
Role of Inductance	20
Capacitor Parasitic Inductance	20
PCB Current Path Inductance	22
Capacitor Mounting Inductance	22
Plane Inductance	23
FPGA Mounting Inductance	24
PCB Stackup and Layer Order	25
Capacitor Effective Frequency	26
Capacitor Anti-Resonance	28
Capacitor Placement Background	28
V _{REF} Stabilization Capacitors	29
Power Supply Consolidation	29
Unconnected V _{CCO} Pins	30

Simulation Methods	30
PDS Measurements	31
Noise Magnitude Measurement	31
Noise Spectrum Measurements	33
Optimum Decoupling Network Design	35
Troubleshooting	35
Possibility 1: Excessive Noise from Other Devices on the PCB	35
Possibility 2: Parasitic Inductance of Planes, Vias, or Connecting Traces	35
Possibility 3: I/O Signals in PCB are Stronger Than Necessary	36
Possibility 4: I/O Signal Return Current Traveling in Sub-Optimal Paths	36

Chapter 3: SelectIO Signaling

Interface Types	37
Single-Ended versus Differential Interfaces	37
SDR versus DDR Interfaces	38
Single-Ended Signaling	38
Modes and Attributes	38
Input Thresholds	38

Chapter 4: PCB Materials and Traces

How Fast is Fast?	39
Dielectric Losses	39
Relative Permittivity	39
Loss Tangent	40
Skin Effect and Resistive Losses	40
Choosing the Substrate Material	40
Traces	41
Trace Geometry	41
Trace Characteristic Impedance Design for High-Speed Transceivers	41
Trace Routing	43
Plane Splits	43
Return Currents	43
Simulating Lossy Transmission Lines	44
Cable	44
Connectors	44
Skew Between Conductors	44

Chapter 5: Design of Transitions for High-Speed Signals

Excess Capacitance and Inductance	45
Time Domain Reflectometry	45
BGA Package	47
SMT Pads	47
Differential Vias	51
P/N Crossover Vias	54
SMA Connectors	54
Backplane Connectors	54
Microstrip/Stripline Bends	54

About This Guide

This guide provides information on PCB design for Spartan®-6 devices, with a focus on strategies for making design decisions at the PCB and interface level.

Guide Contents

This guide contains the following chapters:

- [Chapter 1, “PCB Technology Basics,”](#) discusses the basics of current PCB technology focusing on physical structures and common assumptions.
- [Chapter 2, “Power Distribution System,”](#) covers the power distribution system for Spartan-6 FPGAs, including all details of decoupling capacitor selection, use of voltage regulators and PCB geometries, simulation and measurement.
- [Chapter 3, “SelectIO Signaling,”](#) contains information on the choice of SelectIO™ standards, I/O topographies, and termination strategies as well as information on simulation and measurement techniques.
- [Chapter 4, “PCB Materials and Traces,”](#) provides some guidelines on managing signal attenuation to obtain optimal performance for high-frequency applications.
- [Chapter 5, “Design of Transitions for High-Speed Signals,”](#) addresses the interface at either end of a transmission line. The provided analyses and examples can greatly accelerate the specific design.

Additional Documentation

The following documents are also available for download at <http://www.xilinx.com/support/documentation/spartan-6.htm>.

- **Spartan-6 Family Overview**
This overview outlines the features and product selection of the Spartan-6 family.
- **Spartan-6 FPGA Data Sheet: DC and Switching Characteristics**
This data sheet contains the DC and switching characteristic specifications for the Spartan-6 family.
- **Spartan-6 FPGA Packaging and Pinout Specifications**
This specification includes the tables for device/package combinations and maximum I/Os, pin definitions, pinout tables, pinout diagrams, mechanical drawings, and thermal specifications.
- **Spartan-6 FPGA Configuration User Guide**

This all-encompassing configuration guide includes chapters on configuration interfaces (serial and parallel), multi-bitstream management, bitstream encryption, boundary-scan and JTAG configuration, and reconfiguration techniques.

- **Spartan-6 FPGA SelectIO Resources User Guide**
This guide describes the SelectIO™ resources available in all Spartan-6 devices.
- **Spartan-6 FPGA Clocking Resources User Guide**
This guide describes the clocking resources available in all Spartan-6 devices, including the DCMs and PLLs.
- **Spartan-6 FPGA Configurable Logic Block User Guide**
This guide describes the capabilities of the configurable logic blocks (CLBs) available in all Spartan-6 devices.
- **Spartan-6 FPGA Block RAM Resources User Guide**
This guide describes the Spartan-6 device block RAM capabilities.
- **Spartan-6 FPGA DSP48A1 Slice User Guide**
This guide describes the architecture of the DSP48A1 slice in Spartan-6 FPGAs and provides configuration examples.
- **Spartan-6 FPGA Memory Controller User Guide**
This guide describes the Spartan-6 FPGA memory controller block, a dedicated embedded multi-port memory controller that greatly simplifies interfacing Spartan-6 FPGAs to the most popular memory standards.
- **Spartan-6 FPGA GTP Transceiver User Guide**
This guide describes the GTP transceivers available in the Spartan-6 LXT FPGAs.

