Virtex-5 FPGA
PCB Designer's Guide

XILINX®

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
12/15/2006	1.0	Initial Xilinx release.
08/20/2007	1.1	Updated V_{CCO} capacitors in Table 2-1 to Table 2-6, Table 2-13 to Table 2-18, and Table 2-20. Added SXT V_{CCO} capacitors in Table 2-21 through Table 2-24. Updated ESR instructions in Table 2-29, page 25. Clarified ESR ranges in "Capacitor Specifications."
05/19/2008	1.2	Complete revision of the Chapter 2, "Power Distribution System" section on Decoupling Capacitors: Inside the Package and on the PCB.
12/17/2008	1.3	• Updated Total value in FF323/XC5VLX30T row of Table 2-2, page 14. • Revised Table 2-11 and Table 2-12, page 18. • Added new section "Unconnected V_{CCO} Pins," page 33.
04/20/2009	1.4	• Added new tables, Table 2-5, page 16, and Table 2-10, page 18, with required PCB capacitor quantities and substrate decoupling capacity quantities, respectively, for TXT devices. • Updated Table 2-12, page 18 to include capacitor specifications for FF1156 and FF1759 packages.
02/11/2014	1.5	Updated the disclaimer and copyright. Revised the 4.7µF capacitor value for the XC5VLX220T-FF1738 in Table 2-7, page 17. Updated the link in the third paragraph in "Simulation Methods," page 33.

Table of Contents

Preface

About This Guide

This guide provides information on PCB design for Virtex®-5 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 Virtex-5 FPGAs, including all details of decoupling capacitor selection, use of voltage regulators and PCB geometries, simulation and measurement.

Additional Documentation

The following documents are also available for download at http://www.xilinx.com/virtex5.

- Virtex-5 Family Overview

  The features and product selection of the Virtex-5 family are outlined in this overview.

- Virtex-5 FPGA Data Sheet: DC and Switching Characteristics

  This data sheet contains the DC and Switching Characteristic specifications for the Virtex-5 family.

- Virtex-5 FPGA User Guide

  This user guide includes chapters on:

  - Clocking Resources
  - Clock Management Technology (CMT)
  - Phase-Locked Loops (PLLs)
  - Block RAM and FIFO memory
  - Configurable Logic Blocks (CLBs)
  - SelectIO™ Resources
  - I/O Logic Resources
  - Advanced I/O Logic Resources

- Virtex-5 FPGA RocketIO GTP Transceiver User Guide

  This guide describes the RocketIO™ GTP transceivers available in the Virtex-5 LXT and SXT platforms.

- Virtex-5 FPGA RocketIO GTX Transceiver User Guide

  This guide describes the RocketIO GTX transceivers available in the Virtex-5 TXT and FXT platforms.

- Virtex-5 FPGA Embedded Tri-Mode Ethernet MAC User Guide

  This user guide describes the dedicated Tri-Mode Ethernet Media Access Controller available in the Virtex-5 LXT, SXT, TXT, and FXT platform devices.

- Virtex-5 FPGA Integrated Endpoint Block User Guide for PCI Express Designs

  This user guide describes the integrated Endpoint blocks in the Virtex-5 LXT, SXT, TXT, and FXT platform devices for PCI Express® designs.

- Embedded Processor Block in Virtex-5 FPGAs Reference Guide

  This reference guide describes the embedded processor block available in the Virtex-5 FXT platform.

- Virtex-5 FPGA XtremeDSP Design Considerations User Guide

  This guide describes the XtremeDSP™ slice and includes reference designs for using the DSP48E.

- Virtex-5 FPGA Configuration Guide

  This all-encompassing configuration guide includes chapters on configuration interfaces (serial and SelectMAP), bitstream encryption, Boundary-Scan and JTAG configuration, reconfiguration techniques, and readback through the SelectMAP and JTAG interfaces.

- Virtex-5 FPGA System Monitor User Guide

  The System Monitor functionality available in all the Virtex-5 devices is outlined in this guide.

- Virtex-5 FPGA Packaging and Pinout Specification

  This specification includes the tables for device/package combinations and maximum I/Os, pin definitions, pinout tables, pinout diagrams, mechanical drawings, and thermal specifications.

Additional Support Resources

To find additional documentation, see the Xilinx website at:

http://www.xilinx.com/support/documentation/index.htm.

To search the Answer Database of silicon, software, and IP questions and answers, see the Xilinx website at:

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

Typographical Conventions

This document uses the following typographical conventions. An example illustrates each convention.

Convention	Meaning or Use	Example
Italic font	References to other documents	See the *Virtex-5 Configuration Guide* for more information.
	Emphasis in text	The address (F) is asserted *after* clock event 2.
Underlined Text	Indicates a link to a web page.	http://www.xilinx.com/virtex5

Online Document

The following conventions are used in this document:

Convention	Meaning or Use	Example
Blue text	Cross-reference link to a location in the current document	See the section "Additional Documentation" for details. Refer to "Capacitor Specifications" in Chapter 2. for details.
Blue, underlined text	Hyperlink to a website (URL)	Go to http://www.xilinx.com for the latest documentation.

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 ball pitch (1.0 mm for FF packages) defines the land pad layout. The minimum surface feature sizes of current PCB technology define the via arrangement in the area under the device. Minimum via diameters and *keep-out areas* around those vias are defined by the PCB manufacturer. These diameters limit the amount of space available in-between vias for routing of signals in and out of the via array underneath the device. These diameters define the maximum trace width in these *breakout* traces. PCB manufacturing limits constrain the minimum trace width and minimum spacing.

The total number of PCB layers necessary to accommodate an FPGA is defined by the number of signal layers and the number of plane layers.

- The number of signal layers is defined by the number of I/O signal traces routed in and out of an FPGA package (usually following the total User I/O count of the package).
- The number of plane layers is defined by the number of power and ground plane layers necessary to bring power to the FPGA and to provide references and isolation for signal layers.

Most PCBs for large FPGAs range from 12 to 22 layers.

System compliance often defines the total thickness of the board. Along with the number of board layers, this defines the maximum layer thickness, and therefore, the spacing in the Z direction of signal and plane layers to other signal and plane layers. Z-direction spacing of signal trace layers to other signal trace layers affects crosstalk. Z-direction spacing of signal trace layers to reference plane layers affects signal trace impedance. Z-direction spacing of plane layers to other plane layers affects power system parasitic inductance.

Z-direction spacing of signal trace layers to reference plane layers (defined by total board thickness and number of board layers) is a defining factor in trace impedance.Trace width (defined by FPGA package ball pitch and PCB via manufacturing constraints) is another factor in trace impedance. A designer 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 or 100Ω differential).

