

Virtex-4 FPGA PCB Designer's Guide

UG072 (1.2) June 24, 2008





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.

© 2004–2008 Xilinx, Inc. All rights reserved.

XILINX, the Xilinx logo, the Brand Window, and other designated brands included herein are trademarks of Xilinx, Inc. PowerPC is a trademark of IBM Corp. and is used under license. All other trademarks are the property of their respective owners.

Revision History

The following table shows the revision history for this document.

	Version	Revision
08/02/04	1.0	Initial Xilinx release. Chapter 15 in printed Handbook is Chapter 4 in version 1.1.
09/09/04	1.1	Added Preface and Chapters 1, 2, and 3. Minimal revisions to Chapter 4.
06/24/08	1.2	Removed SSTL3 rows from Table 2-1, page 14 . Updated Appendix A, "References" .

Table of Contents

Revision History	2
Preface: About This Guide	
Guide Contents	5
Additional Documentation	5
Additional Support Resources	6
Typographical Conventions	7
Chapter 1: PCB Technology Basics	
PCB Structures	9
Traces	9
Planes	9
Vias	10
Pads and Anti-Pads	10
Lands	10
Dimensions	10
Transmission Lines	11
Return Currents	12
Chapter 2: SelectIO Signaling	
Interface Types	13
Single-Ended versus Differential Interfaces	13
SDR versus DDR Interfaces	14
Single-Ended Signaling	14
Modes and Attributes	16
Input Thresholds	16
Topologies and Termination	16
Unidirectional Topologies and Termination	17
Bidirectional Topology and Termination	21
Bidirectional Multi-Point Topologies	25
Chapter 3: Multi-Gigabit Serial Signaling	
Chapter 4: Power Distribution System	
Power Distribution Design Overview	29
Basic Decoupling Network Principles	30
What is the Role of Inductance?	31
Capacitor Parasitic Inductance	32
Inductance from PCB Current Paths	33
Mounting Inductance	33
Plane Inductance	35
PCB Stackup and Layer Order	36

Capacitor Effective Frequency	37
Capacitor Anti-Resonance	38
Capacitor Placement	38
PDS Design and Verification	42
Step 1: Determining Critical Parameters of the FPGA	42
Step 2: Designing the Generic Bypassing Network	42
Step 3: Simulation	44
Step 4: Building the Design	45
Step 5: Measuring Performance	45
Noise Magnitude Measurement	46
Noise Spectrum Measurements	48
Step 6: Optimum Bypassing Network Design (Optional)	49
Other Concerns and Causes	50
Possibility 1: Excessive Noise from Other Devices on Board	50
Possibility 2: Parasitic Inductance of Planes, Vias, or Connecting Traces	50
Possibility 3: I/O Signals in PCB are Stronger Than Necessary	51
Possibility 4: I/O Signal Return Current Travelling in Sub-Optimal Paths	51
Calculation of Via Inductance	52
SPICE Simulation Examples	52
HSPICE Netlist	53
HSPICE Output	53
Schematic Circuit	54
EDA Tools for PDS Design and Simulation	54

Appendix A: References

About This Guide

This designer's guide provides information on the design of PCBs for Virtex®-4 devices. It considers all aspects of the PCB from the system level down to the minute details. Instead of providing information on device features and specifications, this design guide focuses on strategies for making design decisions at the PCB and interface level.

Guide Contents

This manual 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, "SelectIO Signaling"](#)
Contains information on the choice of SelectIO™ standards, I/O topologies, and termination strategies as well as information on simulation and measurement techniques.
- [Chapter 3, "Multi-Gigabit Serial Signaling"](#)
Provides PCB information and electrical requirements related to serial transceivers.
- [Chapter 4, "Power Distribution System"](#)
Covers the power distribution system for Virtex-4 FPGAs, including all details of decoupling capacitor selection, use of voltage regulators and PCB geometries, simulation and measurement.
- [Appendix A, "References"](#)
Provides additional information on references listed in this document.