Additional Support Resources

To search the Answer Database of silicon, software, and IP questions and answers, or to create a technical support WebCase, see the Xilinx website at:

<http://www.xilinx.com/support/mysupport.htm>.

PCB Technology Basics

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.

This chapter contains the following sections:

- [PCB Structures](#)
- [Transmission Lines](#)
- [Return Currents](#)

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 microvia 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 pin or ball pitch (1.0 mm for FG 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 for array packages).
- 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.

PCBs for larger FPGAs can range from 4 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 2, Power Distribution System](#). Termination resistor placement constraints must be determined through signal integrity simulation, using IBIS or SPICE.

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 complimentary 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.

Power Distribution System

This chapter documents the power distribution system (PDS) for Spartan®-6 FPGAs, including decoupling capacitor selection, placement, and PCB geometries. A simple decoupling method is provided for each device in the Spartan-6 family. Basic PDS design principles are covered, as well as simulation and analysis methods. This chapter contains the following sections:

- [PCB Decoupling Capacitors](#)
- [Basic PDS Principles](#)
- [Simulation Methods](#)
- [PDS Measurements](#)
- [Troubleshooting](#)

PCB Decoupling Capacitors

Recommended Capacitors per Device

A simple PCB-decoupling network for each Spartan-6 device is listed in [Table 2-1](#).

Decoupling methods other than those presented in [Table 2-1](#) can be used, but the decoupling network should be designed to meet or exceed the performance of the simple decoupling networks presented here. The impedance of the alternate network must be less than or equal to that of the recommended network across frequencies from 100 KHz to 500 MHz.

Because device capacitance requirements vary with CLB and I/O utilization, PCB decoupling guidelines are provided on a per-device basis. V_{CCINT} and V_{CCAUX} capacitors are listed as the quantity per device, while V_{CCO} capacitors are listed as the quantity per I/O bank. Device performance at full utilization is equivalent across all devices when using these recommended networks.

Required PCB Capacitor Quantities

Table 2-1 lists the PCB decoupling capacitor guidelines per V_{CC} supply rail. A dash indicates that data is not currently available.

Table 2-1: Required PCB Capacitor Quantities per Device⁽¹⁾

Package	Device (XC6S)	V_{CCINT} in μF			V_{CCAUX} in μF			V_{CC0} Bank 0 in μF			V_{CC1} Bank 1 in μF			V_{CC2} Bank 2 in μF			V_{CC3} Bank 3 in μF			V_{CC4} Bank 4 in μF			V_{CC5} Bank 5 in μF			Total (2)
		100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	
TQG144	LX4	0	5	1	0	7	3	1	1	2	1	1	2	1	1	2	1	1	2							33
TQG144	LX9	1	1	2	0	7	3	1	1	2	1	1	2	1	1	2	1	1	2							32
CPG196	LX4	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CPG196	LX9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CPG196	LX16	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CSG225	LX4	0	5	1	1	2	4	1	2	3	1	2	3	1	2	3	1	2	3							35
CSG225	LX9	1	1	2	1	2	4	1	2	3	1	2	3	1	2	3	1	2	3							33
CSG225	LX16	1	2	4	1	2	4	1	1	3	1	1	3	1	1	3	1	1	3							35
CSG324	LX9	1	1	2	1	3	5	1	2	3	1	2	4	1	2	3	1	2	4							38
CSG324	LX16	1	2	4	1	3	6	1	2	4	1	2	4	1	2	4	1	2	4							44
CSG324	LX25	1	3	7	1	3	6	1	2	4	1	2	4	1	3	4	1	2	4							48
CSG324	LX25T	1	3	7	1	2	5	1	1	1	1	2	4	1	3	4	1	2	4							43
CSG324	LX45	2	6	12	1	3	6	1	2	3	1	2	4	1	2	4	1	2	4							55
CSG324	LX45T	2	6	12	1	2	5	1	1	1	1	2	4	1	2	4	1	2	4							51
FT(G)256	LX9	1	1	2	1	2	5	1	2	3	1	2	4	1	2	3	1	2	4							36
FT(G)256	LX16	1	2	4	1	2	5	1	1	3	1	2	4	1	1	3	1	2	4							38
FT(G)256	LX25	1	3	7	1	2	5	1	2	3	1	2	4	1	2	3	1	2	4							43
CSG484	LX45	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CSG484	LX45T	2	6	12	1	4	7	1	2	3	1	3	6	1	3	6	1	3	6							67
CSG484	LX75	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CSG484	LX75T	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CSG484	LX100	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CSG484	LX100T	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CSG484	LX150	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
CSG484	LX150T	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
FG(G)484	LX25	1	3	7	1	3	7	1	2	4	1	3	4	1	3	6	1	3	4							54
FG(G)484	LX25T	1	3	7	1	3	6	1	2	3	1	3	4	1	3	6	1	3	4							52
FG(G)484	LX45	2	6	12	1	4	8	1	2	3	1	3	6	1	4	7	1	3	6							70
FG(G)484	LX45T	2	6	12	1	4	7	1	2	3	1	3	6	1	3	6	1	3	6							69