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 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.

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 contributor to ground bounce and Power Distribution System (PDS) noise from simultaneously switching outputs. The importance of return current paths cannot be underestimated.

EXILINX

<div align="right">Chapter 2</div>

Power Distribution System

This chapter documents the power distribution system (PDS) for Virtex®-5 FPGAs, including decoupling capacitor selection, placement, and PCB geometries. A simple decoupling method is provided for each device in the Virtex-5 FPGA family. Basic PDS design principles are covered, as well as simulation and analysis methods. This chapter contains the following sections:

- "Decoupling Capacitors: Inside the Package and on the PCB"
- "Basic PDS Principles"
- "Simulation Methods"
- "PDS Measurements"
- "Troubleshooting"

Decoupling Capacitors: Inside the Package and on the PCB

Recommended PCB Capacitors per Device

A simple PCB-decoupling network for each Virtex-5 device is listed in Table 2-1 through Table 2-5. Many of the devices require very few PCB ceramic capacitors because high-frequency ceramic capacitors are already present inside the device package (mounted on the substrate). Thus PCB decoupling capacitor quantities can be significantly lower in Virtex-5 devices than in previous device families.

Decoupling methods other than those presented in Table 2-1 through Table 2-5 can be used, but the decoupling network should be designed to meet or exceed the performance of the simple decoupling networks presented here.

The amount of in-package decoupling capacitance varies from package to package. There are three package types:

- Packages with high-performance multi-terminal capacitors
- Packages with medium-performance capacitors
- Packages with no capacitors

Table 2-6 through Table 2-10 list the discrete package capacitors present inside the device. This information is intended for use in power system simulation.

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 40-pin I/O bank (two 20-pin I/O banks can share the capacitors specified for one 40-pin 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 through Table 2-5 lists the PCB decoupling capacitor guidelines per V_{CC} supply rail for LX, LXT, SXT, FXT, and TXT devices, respectively.

Table 2-1: **Required PCB Capacitor Quantities per Device: LX Devices[1]**

Package	Device	V_{CCINT}			V_{CCAUX}			V_{CCO} (per I/O Bank)			Total[2]
		330 µF	2.2 µF	0.22 µF	33 µF	2.2 µF	0.22 µF	47 µF	2.2 µF	0.22 µF	
FF324	XC5VLX30	1	4	8	1	2	4	1	2	4	62
FF324	XC5VLX50	1	6	12	1	2	4	1	2	4	68
FF676	XC5VLX30	1			1			1			12
FF676	XC5VLX50	1			1			1			13
FF676	XC5VLX85	2			1			1			14
FF676	XC5VLX110	2			1			1			14
FF1153	XC5VLX50	1			2						3
FF1153	XC5VLX85	2			2						4
FF1153	XC5VLX110	2			2						4
FF1153	XC5VLX155	3			2						5
FF1760	XC5VLX110	2			2						4
FF1760	XC5VLX155	3			2						5
FF1760	XC5VLX220	5			2						7
FF1760	XC5VLX330	7			3						10

Notes:
1. PCB Capacitor specifications are listed in Table 2-11.
2. Total includes all capacitors for all supplies, accounting for the number of I/O banks in the device.

Table 2-2: **Required PCB Capacitor Quantities per Device: LXT Devices[1]**

Package	Device	V_{CCINT}			V_{CCAUX}			V_{CCO} (per I/O Bank)			Total[2]
		330 µF	2.2 µF	0.22 µF	33 µF	2.2 µF	0.22 µF	47 µF	2.2 µF	0.22 µF	
FF323	XC5VLX20T	1	3	6	1	1	3	1	2	4	50
FF323	XC5VLX30T	1	4	8	1	2	4	1	2	4	55
FF665	XC5VLX30T	1			1			1			11
FF665	XC5VLX50T	1			1			1			11
FF1136	XC5VLX50T	1			1						2
FF1136	XC5VLX85T	2			1						3
FF1136	XC5VLX110T	2			2						4
FF1136	XC5VLX155T	3			2						5
FF1738	XC5VLX110T	2			2						4

Table 2-2: **Required PCB Capacitor Quantities per Device: LXT Devices**[1] *(Cont'd)*

Package	Device	V_CCINT			V_CCAUX			V_CCO (per I/O Bank)			Total[2]
		330 µF	2.2 µF	0.22 µF	33 µF	2.2 µF	0.22 µF	47 µF	2.2 µF	0.22 µF	
FF1738	XC5VLX155T	3			2						5
FF1738	XC5VLX220T	5			2						7
FF1738	XC5VLX330T	7			3						10

Notes:
1. PCB Capacitor specifications are listed in Table 2-11.
2. Total includes all capacitors for all supplies, accounting for the number of I/O banks in the device.

Table 2-3: **Required PCB Capacitor Quantities per Device: SXT Devices**[1]

Package	Device	V_CCINT			V_CCAUX			V_CCO (per I/O Bank)			Total[2]
		330 µF	2.2 µF	0.22 µF	33 µF	2.2 µF	0.22 µF	47 µF	2.2 µF	0.22 µF	
FF665	XC5VSX35T	1			1			1			11
FF665	XC5VSX50T	2			1			1			12
FF1136	XC5VSX50T	2			1						3
FF1136	XC5VSX95T	2			2						4
FF1738	XC5VSX240T	7			3						10

Notes:
1. PCB Capacitor specifications are listed in Table 2-11.
2. Total includes all capacitors for all supplies, accounting for the number of I/O banks in the device.

Table 2-4: **Required PCB Capacitor Quantities per Device: FXT Devices**[1]

Package	Device	V_CCINT			V_CCAUX			V_CCO (per I/O Bank)			Total[2]
		330 µF	2.2 µF	0.22 µF	33 µF	2.2 µF	0.22 µF	47 µF	2.2 µF	0.22 µF	
FF665	XC5VFX30T	1			1			1			11
FF665	XC5VFX70T	2			1			1			12
FF1136	XC5VFX70T	2			2						4
FF1136	XC5VFX100T	2			2						4
FF1738	XC5VFX100T	2			2						4
FF1738	XC5VFX130T	3			2						5
FF1738	XC5VFX200T	5			3						8

Notes:
1. PCB Capacitor specifications are listed in Table 2-11.
2. Total includes all capacitors for all supplies, accounting for the number of I/O banks in the device.