Additional Documentation

The following documents are also available for download at <http://www.xilinx.com/virtex4>.

- Virtex-4 Family Overview
The features and product selection of the Virtex-4 family are outlined in this overview.
- Virtex-4 FPGA Data Sheet: DC and Switching Characteristics
This data sheet contains the DC and Switching Characteristic specifications for the Virtex-4 family.

- Virtex-4 FPGA User Guide
This guide includes chapters on:
 - ◆ Clocking Resources
 - ◆ Digital Clock Manager (DCM)
 - ◆ Phase-Matched Clock Dividers (PMCD)
 - ◆ Block RAM and FIFO memory
 - ◆ Configurable Logic Blocks (CLBs)
 - ◆ SelectIO Resources
 - ◆ SelectIO Logic Resources
 - ◆ Advanced SelectIO Logic Resources
 - ◆ System Monitor
- XtremeDSP™ Design Considerations
This guide describes the XtremeDSP slice and includes reference designs for using DSP48 math functions and various FIR filters.
- Virtex-4 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-4 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.
- Virtex-4 RocketIO Multi-Gigabit Transceiver User Guide
This guide describes the RocketIO Multi-Gigabit Transceivers available in the Virtex-4-FX family.
- Virtex-4 FPGA Embedded Tri-Mode Ethernet MAC User Guide
This guide describes the Tri-mode Ethernet Media Access Controller available in the Virtex-4 FX family.
- PowerPC® 405 Processor Block Reference Guide
This guide is updated to include the PowerPC 405 processor block available in the Virtex-4 FX family.

Additional Support Resources

To search the database of silicon and software questions and answers, or to create a technical support case in WebCase, see the Xilinx website at:

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

Typographical Conventions

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

Convention	Meaning or Use	Example
<i>Italic font</i>	References to other documents	See the <i>Virtex-4 FPGA Configuration Guide</i> for more information.
	Emphasis in text	The address (F) is asserted <i>after</i> clock event 2.
<u>Underlined Text</u>	Indicates a link to a web page.	http://www.xilinx.com/virtex4

PCB Technology Basics

Printed Circuit Boards (PCBs) are electrical systems, having electrical properties just as complicated as the discrete components and devices mounted to them. The PCB designer has complete control over many aspects of the PCB. At the same time, current technology places constraints and limits on the geometries and resulting electrical properties. The following information is provided first as a primer, but also 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 insulative substrate material (usually FR4, an epoxy/glass composite) with copper plating on both sides has portions of copper etched away to form conductive paths. A number of layers of plated and etched substrates are glued together in a stack with additional insulative substrates in-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 the areas of material properties, number of stacked layers, geometries, and drilling techniques (allowing holes that only penetrate only a portion of the stackup), the basic structures of PCBs have not changed. These structures, formed through the processes outlined above, 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. Traces carry signals between these points.

Planes

A plane is an uninterrupted area of metal covering the entirety of a 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.

Vias

A via is a physical 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. A via is formed in current microvia technology (also known as High Density Interconnect or HDI) with a laser by ablating the substrate material and deforming the conductive plating, in the process forming a conductive connection. These types of Microvias formed 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 Anti-Pads

Pads are small areas of copper in prescribed shapes. Anti-pads 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. Anti-pads are used mainly with vias.

Since plated through-hole vias are conductive over the whole length, a method is needed to selectively making electrical connections to traces, planes and planelets of the various layers of a PCB. This is the function of pads and anti-pads.

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.

Anti-pads are used in a similar capacity but in reverse. Since plane and planelet copper is otherwise uninterrupted, any via travelling through it makes an electrical connection to it. Where vias are not intended to make an electrical connection to the planes or planelets passed through, an anti-pad 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, and minimum dimension specifications exist defining the minimum length of the connecting trace. 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 FPGA package geometries, PCB manufacturing limits, and system compliance. Other factors such as Design For Manufacturing (DFM) and reliability impose further limits, but as these are application specific, they are not covered here.