Table 2-1: Required PCB Capacitor Quantities per Device⁽¹⁾ (Continued)

Package	Device (XC6S)	V _{CCINT} in μ F			V _{CCAUX} in μ F			V _{CCO} Bank 0 in μ F			V _{CCO} Bank 1 in μ F			V _{CCO} Bank 2 in μ F			V _{CCO} Bank 3 in μ F			V _{CCO} Bank 4 in μ F			V _{CCO} Bank 5 in μ F			Total (2)
		100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	100	4.7	0.47	
FG(G)484	LX75	3	10	20	1	4	9	1	3	5	1	3	6	1	4	7	1	4	6							88
FG(G)484	LX75T	3	10	20	1	4	7	1	2	3	1	3	6	1	3	6	1	3	6							82
FG(G)484	LX100	4	13	27	1	4	8	1	3	5	1	3	6	1	4	6	1	4	6							97
FG(G)484	LX100T	4	13	27	1	4	7	1	2	3	1	3	6	1	3	6	1	3	6							93
FG(G)484	LX150	7	20	40	1	4	8	1	2	5	1	3	6	1	3	6	1	3	6							118
FG(G)484	LX150T	7	20	40	1	4	7	1	2	3	1	3	6	1	3	6	1	3	6							114
FG(G)676	LX45	2	6	12	1	5	9	1	3	5	1	3	6	1	3	6	1	4	8							76
FG(G)676	LX75	3	10	20	2	6	12	1	4	6	1	4	6	1	4	7	1	4	7	1	2	4	1	2	4	111
FG(G)676	LX75T	3	10	20	1	5	10	1	3	6	1	2	4	1	3	6	1	3	5	1	2	4	1	2	4	97
FG(G)676	LX100	4	13	27	2	6	12	1	4	6	1	4	6	1	4	7	1	4	7	1	2	4	1	2	4	122
FG(G)676	LX100T	4	13	27	1	5	10	1	3	5	1	2	4	1	3	5	1	3	5	1	2	3	1	2	4	106
FG(G)676	LX150	7	20	40	2	6	12	1	3	6	1	3	6	1	3	7	1	3	7	1	2	4	1	2	4	142
FG(G)676	LX150T	7	20	40	1	5	10	1	3	5	1	2	4	1	3	5	1	2	5	1	2	3	1	2	4	126
FG(G)900	LX100T	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
FG(G)900	LX150	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
FG(G)900	LX150T	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Notes:

1. PCB Capacitor specifications are listed in [Table 2-2](#).
2. Total includes all capacitors for all supplies, accounting for the number of I/O banks in the device.

Capacitor Specifications

The electrical characteristics of the capacitors in [Table 2-1](#) are described in this section. Characteristics of the PCB bulk and high-frequency capacitors are specified in [Table 2-2](#), followed by guidelines on acceptable substitutions. The equivalent series resistance (ESR) ranges specified for these capacitors can be over-ridden. However, this requires analysis of the resulting power distribution system impedance to ensure that no resonant impedance spikes result.

Table 2-2: PCB Capacitor Specifications

Ideal Value	Value Range ⁽¹⁾	Body Size ⁽²⁾	Type	ESL Maximum	ESR Range ⁽³⁾	Voltage Rating ⁽⁴⁾	Suggested Part Number
100 μ F	$C > 100 \mu\text{F}$	1210	2-Terminal Ceramic X7R or X5R	5 nH	$10 \text{ m}\Omega < \text{ESR} < 60 \text{ m}\Omega$	6.3V	GRM32ER60J107ME20L
4.7 μ F	$C > 4.7 \mu\text{F}$	0805	2-Terminal Ceramic X7R or X5R	2 nH	$10 \text{ m}\Omega < \text{ESR} < 60 \text{ m}\Omega$	6.3V	
0.47 μ F	$C > 0.47 \mu\text{F}$	0204 or 0402	2-Terminal Ceramic X7R or X5R	1.5 nH	$10 \text{ m}\Omega < \text{ESR} < 60 \text{ m}\Omega$	6.3V	

PCB Capacitor Substitution Rules:

1. Values can be larger than specified.
2. Body size can be smaller than specified.
3. ESR must be within the specified range.
4. Voltage rating can be higher than specified.