Table 2-5: **Required PCB Capacitor Quantities per Device: TXT Devices[1]**

Package	Device	V_CCINT			V_CCAUX			V_CCO (per I/O Bank)			Total[2]
		330 μF	2.2 μF	0.22 μF	33 μF	2.2 μF	0.22 μF	47 μF	2.2 μF	0.22 μF	
FF1156	XC5VTX150T	4			1				1		14
FF1759	XC5VTX150T	4			1				1		22
FF1759	XC5VTX240T	6			2				1		25

Notes:

1. PCB Capacitor specifications are listed in Table 2-11.
2. Total includes all capacitors for all supplies, accounting for the number of I/O banks in the device.

Discrete Package Capacitors

Table 2-6 through Table 2-10 lists the capacitors present inside each device for LX, LXT, SXT, FXT, and TXT devices, respectively.

Table 2-6: **Substrate Decoupling Capacitor Quantities Inside Device: LX Devices**

Package	Device	V_CCINT		V_CCAUX		V_CCO (per I/O Bank)	
		2.2 μF	4.7 μF	0.1 μF	1.0 μF	0.1 μF	1.0 μF
FF324	XC5VLX30						
FF324	XC5VLX50						
FF676	XC5VLX30	2		2		1	
FF676	XC5VLX50	2		2		1	
FF676	XC5VLX85	4		2		1	
FF676	XC5VLX110	4		2		1	
FF1153	XC5VLX50		2		2		1
FF1153	XC5VLX85		2		2		1
FF1153	XC5VLX110		4		2		1
FF1153	XC5VLX155		4		2		1
FF1760	XC5VLX110		4		2		1
FF1760	XC5VLX155		4		2		1
FF1760	XC5VLX220		8		2		1
FF1760	XC5VLX330		8		2		1

Notes:

1. Package capacitor specifications are listed in Table 2-12.

Table 2-7: **Substrate Decoupling Capacitor Quantities Inside Device: LXT Devices**

Package	Device	V_{CCINT}		V_{CCAUX}		V_{CCO} (per I/O Bank)	
		2.2 μF	4.7 μF	0.1 μF	1.0 μF	0.1 μF	1.0 μF
FF323	XC5VLX20T						
FF323	XC5VLX30T						
FF665	XC5VLX30T	2		2		1	
FF665	XC5VLX50T	2		2		1	
FF1136	XC5VLX50T		2		2		1
FF1136	XC5VLX85T		2		2		1
FF1136	XC5VLX110T		4		2		1
FF1136	XC5VLX155T		4		2		1
FF1738	XC5VLX110T		4		2		1
FF1738	XC5VLX155T		4		2		1
FF1738	XC5VLX220T		4		2		1
FF1738	XC5VLX330T		8		2		1

Notes:

1. Package capacitor specifications are listed in Table 2-12.

Table 2-8: **Substrate Decoupling Capacitor Quantities Inside Device: SXT Devices**

Package	Device	V_{CCINT}		V_{CCAUX}		V_{CCO} (per I/O Bank)	
		2.2 μF	4.7 μF	0.1 μF	1.0 μF	0.1 μF	1.0 μF
FF665	XC5VSX35T	2		2		1	
FF665	XC5VSX50T	2		2		1	
FF1136	XC5VSX50T		2		2		1
FF1136	XC5VSX95T		4		2		1
FF1738	XC5VSX240T		8		2		1

Notes:

1. Package capacitor specifications are listed in Table 2-12.

Table 2-9: **Substrate Decoupling Capacitor Quantities Inside Device: FXT Devices**

Package	Device	V_{CCINT}		V_{CCAUX}		V_{CCO} (per I/O Bank)	
		2.2 μF	4.7 μF	0.1 μF	1.0 μF	0.1 μF	1.0 μF
FF665	XC5VFX30T	4		2		1	
FF665	XC5VFX70T	4		2		1	
FF1136	XC5VFX70T		4		2		1
FF1136	XC5VFX100T		4		2		1
FF1738	XC5VFX100T		6		2		1
FF1738	XC5VFX130T		6		2		1
FF1738	XC5VFX200T		6		2		1

Notes:

1. Package capacitor specifications are listed in Table 2-12.

Table 2-10: **Substrate Decoupling Capacitor Quantities Inside Device: TXT Devices**

Package	Device	V$_{CCINT}$		V$_{CCAUX}$			V$_{CCO}$ (per I/O Bank)		
		2.2 µF	4.7 µF	0.1 µF	1.0 µF	4.7 µF	0.1 µF	1.0 µF	0.047 µF
FF1156	XC5VTX150T		4			2			1
FF1759	XC5VTX150T		6			2			1
FF1759	XC5VTX240T		6			2			1

Notes:

1. Package capacitor specifications are listed in Table 2-12.

Capacitor Specifications

The electrical characteristics of the capacitors in Table 2-1 through Table 2-5 are described in this section. Characteristics of the PCB tantalum and ceramic capacitors are specified in Table 2-11, followed by guidelines on acceptable substitutions. Characteristics of the discrete package capacitors inside the device are described in Table 2-12. The 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 2-11: **PCB Capacitor Specifications**

Ideal Value	Value Range	Body Size[1]	Type	ESR Range[2]	Voltage Rating[3]	Suggested Part Number
330 µF	C > 330 µF	V-Case	2-Terminal Tantalum	15 mΩ < ESR < 40 mΩ	2.5V	T520V337M2R5ATE025
47 µF	C > 47 µF	B-Case	2-Terminal Tantalum	50 mΩ < ESR < 90 mΩ	6.3V	T520B476M006ATE070
33 µF	C > 33 µF	B-Case	2-Terminal Tantalum	30 mΩ < ESR < 60 mΩ	6.3V	T520B336M006ATE040
2.2 µF	C > 2.2 µF	0805	2-Terminal Ceramic X7R or X5R	5 mΩ < ESR < 60 mΩ	6.3V	
0.22 µF	C > 0.22 µF	0204 or 0402	2-Terminal Ceramic X7R or X5R	15 mΩ < ESR < 60 mΩ	6.3V	

Notes:

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

Table 2-12: **Discrete Substrate Capacitor Specifications**

Value	Body Size	Type	ESR	Voltage Rating	Used On Supplies	Used in Packages
4.7 µF	0805	8-Terminal Ceramic X7S	4 mΩ	4.0V	V$_{CCINT}$	FF1760, FF1738, FF1153, FF1136
2.2 µF	0603	8-Terminal Ceramic X7S	5 mΩ	4.0V	V$_{CCINT}$	FF676, FF665
1.0 µF	0603	8-Terminal Ceramic X7S	6 mΩ	4.0V	V$_{CCO}$, V$_{CCAUX}$	FF1760, FF1738, FF1153, FF1136

Table 2-12: **Discrete Substrate Capacitor Specifications** *(Cont'd)*

Value	Body Size	Type	ESR	Voltage Rating	Used On Supplies	Used in Packages
0.1 µF	0204	2-Terminal Ceramic X6S	15 mΩ	6.3V	V_{CCO}, V_{CCAUX}	FF676, FF665
0.047 µF	0603	8-Terminal Ceramic X7S	135 mΩ	4.0V	V_{CCO}	FF1156, FF1759

PCB Tantalum Capacitors

The purpose of the tantalum capacitors is to cover the low-frequency range between where the voltage regulator stops working and where the ceramic capacitors start working. As specified in Table 2-1 through Table 2-5, some FPGA supplies do not require any explicit tantalum capacitors. On these supplies, adequate low frequency bulk capacitance at the output of the voltage regulator is all that is necessary. See the regulator manufacturer's guidelines for output capacitor requirements.