The dimensions of the FPGA package, in combination with PCB manufacturing limits define most of the geometric aspects of the above structures, both directly and indirectly. This in itself constrains the PCB designer significantly. The package ball pitch (1.0 mm for FF packages; 0.8 mm for SF 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 defined by the

PCB manufacturer limit the amount of space available in-between vias for routing of signals in and out of the via array underneath the device. This defines the maximum trace width in these “breakout” traces. PCB manufacturing limits constrain the minimum trace width and minimum spacing.

The number of signal 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 are defined by the number of I/O signal traces to 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 50 Ω single-ended or 100 Ω differential).

Bypass 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 may be placed. The purpose of this is to allow room for rework where necessary. For this reason, the area just outside the keep-out area is one where components will compete for placement. It is up to the PCB designer to determine the high priority components. Bypass capacitor placement constraints are described in [Chapter 4, “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.

SelectIO Signaling

The SelectIO™ term in the Virtex-4 device family refers to the general-purpose I/O and its various settings. With over 20 I/O standards and over 110 variants within these standards, the highly flexible SelectIO resource offers a wide array of choices for designing I/O interfaces.

This chapter provides some strategies for choosing I/O standard, topology, 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 present, 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 interfaces versus Differential interfaces
- Single Data Rate (SDR) interfaces versus Double Data Rate (DDR) interfaces

Single-Ended versus Differential Interfaces

Traditional digital logic uses single-ended signaling – a convention that transmits a signal and assumes a Ground 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 Ground. 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.

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 relative to Ground, 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 may be SDR or DDR, and a differential interface may also be SDR or DDR.

Single-Ended Signaling

A variety of single-ended I/O standards are available in the Virtex-4 FPGA IOB configuration options. Table 2-1 lists all available single-ended I/O standards. The checkmarks indicate which features are available or can be configured for each I/O standard.

Table 2-1: Virtex-4 FPGA Single-Ended I/O Standards

I/O Standard	V _{REF} Input Threshold	Configurable Attribute					Recommended Mode	
		Drive Strength	Slew Rate	Pull-up	Pull-down	Weak Keeper	Unidirectional	Bidirectional
LVTTTL		√	√	√	√	√	√	√
LVC MOS33		√	√	√	√	√	√	√
LVC MOS25		√	√	√	√	√	√	√
LVC MOS18		√	√	√	√	√	√	√
LVC MOS15		√	√	√	√	√	√	√
LVDCI33							√	√
LVDCI25							√	√
LVDCI18							√	√
LVDCI15							√	√
HSLVDCI33	√							√
HSLVDCI25	√							√
HSLVDCI18	√							√
HSLVDCI15	√							√
HSTL Class I	√						√	
HSTL Class II	√						√	√
HSTL Class III	√						√	
HSTL Class IV	√						√	√
HSTL18 Class I	√						√	
HSTL18 Class II	√						√	√
HSTL18 Class III	√						√	

Table 2-1: Virtex-4 FPGA Single-Ended I/O Standards (Continued)

I/O Standard	V _{REF} Input Threshold	Configurable Attribute					Recommended Mode	
		Drive Strength	Slew Rate	Pull-up	Pull-down	Weak Keeper	Unidirectional	Bidirectional
HSTL18 Class IV	√						√	√
HSTL Class I DCI	√						√	
HSTL Class II DCI	√						√	√
HSTL Class III DCI	√						√	
HSTL Class IV DCI	√						√	√
HSTL18 Class I DCI	√						√	
HSTL18 Class II DCI	√						√	√
HSTL18 Class III DCI	√						√	
HSTL18 Class IV DCI	√						√	√
SSTL2 Class 1	√						√	
SSTL2 Class 2	√						√	√
SSTL18 Class 1	√						√	
SSTL18 Class 2	√						√	√
SSTL2 Class 1 DCI	√						√	
SSTL2 Class 2 DCI	√						√	√
SSTL18 Class 1 DCI	√						√	
SSTL18 Class 2 DCI	√						√	√
GTL	√						√	√
GTL+	√						√	√
GTL DCI	√						√	√
GTL+ DCI	√						√	√
PCI-33				√	√	√	√	√
PCI-66				√	√	√	√	√
PCI-X				√	√	√	√	√

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). Drive strength and slew rate can be used to tune an interface for adequate speed while not overdriving the signals. Weak pull-ups, weak pull-downs, and weak keepers can be used to ensure a known or steady level on a floating or 3-stated signal. [Table 2-1](#) denotes which standards support these attributes. See the SelectIO chapter in the *Virtex-4 FPGA User Guide* for more information.