PCB Bulk Capacitors

The purpose of the bulk capacitors is to cover the low-frequency range between where the voltage regulator stops working ($\sim 100 \text{ KHz}$) and where the high-frequency capacitors start working ($\sim 2 \text{ MHz}$). As specified in [Table 2-1](#), all FPGA supplies require bulk capacitors.

The bulk capacitors in [Table 2-1](#) and [Table 2-2](#) are not necessarily in addition to the voltage regulator output capacitors required by the regulator manufacturer, provided there is no inductor, ferrite bead, choke, or other filter between the FPGA and the bulk capacitors. However, if the FPGA bulk and regulator output requirements are merged, the total capacitance of this network must not be less than the total bulk specified in [Table 2-1](#) and [Table 2-2](#), and must comply with the regulator manufacturer's output capacitor requirements.

The bulk PCB capacitors specified in [Table 2-1](#) are ceramic capacitors from Murata, a capacitor manufacturer. This capacitor was selected for its value, size, and low-cost. It is also RoHS compliant. If another manufacturer's capacitors or another type of capacitor (e.g., tantalum or high-performance electrolytic) meet the specifications listed in [Table 2-2](#), substitution is acceptable.

PCB High-Frequency Capacitors

There are two high-frequency capacitor values in [Table 2-2](#): the 4.7 μ F capacitor in an 0805 package and the 0.47 μ F capacitor in an 0402 or 0204 package. Substitutions can be made for some characteristics, but not others; see the notes attached to [Table 2-2](#) for details.

Capacitor Consolidation Rules

Sometimes a number of I/O banks are powered from the same voltage (e.g., 1.8V) and the recommended guidelines call for multiple bulk capacitors. This is also the case for V_{CCINT} and V_{CCAUX} in the larger devices. These many smaller capacitors can be consolidated into fewer (larger value) bulk capacitors provided the electrical characteristics of the consolidated capacitors (ESR and ESL) are equal to the electrical characteristics of the parallel combination of the recommended capacitors.

For most consolidations of V_{CCO} , V_{CCINT} , and V_{CCAUX} capacitors, large bulk capacitors (ceramic, tantalum, or high-performance electrolytic) with sufficiently low ESL and ESR are readily available. High-frequency capacitors cannot be consolidated as the usefulness of high-frequency capacitors depends on the number of PCB vias accessed.

Example

This example is of an FPGA with a single interface spanning three I/O banks, all powered from the same voltage. The required PCB capacitor table ([Table 2-1](#)) calls for one 100 μF capacitor per bank. These three capacitors can be consolidated into one capacitor since three 100 μF capacitors can be covered by one 330 μF capacitor. The following is then true:

- The ESL of the combination must be one-third of the specified capacitor. Three capacitors at 5 nH are equivalent to one capacitor at 1.7 nH. This implies that a 330 μF capacitor is acceptable provided its ESL is less than 1.7 nH.
- The ESR of the combination must be one-third of the specified capacitor. Three capacitors each in the range of 10 m Ω to 60 m Ω are equivalent to one capacitor in the range of 3.3 m Ω to 20 m Ω . A 330 μF capacitor is acceptable provided its ESR is in this range.
- Three 100 μF capacitors with 3 nH ESL and 20 m Ω ESR are replaced by one 330 μF capacitor with a 0.5 nH ESL and a 15 m Ω ESR.

PCB Capacitor Placement and Mounting Techniques

Placement and mounting restrictions presented in this section are unique to each capacitor type listed in the [Capacitor Specifications](#) section.

PCB Bulk Capacitors

Bulk capacitors can be large and difficult to place very close to the FPGA. Fortunately, this is not a problem because the low-frequency energy covered by bulk capacitors is not as sensitive to capacitor location. Bulk capacitors can be placed almost anywhere on the PCB, but the best placement is as close as possible to the FPGA. Capacitor mounting should follow normal PCB layout practices, tending toward short and wide shapes connecting to power planes with multiple vias.

0805 Ceramic Capacitor

The 4.7 μF 0805 capacitor covers the middle frequency range. Placement has some impact on its performance. The capacitor should be placed as close as possible to the FPGA. Any placement within two inches of the device's outer edge is acceptable.

The capacitor mounting (solder lands, traces, and vias) should be optimized for low inductance. Vias should be butted directly against the pads. Vias can be located at the ends of the pads (see Figure 2-1B), but are more optimally located at the sides of the pads (see Figure 2-1C). Via placement at the sides of the pads decreases the mounting's overall parasitic inductance by increasing the mutual inductive coupling of one via to the other. Dual vias can be placed on both sides of the pads (see Figure 2-1D) for even lower parasitic inductance, but with diminishing returns.