The tantalum PCB capacitors specified in Table 2-1 through Table 2-5 are from Kemet, a capacitor manufacturer. These parts were selected for their value and controlled equivalent series resistance (ESR). They are also RoHS compliant. If another manufacturer's tantalum capacitors or high-performance electrolytic capacitors meet the specifications listed in Table 2-11, substitution is acceptable.

PCB Ceramic Capacitors

There are two ceramic capacitor values in Table 2-11: the 2.2 µF capacitor in an 0805 package and the 0.22 µF capacitor in an 0402 or 0204 package. Substitutions can be made for some characteristics, but not others; see the notes attached to Table 2-11 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 tantalum capacitors. This is also the case for V_{CCINT} and V_{CCAUX} in the larger Virtex-5 devices. These many smaller capacitors can be consolidated into fewer (larger value) tantalum 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 tantalum capacitors with sufficiently low ESL and ESR are readily available. Ceramic capacitors cannot be consolidated as the usefulness of ceramic capacitors depends on the number of PCB vias accessed.

Example

This example is of an FPGA with a single memory interface spanning three I/O banks, all powered from the same voltage. The Required PCB Capacitor table (Table 2-1 through Table 2-5) calls for one 47 µF capacitor per bank. These three capacitors can be consolidated into one capacitor since three 47 µF capacitors can be covered by one 150 µF capacitor. The following is then true:

- The ESL of the combination must be one-third of the specified capacitor. Three capacitors, each at 3 nH equals one capacitor at 1 nH.
- The ESR of the combination must be one-third of the specified capacitor. Three capacitors, each at 30-60 mΩ equals one capacitor at 10-20 mΩ.

Three 47 µF capacitors with 3 nH ESL and 50 mΩ ESR are replaced by one 150 µF capacitor with a 1 nH ESL and a 15 mΩ ESR.

PCB Capacitor Placement and Mounting Techniques

Placement and mounting restrictions presented in this section are different for each capacitor type listed in the "Capacitor Specifications" section.

PCB Tantalum Capacitors

Tantalum 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 tantalum capacitors is not as sensitive to capacitor location. Tantalum capacitors can be placed almost anywhere on the PCB, but the best placement is close as possible to the FPGA. Capacitor mounting should follow normal PCB layout practices, tending toward short and wide shapes connecting to power planes.

0805 Ceramic Capacitor

The 2.2 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) or 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.
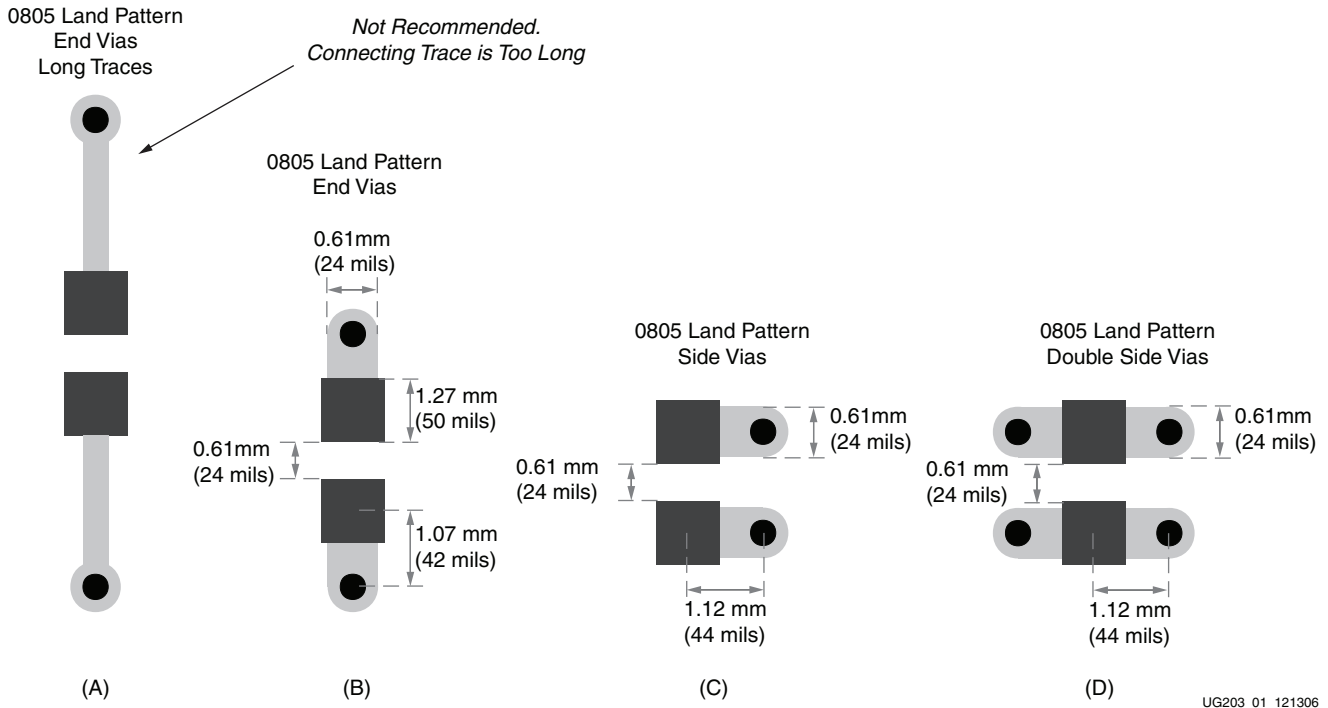
Figure 2-1: **Example 0805 Capacitor Land and Mounting Geometries**