LVC MOS_{6F} drivers and other “weak” drivers have an output impedance close to 50 Ω, allowing them to be used as a crude approximation of a controlled-impedance driver. Note that the impedance match of the weak driver to the transmission line is only approximate and varies with voltage and temperature. LVDCI, a true controlled-impedance driver, is adaptive, maintains a much closer impedance match, and remains constant over voltage and temperature.

Input Thresholds

The input circuitry of the single-ended standards listed in [Table 2-1](#) 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 Ground 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 2.5V I/O standards that illustrate this are LVC MOS₂₅ and SSTL2 Class 1. The thresholds for 2.5V LVC MOS are set at 0.7V and 1.7V (necessitating that the signal at the receiver swing a full 1.0V at minimum to make a logic transition). The thresholds for SSTL2 Class 1 are set at $V_{REF} - 0.15V$ and $V_{REF} + 0.15V$, or for a nominal V_{REF} of 1.25V, set at 1.1V and 1.4V (necessitating that the signal at the receiver only swing 0.3V at minimum to make a logic transition). This smaller required swing allows for higher frequency of operation in the overall link. A smaller swing at the driver means reduced DC power is required with less transient current. 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 4, “Power Distribution System.”](#)

Topologies and Termination

Topology generally refers to the arrangement of drivers, receivers, interconnect and terminations in an interface. The techniques used in unidirectional topologies are different from those used in bidirectional topologies, so these are treated separately.

The SelectIO standards can be used in countless topologies depending on the requirements of the system. SelectIO drivers and receivers adhering to a standard (SSTL, LVC MOS, etc.) either can be used according to the letter of the standard (published by a standards body such as EIA/TIA or JEDEC) or they can be mixed and matched with drivers or receivers from another standard or hybrid I/O. An I/O standard specification might define

something as limited as the V_{IL} and V_{IH} of the receiver, or it might define every aspect of the interface, including driver impedance and slew rate, PCB trace length and topology, value and position of passive termination, the maximum input capacitance of a receiving device, and even the maximum number of receivers.

It is up to the designer to apply the standard in question to the system in which it is working. There are many decisions to make with respect to topologies and termination, which determine the signal integrity of the interface. It is of utmost importance that the signal integrity of each interface is verified through both simulation and measurement.

Termination generally refers to impedance-matching or impedance-compensating devices that are used to maintain signal integrity in an interface. While many types of elements can be used as *terminators* (such as, resistors, capacitors, diodes), this discussion is limited to resistive termination. In general, capacitor and diode termination techniques are more complicated and are not covered here.

Unidirectional Topologies and Termination

The two basic subsets of unidirectional topologies are point-to-point and multi-drop. A point-to-point topology has one driver and one receiver, while a multi-drop topology has one driver and many receivers. Whether or not a topology is point-to-point or multi-drop defines important aspects of the interface that determine which termination strategies are appropriate and which are not.

Unidirectional Point-to-Point Topologies

The simplest unidirectional topology is point-to-point. That is, there is one driver and one receiver. Termination, if present, can consist of parallel termination at the receiver (Figure 2-1), series termination at the driver (Figure 2-2), or a controlled-impedance driver (Figure 2-3 and Figure 2-4). Always use IBIS simulation to determine the optimal resistor values, V_{TT} voltage level, and VRN/VRP reference resistors for these terminations.