Figure 2-1: Example 0805 Capacitor Land and Mounting Geometries

0402 Ceramic Capacitor

The 0.47 μF 0402 capacitor covers the high-middle frequency range. Placement and mounting are *critical* for these capacitors.

The capacitor should be mounted as close to the FPGA as possible (achieves the least parasitic inductance possible).

For PCBs with a total thickness of < 1.575 mm (62 mils), the best placement location is on the PCB *backside*, within the device footprint (in the empty cross with an absence of vias). V_{CC} and GND vias corresponding to the supply of interest should be identified in the via array. Where space is available, 0402 mounting pads should be added and connected to these vias.

For PCBs with a total thickness > 1.575 mm (62 mils), the best placement location could be on the PCB top surface. The depth of the V_{CC} plane of interest in the PCB stackup is the key factor: if the V_{CC} plane is in the PCB stackup's top half, capacitor placement on the top PCB surface is optimal; if the V_{CC} plane is in the PCB stackup's bottom half, capacitor placement on the bottom PCB surface is optimal.

Any 0402 capacitors placed outside the device footprint (whether on the top or bottom surface) should be within 0.5 inch of the device's outer edge.

The capacitor mounting (solder lands, traces, and vias) must be optimized for low inductance. Vias should be butted against the pads with no trace length in-between. These vias should be at the sides of the pads if at all possible (see [Figure 2-2C](#)). Via placement at the sides of the pads decreases the mounting's parasitic inductance by increasing the mutual inductive coupling of one via to the other. Dual vias can be placed on both sides of the pads (see [Figure 2-2D](#)) for even lower parasitic inductance, but with diminishing returns.

Many manufacturing rules prevent mounting any device within 0.1 inch of the FPGA on the PCB top surface. Manufacturing rules can also prevent capacitor placement on the PCB backside within the device footprint, whether because backside mounting is prohibited or geometries necessary to fit mounting pads in the tight spaces between vias are too small for reliable soldering. These rules decrease the available options for capacitor placement but do not preclude meeting the Xilinx placement recommendations. Discuss any specific concerns with a PCB fabrication, assembly, and/or quality department.




Figure 2-2: Example 0402 Capacitor Land and Mounting Geometries

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 Spartan-6 FPGAs, require that V_{CC} supplies not fluctuate more than $\pm 5\%$ of the nominal V_{CC} value. This means that the peak-to-peak V_{RIPPLE} must be no more than 10% of the nominal V_{CC} . In this document the term V_{CC} is used generically for the following FPGA power supplies: V_{CCINT} , V_{CCO} , V_{CCAUX} , and V_{REF} . This assumes that nominal V_{CC} is exactly the nominal value provided in the data sheet. If not, then V_{RIPPLE} must be adjusted to a value correspondingly less than 10%.

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. 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 1 GHz.

Because the voltage level of V_{CC} 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). 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.

Figure 2-3 shows the major PDS components: the voltage regulator, the decoupling capacitors, and the active device being powered (FPGA).



UG393_c2_03_091809

Figure 2-3: Simplified PDS Circuit

Figure 2-4 shows a simplified PDS circuit with all reactive components represented by a frequency-dependent resistor.




Figure 2-4: Further Simplified PDS Circuit

Role of Inductance

Inductance is the property of the capacitors and the PCB current paths that slows down changes in current flow. 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 possible. 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 device and capacitors and between the capacitors and the voltage regulator (see Figure 2-3). 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 the bypass capacitors'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 , 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 100 μF to 1000 μF , tantalum capacitors are 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).

A *real* capacitor of any type then not only has capacitance characteristics but also inductance and resistance characteristics. Figure 2-5 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 2-5: Parasitics of a Real, Non-Ideal Capacitor

Figure 2-6 shows a *real* capacitor's impedance characteristic. Overlaid on this plot are dashed-line 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.

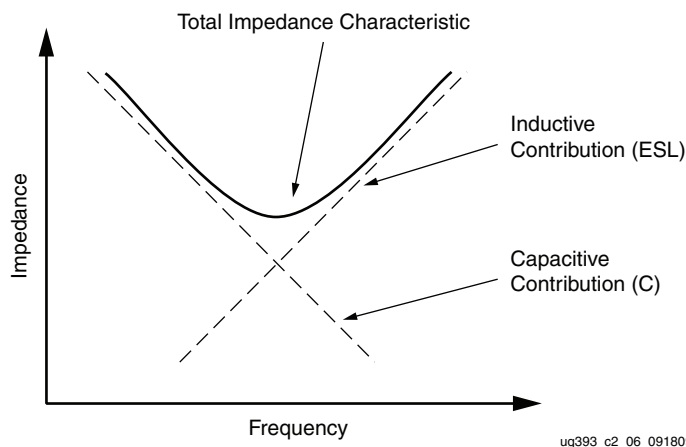


Figure 2-6: Contribution of Parasitics to Total Impedance Characteristics

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 Figure 2-8.