0402 Ceramic Capacitor

The 0.22 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.
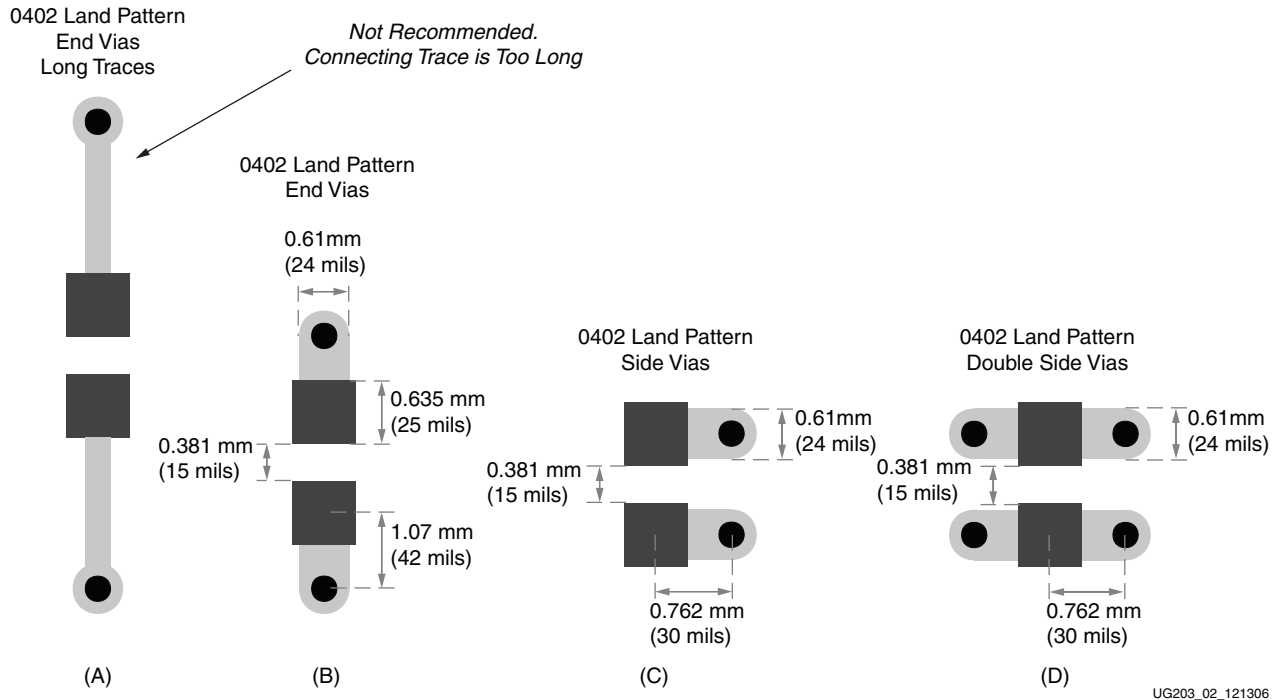

Figure 2-2: **Example 0402 Capacitor Land and Mounting Geometries**

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 Virtex-5 FPGAs, require that V_{CC} supplies not fluctuate more than ±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 all 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.

- 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.

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.

The 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.
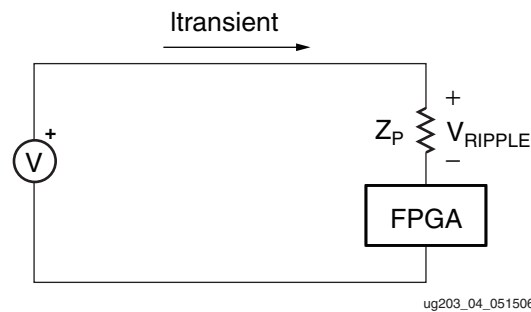
Figure 2-3 shows the major PDS components: the power supply, the decoupling capacitors, and the active device being powered (FPGA).
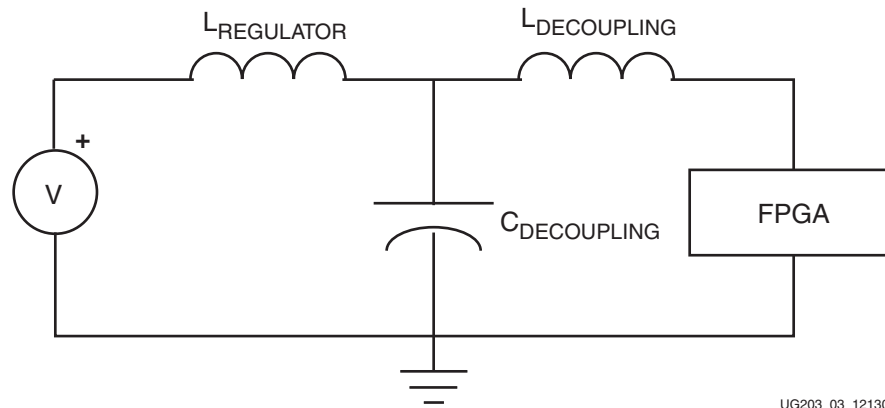
Figure 2-3: **Simplified PDS Circuit**

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 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 the change occurs.

The PDS 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 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 10 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 47 µF to 680 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 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).
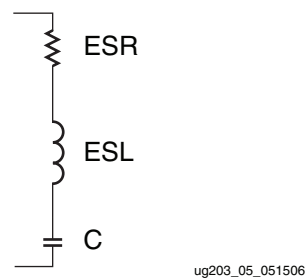

ug203_05_051506

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 curves corresponding to the capacitor's capacitance and parasitic inductance (ESL). These two curves combine to form the RLC circuit's total impedance characteristic.
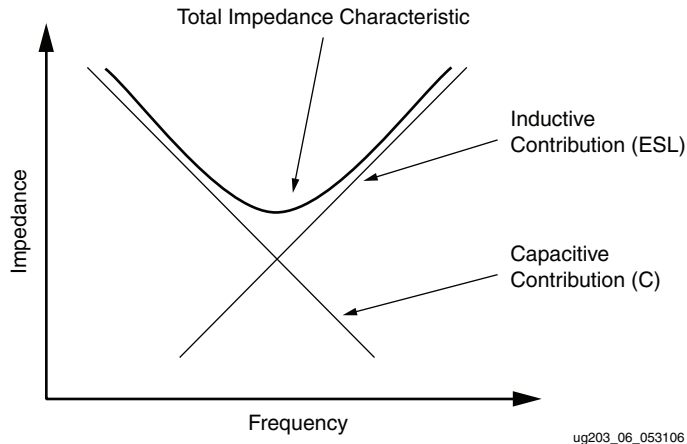
Figure 2-6: **Contribution of Parasitics to Total Impedance Characteristics**

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.