Figure 2-1: **Parallel-Terminated Unidirectional, Point-to-Point Topology**




Figure 2-2: **Series-Terminated Unidirectional, Point-to-Point Topology**




Figure 2-3: **DCI-Controlled Impedance Driver Unidirectional, Point-to-Point Topology**




Figure 2-4: **“Weak Driver” Unidirectional, Point-to-Point Topology**

In general, parallel resistive termination (R_P) has a value equal to the characteristic impedance (Z_0) of the transmission line it is terminating. Series resistive terminations (R_S) have a value equal to the characteristic impedance of the transmission line (Z_0) minus the output impedance of the driver (R_O) to which they are connected. Controlled-impedance drivers are tuned such that the driver output impedance (R_O) is equal to the characteristic impedance (Z_0) of the transmission line it is terminating.

Assuming transmission lines with 50 Ω characteristic impedance and a driver output impedance (R_O) of 25 Ω , a 25 Ω series termination (Figure 2-2) or a 50 Ω parallel termination (Figure 2-1) is appropriate. Controlled-impedance drivers, whether implemented with DCI or with weak LVC MOS drivers, should be sized to have an output impedance (R_O) of 50 Ω . This corresponds to VRN and VRP resistors equal to 50 Ω for DCI. Weak LVC MOS drivers of 6 mA to 8 mA drive strength have an output impedance approximately equal to 50 Ω (Figure 2-3).

Typically, parallel terminations have best performance when V_{TT} (the voltage source connected to the parallel termination resistor) is equal to half of the signaling voltage. For 2.5V signals ($V_{CCO} = 2.5V$), V_{TT} is ideally 1.25V. In cases where this voltage is not available, it is advisable to use a Thevenin parallel termination. Thevenin parallel termination consists of a voltage divider with a parallel equivalent resistance (R_{PEQ}) equal to the characteristic impedance of the transmission line (50 Ω in most cases). The divided voltage point is designed to be at V_{TT} . Figure 2-5 illustrates a Thevenin parallel termination powered from 2.5V V_{CCO} , made up of two 100 Ω resistors, resulting in a V_{TT} of 1.25V and a parallel equivalent resistance (R_{PEQ}) of 50 Ω .

Parallel termination can be less desirable than series termination or controlled-impedance drivers because it dissipates more power. This trade-off must be weighed against other trade-offs to determine the optimum termination topology for an interface.




Figure 2-5: Thevenin Parallel Termination

Table 2-2 lists I/O interface types that can be used with the unidirectional point-to-point topology.

Table 2-2: I/O Interface Type for Unidirectional Point-to-Point Topologies

LVTTTL
LVC MOS
LVDCI
SSTL Class I
HSTL Class I
HSTL Class III
GTL
GTL+
GTLDCI
GTL+DCI

LVTTTL and LVC MOS do not specify any canonical termination method. Series termination at the driver or parallel termination at the receiver are both appropriate.

LVDCI implicitly uses controlled-impedance driver termination. No form of termination is needed at the receiver.

HSTL Class I and HSTL Class III specify a parallel termination at the receiver. In the case of HSTL Class I, the termination voltage V_{TT} is defined as half of the supply voltage V_{CC} . In the case of HSTL Class III, the termination voltage V_{TT} is defined as equal to the supply voltage V_{CC} . The designer may elect either not to use termination at all or to use a different termination, such as series termination at the driver. There are a number of reasons why this selection might be advantageous in a given system. It is up to the designer to verify through simulation and measurement that the signal integrity at the receiver is adequate.

SSTL Class I specifies both series termination at the driver and parallel termination at the receiver. The termination voltage V_{TT} is defined as half of the supply voltage V_{CC} . The designer may elect either not to use termination at all or to use a different termination, such as only series termination at the driver or only parallel termination at the receiver. There are a number of reasons why this might be advantageous in a given system. It is up to the

designer to verify through simulation and measurement that the signal integrity at the receiver is adequate.