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 geometry.

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 Figure 2-7.




Figure 2-7: 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 (Figure 2-1A) and the via should butt up against the land (Figure 2-1B). Placing vias to the side of the capacitor lands (Figure 2-1C) or doubling the number of vias (Figure 2-1D), 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, reduce the total number of capacitors and maintain a one-to-one ratio of lands to vias.

The capacitor mounting (lands, traces, and vias) typically contributes about the same amount or more inductance than the capacitor's own parasitic self-inductance.

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 [Table 2-3](#).

Table 2-3: Capacitance and Spreading Inductance Values for Different Thicknesses of FR4 Power-Ground Plane Sandwiches

Dielectric Thickness		Inductance	Capacitance	
(micron)	(mil)	(pH/square)	(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 V_{CC} and GND planes. When possible, place the V_{CC} planes directly adjacent to the GND planes in the PCB stackup. Facing V_{CC} and GND planes are sometimes referred to as *sandwiches*. While the use of V_{CC} – 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.

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 [Table 2-3](#).

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:

- V_{CCINT} and GND are placed in a checkerboard arrangement in the center area of the BGA packages.
- V_{CCO} and GND pins are distributed among the I/O pins.

In BGA packages, 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 must 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_{SELF} . 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 2-5](#)). The RLC circuit's resonant frequency is the capacitor's self-resonant frequency.

To determine the RLC circuit's resonant frequency, use [Equation 2-1](#):

$$F = \frac{1}{2\pi\sqrt{LC}} \quad \text{Equation 2-1}$$

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 the [Simulation Methods](#) section 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_{SELF} , (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} :

$$L_{\text{IS}} = L_{\text{SELF}} + L_{\text{MOUNT}} \quad \text{Equation 2-2}$$

For example, using X7R Ceramic Chip capacitor in 0402 body size:

$C = 0.01 \mu\text{F}$ (selected by user)

$L_{\text{SELF}} = 0.9 \text{ nH}$ (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:

$$\begin{aligned} L_{IS} &= L_{SELF} + L_{MOUNT} = 0.9 \text{ nH} + 0.8 \text{ nH} \\ L_{IS} &= 1.7 \text{ nH} \end{aligned} \quad \text{Equation 2-3}$$

The values from the example are used to determine the mounted capacitor resonant frequency (F_{RIS}). Using Equation 2-1:

$$F_{RIS} = \frac{1}{2\pi\sqrt{L_{IS}C}} \quad \text{Equation 2-4}$$

$$F_{RIS} = \frac{1}{2\pi\sqrt{(1.7 \times 10^{-9} \text{ H}) \cdot (0.01 \times 10^{-6} \text{ F})}} = 38 \times 10^6 \text{ Hz} \quad \text{Equation 2-5}$$

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 2-8, for any given frequency along the inductive portion of the curve, the capacitors are equally effective.




Figure 2-8: Effective Frequency Example

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 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 V_{CC} 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 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 is 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). To counteract this, the decoupling capacitor must first sense a voltage difference.

A finite time delay (Equation 2-6) occurs between the start of the disturbance at the FPGA power pins and the point when the capacitor senses the disturbance.

$$\text{Time Delay} = \frac{\text{Distance from the FPGA power pins to the capacitor}}{\text{Signal propagation speed through FR4 dielectric}} \quad \text{Equation 2-6}$$

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. Since no low-frequency energy is needed, larger capacitors are not necessary.

Power Supply Consolidation

To design for reduced noise-levels on the V_{CCAUX} supply, separate supplies should be provided for V_{CCAUX} and any of the V_{CCO} voltage rails. However, a common supply to both V_{CCAUX} and V_{CCO} pins (at either 2.5V or 3.3V) is allowed as long as the recommended operating conditions for the V_{CC} inputs are met, as specified in the *Spartan-6 FPGA Data Sheet*. Similarly, this is also true if V_{CCO} and V_{CCINT} are being considered for combination, when using a 1.2V V_{CCO}. The supplies for the GTP transceiver V_{CC} pins should never be combined with other rails on the board.

Unconnected V_{CCO} Pins

In some cases, one or more I/O 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_{CCO} 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_{CCO} pins of unused I/O banks floating reduces the level of ESD protection on these pins and the I/O pins in the bank. ESD events at the unconnected solder balls in the inner rows of a BGA pinout array are unlikely and not considered a high risk. However, ESD events at exposed pins on the perimeter of a QFP-type package are likely. In these packages, the V_{CCO} pins of unused I/O banks should be connected to the V_{CCO} of a neighboring I/O bank.