For a fixed capacitor package, the low-frequency 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 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 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.
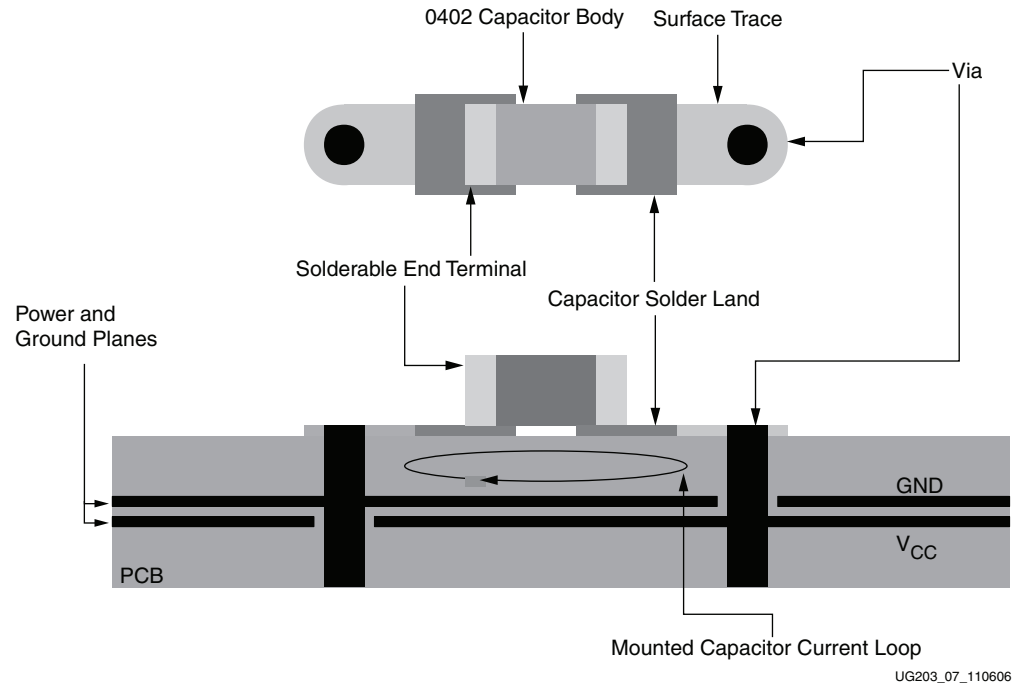
Figure 2-7: **Example Cutaway View of PCB with Capacitor Mounting**

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 (AVX's LICC).

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 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 2-13.

Table 2-13: **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, the speeds involved and the sheer amount of power required for fast, dense devices 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 2-13.

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 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, carried by a V_{CCO} or GND pin.

To reduce parasitic inductance:

- Core V_{CC} pins such as V_{CCINT} and V_{CCAUX} are placed in a checkerboard arrangement in the pinout.

- V_{CCO} and GND pins are distributed among the I/O pins.

Every I/O pin in the Virtex-5 FPGA pinout is adjacent to a return-current pin.

FPGA pinout arrangement determines the PCB via arrangement. The PCB designer cannot control the proximity of opposing current paths but has control over the trade-offs between the capacitor's mounting inductance and FPGA's mounting inductance:

- Both mounting inductances are reduced by placing power planes close to the PCB stackup's top half and placing the capacitors on the top surface (reducing the capacitor's via length).

- If power planes are placed in the PCB stackup's bottom half, the capacitors must be mounted on the PCB backside. In this case, FPGA mounting vias are already long, and making the capacitor vias long as well is a bad practice. A better practice is to take advantage of the short distance between the underside of the PCB and the power plane of interest, mounting capacitors on the underside.

PCB Stackup and Layer Order

V_{CC} and ground plane placement in the PCB stackup (the layer order) has a significant impact on the parasitic inductances of power current paths. Layer order must be considered early in the design process:

- High-priority supplies should be placed closer to the FPGA (in the PCB stackup's top half)

- Low-priority supplies should be placed farther from the FPGA (in the PCB stackup's bottom half)

Power supplies with high transient current should have the associated V_{CC} planes close to the top surface (FPGA side) of the PCB stackup. This decreases the vertical distance (V_{CC} and GND via length) that currents travel before reaching the associated V_{CC} and GND planes. To reduce spreading inductance, every V_{CC} plane should have an adjacent GND plane in the PCB stackup. The skin effect causes high-frequency currents to couple tightly and the GND plane adjacent to a specific V_{CC} plane tends to carry the majority of the current complementary to that in the V_{CC} plane. Thus, adjacent V_{CC} and GND planes are treated as a pair.

Not all V_{CC} and GND plane pairs reside in the PCB stackup's top half because manufacturing constraints typically require a symmetrical PCB stackup around the center (with respect to dielectric thicknesses and etched copper areas). The PCB designer chooses the priority of the V_{CC} and GND plane pairs: high priority pairs carry high transient currents and are placed high in the stackup, while low priority pairs carry lower transient currents (or can tolerate more noise) and are placed in the lower part of the stackup.

Capacitor Effective Frequency

Every capacitor has a narrow frequency band where it is most effective as a decoupling capacitor. This band is centered at the capacitor's self-resonant frequency F_{RSELF}. The effective frequency bands of some capacitors are wider than others. A capacitor's ESR determines the capacitor's quality (Q) factor, and the Q factor determines 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:

$$F = \frac{1}{2\pi\sqrt{LC}}$$

Equation 2-1

Another way to determine 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_{RSELF}, (capacitor data sheet value) is much higher than its effective mounted resonant frequency in the system, F_{RIS}. Because the mounted capacitor's performance is most important, the mounted resonant frequency is used when evaluating a capacitor as part of the greater PDS.

Mounted parasitic inductance is caused by 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. Vias in thicker boards have higher inductance.

To determine the capacitor's total parasitic inductance in the system, L_{IS}, the capacitor's parasitic inductance, L_{SELF}, is added to the mounting's parasitic inductance, L_{MOUNT}:

$$L_{IS} = L_{SELF} + L_{MOUNT}$$

Equation 2-2

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

C = 0.01 µF (selected by user)

L_{SELF} = 0.9 nH (AVX data sheet parameter)

F_{RSELF} = 53 MHz (AVX data sheet parameter)

L_{MOUNT} = 0.8 nH (based on PCB mounting geometry)

To determine the effective in-system parasitic inductance (L_{IS}), add the via parasitics:

$$L_{IS} = L_{SELF} + L_{MOUNT} = 0.9 \ nH + 0.8 \ nH$$
$$L_{IS} = 1.7 \ nH$$

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

$$F_{RIS} = \frac{1}{2\pi\sqrt{L_{IS}C}}$$

$$F_{RIS} = \frac{1}{2\pi\sqrt{(1.7\times10^{-9}H) \cdot (0.01\times10^{-6}F)}} = 38\times10^{6} Hz$$

F_{RSELF} is 53 MHz, but F_{RIS} is lower at 38 MHz. The addition of mounting inductances shifts the effective-frequency band down. A decoupling capacitor is most effective at the narrow-frequency band around its resonant frequency, and thus, the resonant frequency must be reviewed when choosing a capacitor collection to build up a decoupling network.