GTL and GTL+ specify and require parallel termination both at the driver and at the receiver. The designer **may not** elect to omit either of these terminations because the driver is open drain, meaning that the driver is not capable of driving a High signal and depends on the parallel termination to generate High logic levels.

Unidirectional Multi-Drop Topologies

In more complex topologies, a single driver can drive multiple receivers. The receivers represent loads that must be fed by individual transmission line stubs. From a signal integrity standpoint, the best topology to use in this case is a single long transmission line with the driver at one end and parallel termination at the other, with receivers connected to the main trace by short stubs in between. This type of topology is often referred to as a *flyby multi-drop* topology.

There are two critical aspects of this topology. The first is the presence of parallel termination at the far end of the transmission line. Series termination at the driver must *never* be used. Parallel termination is the only applicable termination type for this topology. The second critical aspect is the length of the connecting stubs at each receiver. These must remain short: no more than 8 mm in length. As the stubs become longer, they present a larger impedance discontinuity to the signal travelling down the transmission line, and can support significant reflections. These impedance discontinuities corrupt the signal. With increasing numbers of loads and increasing length of stubs, the signal is corrupted to the point where it is no longer usable.

Star topologies are not recommended. The constraints involved in designing a star topology with good signal integrity are beyond the scope of this document.

As stated in “[Unidirectional Point-to-Point Topologies](#)”, ideal parallel resistive termination has a value equal to the characteristic impedance of the transmission line it is terminating. The best performance is achieved when V_{TT} is equal to half of the signaling voltage, and when this voltage is not available, a Thevenin parallel termination is recommended, as defined in the previous subsection.

[Figure 2-6](#) illustrates a Thevenin parallel termination powered from V_{CC0} , made up of two 100 Ω resistors, resulting in a V_{TT} of $V_{CC0}/2$ and a parallel equivalent resistance of 50 Ω . This figure shows a topology with one driver (an LVCMOS driver) and four receivers. The driver is on the left side, the receivers are spaced at interim points across the 50 Ω transmission line, and the Thevenin parallel termination of two 100 Ω resistors is on the right side.

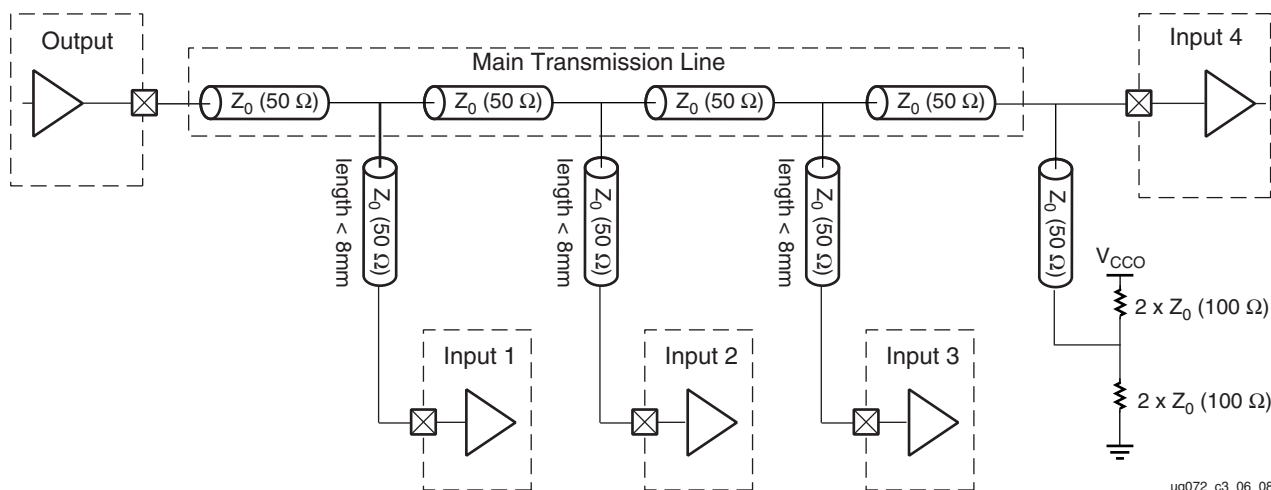


Figure 2-6: Basic Multi-Drop Topology

The main transmission line should be kept as short as possible. Lengths up to 20 inches or more are practical for most I/O standards as long as precise trace impedance is maintained and crosstalk sources are avoided. The lengths of interim segments of the main transmission line need not be equal. Their relative lengths may be arbitrary. Receivers at different points along the main transmission line receive the signal with varying amounts of delay, but all signal rise times are similar.