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 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 impedance calculation (among other useful tools for PDS simulation) on his website:

<http://www.electrical-integrity.com>

Table 2-4 also lists a few EDA tool vendors for PDS design and simulation. These tools span a wide range of sophistication levels.

Table 2-4: EDA Tools for PDS Design and Simulation

Tool	Vendor	Website URL
ADS	Agilent	http://www.agilent.com
SIwave, HFSS	Ansoft	http://www.ansoft.com
Specctraquest Power Integrity	Cadence	http://www.cadence.com
Speed 2000, PowerSI, PowerDC, OptimizePI	Sigrity	http://www.sigrity.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 or at an unused I/O driven High or Low (referred to as a *spyhole measurement*).

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 generally in opposition to the noise, and therefore 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 "spyhole" 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.

Figure 2-9 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 2-9: Averaged Measurement of V_{CCO} Supply with Multiple I/O Sending Patterns at 250 Mb/s

Figure 2-10 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 2-10: Infinite Persistence Measurement of Same Supply

The measurement shown in Figure 2-9 and Figure 2-10 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 the user is 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, the user sees 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.

Figure 2-11 is an example of a noise spectrum measurement of the V_{CCO} power-supply noise, with multiple I/O sending patterns at 100 MHz.



UG393_c2_11_091809

Figure 2-11: Screenshot of Spectrum Analyzer Measurement of V_{CCO}

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 (Equation 2-7):

$$I(f) = \frac{V(f) \text{ From Spectrum Analyzer}}{Z(f) \text{ From Network Analyzer}} \quad \text{Equation 2-7}$$

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 (see [Figure 2-1A](#)).

To reduce the current loop area, vias should be placed directly against capacitor solder lands (see [Figure 2-1B](#)). *Never* connect vias to the lands with a section of trace (see [Figure 2-1A](#)).

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), shown in [Figure 2-1C](#). Double vias also improve connecting trace geometry and capacitor land geometry (see [Figure 2-1D](#)).

Exceptionally thick boards (> 2.3 mm or 90 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 highest frequency 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 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 (decoupling capacitors at PCB locations where layer transitions occur).

SelectIO Signaling

The Spartan-6 FPGA 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.

This chapter contains the following sections:

- [Interface Types](#)
- [Single-Ended Signaling](#)

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 Spartan-6 FPGA IOB configuration options.

Modes and Attributes

Some of these 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), and stronger input-termination resistors. Drive strength, slew rate, and in some cases specifying untuned output driver impedance 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. See the *Spartan-6 FPGA SelectIO Resources User Guide* for more information.

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 LVCMOS18 and SSTL18 Class 1. When a Spartan-6 FPGA is receiving, the input thresholds, V_{IL} and V_{IH} , are much closer together for the SSTL18 standard.

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. The one drawback to the use of V_{REF} is that the semi-dedicated V_{REF} pins of the bank cannot be used as I/Os – they must all be connected to an external reference voltage with a decoupling capacitor for each V_{REF} pin. For more information on V_{REF} decoupling and decoupling of all other supplies, see [Chapter 2, “Power Distribution System.”](#)

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 [Equation 4-1](#):

$$f \approx 0.35 / T \quad \text{Equation 4-1}$$

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 [Equation 4-1](#). 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 materials 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 $0.06\Omega/\text{inch}$, while the same trace at 10 Gb/s has a resistance of just over $1\Omega/\text{inch}$. 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 $Z_{DIFF} = 2 \times Z_{0O}$.

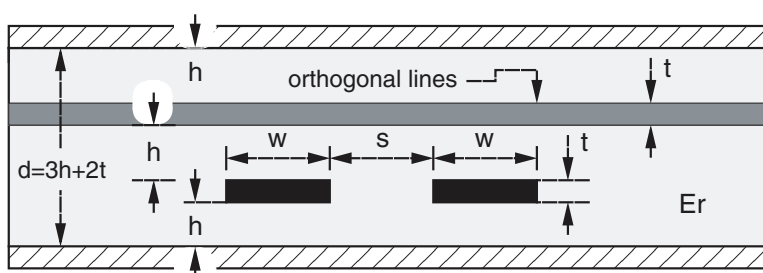
The same W/S ratio also must be less than 0.8, otherwise strong coupling between the traces requires narrower, lossier traces for a Z_{0O} of 50Ω . To clarify, with Z_{0O} at 50Ω , an even mode impedance (Z_{0E}) of 60Ω or below is desired.

Figure 4-1 through Figure 4-4 show example cross sections of differential structures.