Capacitor Anti-Resonance

One common problem associated with capacitors in a PDS of an FPGA is anti-resonant spikes in the PDS aggregate impedance. The possible cause for these spikes is the bad combination of energy storage elements in the PDS (discrete capacitors, parasitic inductances, and power and ground planes).

- If the inter-plane capacitance of the power and ground planes has a high-Q factor, 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 (a stimulus), a large noise voltage occurs.

To correct this problem, the characteristics of the high-frequency discrete capacitors or the characteristics of the V_{CC} and ground planes must be changed.

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 also 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 2-3) occurs between the start of the disturbance at the FPGA power pins and the point where the capacitor starts sensing the disturbance.

$$Time\ Delay\ =\ \frac{Distance\ from\ the\ FPGA\ power\ pins\ to\ the\ capacitor}{Signal\ propagation\ speed\ through\ FR4\ dielectric} \qquad \textit{Equation 2-3}$$

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 specific frequency's wavelength.

- Energy transferred to the FPGA increases from 0% at one-quarter of a wavelength distance 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 discrete 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, only one capacitor per pin is used. 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 little low-frequency energy is needed, larger capacitors are not necessary.

Power Supply Consolidation

Powering 2.5V V_{CCO} and V_{CCAUX} from a common PCB plane is allowed in Virtex-5 FPGA designs.

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. However, ESD events at the unconnected solder balls in the inner rows of the pinout array are unlikely and not considered a high risk.

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, and though it does not account for the distributed behavior of a PDS, it is a useful tool for selecting and verifying combinations of decoupling capacitor values.

Lumped RLC simulation is performed either in a version of SPICE or other circuit simulator, or by using a mathematical tool like MathCAD or Microsoft Excel. Istvan Novak publishes a free Excel spreadsheet for lumped RLC simulation (among other useful tools for PDS simulation) on his website:
http://www.electrical-integrity.com/

Table 2-14 also lists a few EDA tool vendors for PDS design and simulation. These tools span a wide range of sophistication levels.

Table 2-14: **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	Sigrity	http://www.sigrity.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 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. This measurement technique shows the crosstalk (via field, PCB routing, package routing) induced on the victim as well as the die-level power system noise.

Oscilloscope Measurement Methods:

- 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).

Noise measurements should be made at a few different FPGA locations to ensure that any local noise phenomena are captured.

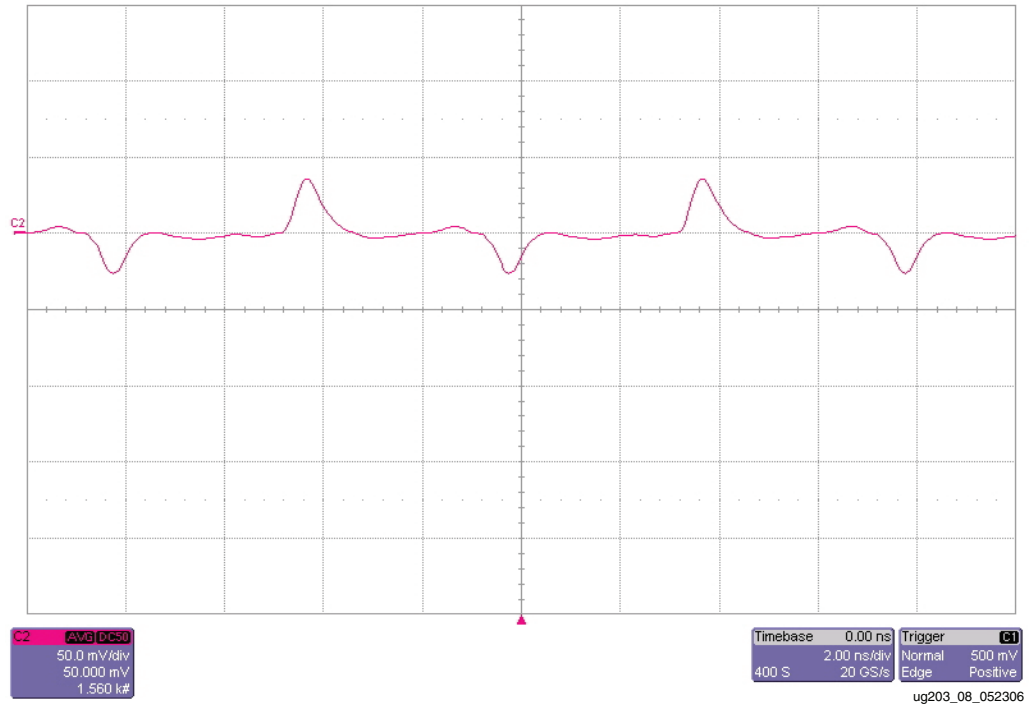Figure 2-8 shows an averaged noise measurement taken at the V_{CCINT} pins of a sample design.

ug203_08_052306

Figure 2-8: **Averaged Measurement of V$_{CCO}$ Supply with Multiple I/O Sending Patterns at 100 MHz**