Stubs stretching from the main transmission line to the individual receivers must be kept as short as possible. The longer these stubs become, the more corrupted the received waveforms are. Simulation and measurement are required to assess signal integrity at the individual receivers.

Table 2-3 lists I/O interface types that can be used with the unidirectional point-to-point multi-drop topology.

Table 2-3: I/O Interface Types for Multi-Drop Point-to-Point I/O Topologies

LVTTL
LVC MOS

LVTTL and LVC MOS do not specify any canonical termination method. Parallel termination at the end of the long t-line is the only appropriate termination method.

Bidirectional Topology and Termination

The two basic subsets of bidirectional topologies are point-to-point and multi-point. A point-to-point topology has two transceivers (driver and receiver sharing one device pin), while a multi-point topology may have many transceivers. Whether or not a topology is point-to-point or multi-point defines important aspects of the interface that determine which termination strategies are appropriate and which are not.

Bidirectional Point-to-Point Topologies

The simplest bidirectional topology is point to point. That is, there are two transceivers connected by a transmission line. Because bidirectional interfaces need to operate equally well in both directions, symmetry of the topology is desirable. While asymmetrical topologies can be designed with reasonably good signal integrity, the easiest way to ensure

good signal integrity is to keep the topology symmetrical. Thus any termination used on one side of the link should also be used on the other side of the link. Series termination (Figure 2-8) is rarely appropriate for bidirectional interfaces as incoming signals are attenuated by the series resistor of the receiving transceiver. Parallel termination (Figure 2-7) always achieves better signal levels at both receivers. Controlled-impedance drivers, whether crudely controlled in the form of a weak LVCMOS driver or adaptively controlled in the form LVDCI or HSLVDCI, also can have good results as shown in Figure 2-9, Figure 2-10, and Figure 2-11 (implemented with a low-drive strength LVCMOS driver). Always use IBIS simulation to determine the optimal resistor, V_{TT} voltage level and VRN/VRP reference resistor values for these terminations.




Figure 2-7: Parallel Terminated Bidirectional Point-to-Point Topology




Figure 2-8: Series Terminated Bidirectional Point-to-Point Topology: Not Recommended

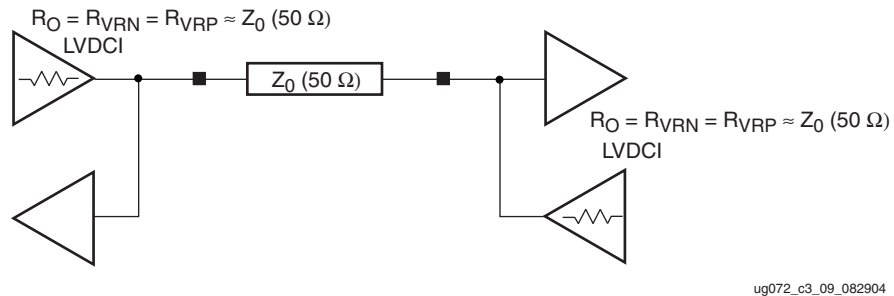
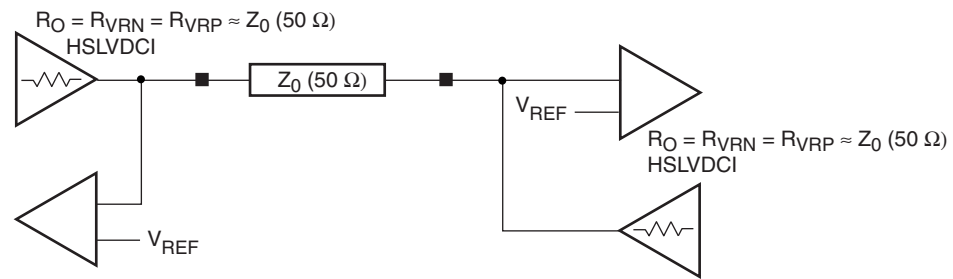