UG393_c4_01_091809

Figure 4-1: Differential Edge-Coupled Centered Stripline




UG393_c4_02_091809

Figure 4-2: Differential Edge-Coupled Offset Stripline



UG393_c4_03_091809

Figure 4-3: Centered Broadside-Coupled Stripline



UG393_c4_04_091809

Figure 4-4: Differential Microstrip

A good PCB manufacturer understands controlled impedance and allows fine adjustments for line widths to produce a Z_{0O} 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_{0O} 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, page 51](#).

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, page 39](#). 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.

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 [Figure 5-1](#)) causes a momentary dip in the impedance, while a series inductance (see [Figure 5-2](#)) 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.



UG393_c5_01_091809

Figure 5-1: TDR Signature of Shunt Capacitance



UG393_c5_02_091809


Figure 5-2: 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:

$$C = -\frac{2}{Z_0} \int_{t_1}^{t_2} \frac{V_{\text{tdr}}(t) - V_{\text{step}}}{V_{\text{step}}} dt \quad \text{Equation 5-1}$$

$$L = 2Z_0 \int_{t_1}^{t_2} \frac{V_{\text{tdr}}(t) - V_{\text{step}}}{V_{\text{step}}} dt \quad \text{Equation 5-2}$$

Figure 5-3 shows the integration of the normalized TDR area.



UG393_c5_03_091809

Figure 5-3: 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

The transceiver signal paths within the BGA package are optimized using a 3D full-wave solver. Package traces are designed to be 50Ω high-speed transmission lines, while solder ball and bump regions are tuned to 50Ω .

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 [Figure 5-4](#) 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 5-4: 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 [Figure 5-5](#), 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 5-5: 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.

Figure 5-6 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 5-6: Ansoft HFSS Model of Pad Clear-Out

Figure 5-7 shows the return loss comparison between 0402 pad structures with linear scale.




Figure 5-7: Return Loss Comparison Between 0402 Pad Structures

The approximately -40 dB/decade slope in Figure 5-8 shows good fit to the frequency response of a lumped capacitor.




Figure 5-8: 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 Figure 5-9 and Figure 5-10, 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 5-1 and Equation 5-2.




Figure 5-9: TDR Results Comparing 0402 Pad Structures




Figure 5-10: TDR Results Comparing 0402 Pad Structures

As shown from Figure 5-11 and Figure 5-12, 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 5-11: 840 fF Excess Capacitance with Ground Plane Intact




Figure 5-12: 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.

Figure 5-13 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 5-13: Differential Via Design Example

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 [Figure 5-13](#) 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.

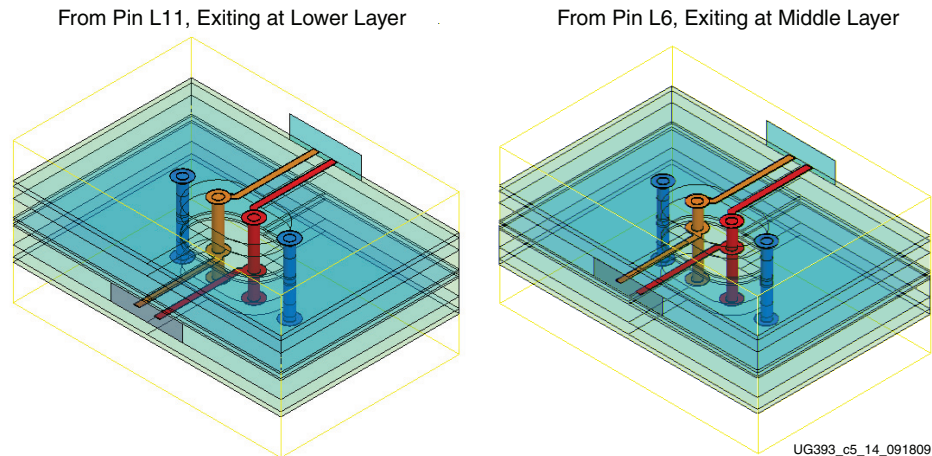


Figure 5-14: Differential GSSG Via in 16-Layer PCB from Pins L11 and L6

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 Figure 5-15 compares the S-parameter return loss for common-mode (SCC11) and differential (SDD11) responses.




Figure 5-15: Simulated Return Loss Comparing Differential and Common-Mode Losses for L11 and L6 GSSG Vias

From the graph in Figure 5-15, 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.

Xilinx 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. [Figure 5-16](#) shows example bends in traces.




Figure 5-16: Example Design for 90 Degree Bends in Traces

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.

Figure 5-17 through Figure 5-19 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 5-17: Simulated TDR of 45 Degree Bends with Jog-Outs




Figure 5-18: Simulated TDR of 45 Degree Bends with Jog-Outs




Figure 5-19: Simulated Phase Response of 45 Degree Bends with Jog-Outs

For wide traces, curved routing can also be helpful as shown in Figure 5-20.




Figure 5-20: Measured TDR of 45 Degree Bends with and without Jog-Outs