Figure 2-9 shows an infinite persistence noise measurement of the same design. 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.
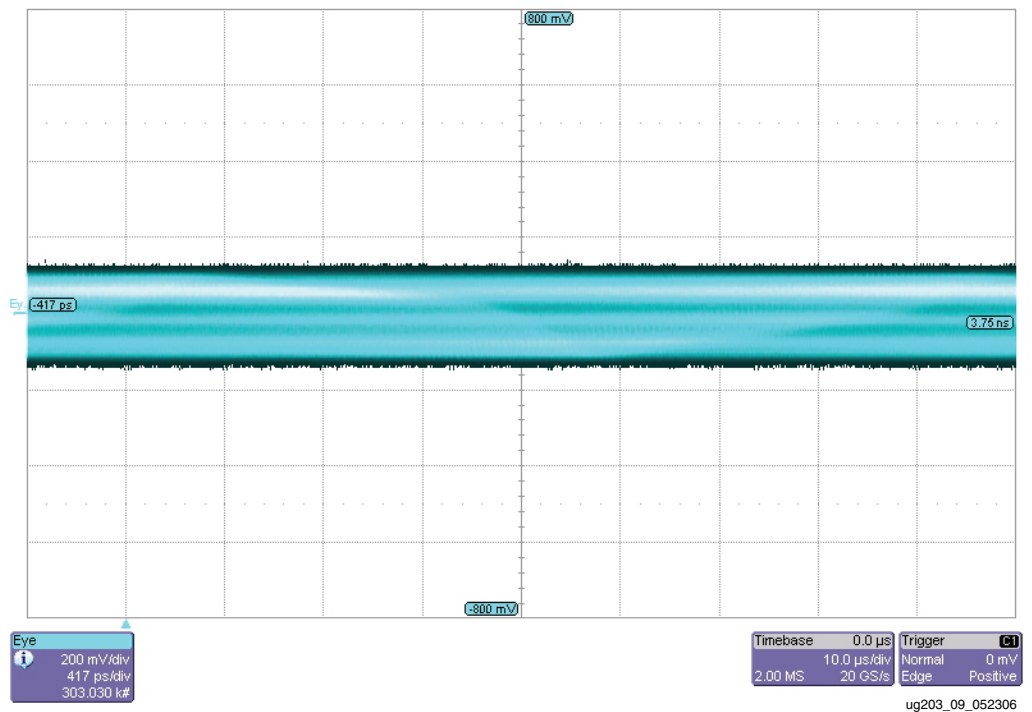

ug203_09_052306

Figure 2-9: **Infinite Persistence Measurement of Same Supply**

The measurement shown in Figure 2-8 and Figure 2-9 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 with FFT math functionality work well.

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. The noise frequency content can be approximated by looking at 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. The user sees the exact frequencies where the PDS is inadequate when a spectrum analyzer is used.

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 resonant frequencies close to the noise frequency or otherwise lowering the PDS impedance at the critical frequency.

The noise spectrum measurement should be taken at the same place 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 (of noise in the power planes) by removing a decoupling capacitor in the device vicinity and soldering the cable's center conductor and shield directly to the capacitor lands. Alternatively, a probe station with 50Ω RF probes can be used.

- Add a DC blocking capacitor or attenuator to protect the spectrum analyzer from the device supply voltage.

Figure 2-10 shows an example of a noise spectrum measurement. It is a spectrum-analyzer measurement screenshot of the V_{CCO} power-supply noise, with multiple I/O sending patterns at 150 MHz.
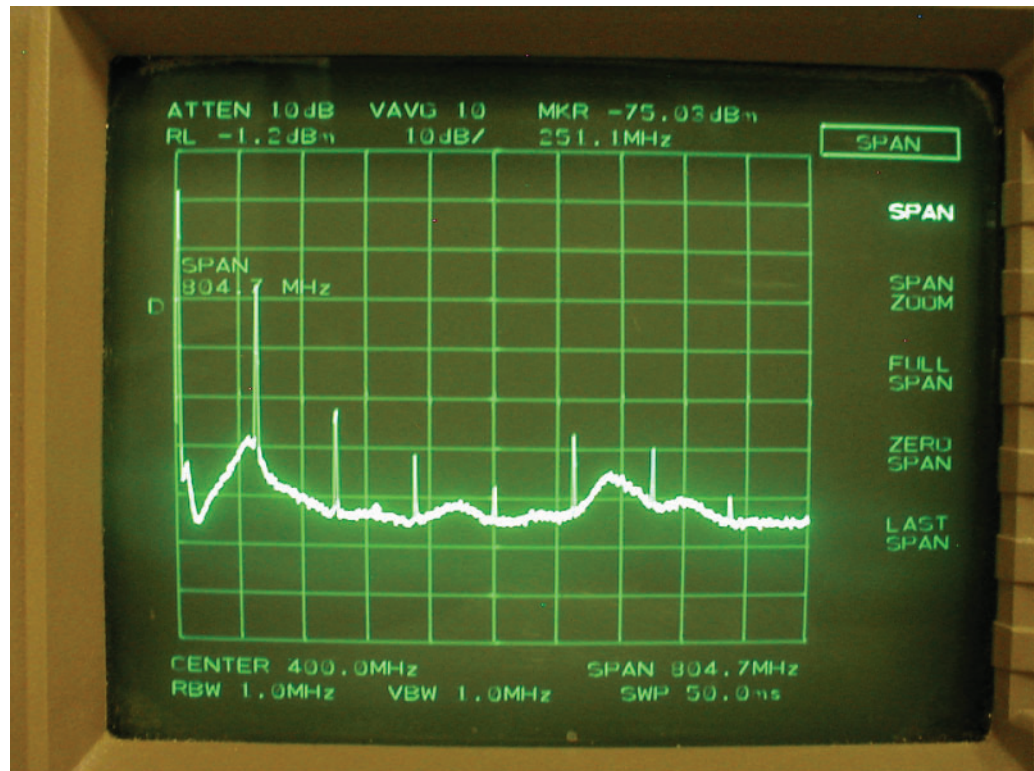
UG203_10_112806

Figure 2-10: **Screenshot of Spectrum Analyzer Measurement of V$_{CCO}$**

Optimum Decoupling Network Design

If a highly optimized PDS is needed, more measurements help create a customized decoupling network design.

- A network analyzer measures the impedance profile of a prototype PDS. The network analyzer sweeps a stimulus across a range of frequencies and measures the PDS impedance at each frequency.

  The network-analyzer measurement output can be converted to impedance as a function of frequency.

- A spectrum-analyzer measurement output is voltage as a function of frequency.

Using the previous two measurements, transient current as a function of frequency (Equation 2-4) is determined:

$$I(f) \; = \; \frac{V(f) \; From \; Spectrum \; Analyzer}{Z(f) \; From \; Network \; Analyzer} \qquad \qquad Equation \; 2\text{-}4$$

Understanding the design's transient current requirements improves the designer's PDS choices. 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

Sometimes a design with the required noise specifications cannot be created by using the methods described in this document, and system aspects should be analyzed for possible changes.

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 because of temporary periodic contention and high-current drivers

- Large microprocessors

When unacceptable noise amounts 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

In this possibility, 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 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), see 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.

Both changes used together reduce the parasitic inductance of the relevant current path.

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 also applies to 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 *and* output paths).

Possibility 4: I/O Signal Return Current Traveling in Less 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 less optimal path:

- Restrict signals to fewer available routing layers with verified continuous return current paths.
- Provide low-impedance paths for AC currents to travel between reference planes (decoupling capacitors at specific PCB locations).