Figure 2-9: DCI Controlled Impedance Bidirectional Point-to-Point Topology



ug072_c3_10_082904

Figure 2-10: HSLVDCI Controlled Impedance Driver Bidirectional Point-to-Point Topology



ug072_c3_11_082904

Figure 2-11: "Weak Driver" Bidirectional Point-to-Point Topology

In general, parallel resistive termination (R_P) has a value equal to the characteristic impedance Z_0 of the transmission line it is terminating. Controlled-impedance drivers are tuned such that the driver output impedance (R_O) is equal to the characteristic impedance (Z_0) of the transmission line it is terminating.

Assuming transmission lines with 50 Ω characteristic impedance and a driver output impedance of 25 Ω , 50 Ω parallel terminations are appropriate (Figure 2-7). Controlled-impedance drivers, whether implemented with DCI or with weak LVC MOS drivers, should be sized to have an output impedance (R_O) of 50 Ω . This corresponds to VRN and VRP resistors equal to 50 Ω for DCI (Figure 2-9 and Figure 2-10). Weak LVC MOS drivers of 6 mA to 8 mA drive strength have an output impedance approximately equal to 50 Ω (Figure 2-11).

Typically, parallel terminations have the best performance when V_{TT} (the voltage source connected to the parallel termination resistor) is equal to half of the signaling voltage. For 2.5V signals ($V_{CC0} = 2.5V$), V_{TT} is ideally 1.25V. In cases where this voltage is not available, it is advisable to use a Thevenin parallel termination. Thevenin parallel termination consists of a voltage divider with a parallel resistance equal to the characteristic impedance of the transmission line (50 Ω in most cases). The divided voltage point is designed to be at V_{TT} . Figure 2-12 illustrates a Thevenin parallel termination powered from 2.5V V_{CC0} , made up of two 100 Ω resistors, resulting in a V_{TT} of 1.25V and a parallel equivalent resistance (R_{PEQ}) of 50 Ω .

Parallel termination can be less desirable than series termination or controlled-impedance drivers because it dissipates more power. This trade-off must be weighed against other trade-offs to determine the optimum termination topology for an interface.




Figure 2-12: Thevenin Parallel Termination (Bidirectional Point-to-Point Topology)

Table 2-4 lists I/O interface types that can be used with the bidirectional point-to-point topology.

Table 2-4: I/O Interface Types for Bidirectional Point-to-Point I/O Topologies

LVTTTL
LVC MOS
LVDCI
HSLVDCI
SSTL CLASS II
SSTL CLASS II DCI
HSTL CLASS II
HSTL CLASS II DCI
HSTL CLASS IV
HSTL CLASS IV DCI
GTL
GTL+
GTL DCI
GTL + DCI

LVTTTL and LVC MOS do not specify any canonical termination method. Series termination is not recommended for bidirectional interfaces. Parallel termination and weak drivers, however, are both appropriate.

LVDCI and HSLVDCI both implicitly use controlled-impedance driver termination.

HSTL Class II and HSTL Class IV specify parallel termination at both transceivers. In the case of HSTL Class II, the termination voltage V_{TT} is defined as half of the supply voltage V_{CC0} . In the case of HSTL Class IV, the termination voltage V_{TT} is defined as equal to the supply voltage V_{CC0} . The designer may elect either not to use termination at all or to use a different termination. It is up to the designer to verify through simulation and measurement that the signal integrity at the receiver is adequate.

SSTL Class II specifies both series termination and parallel termination. The termination voltage V_{TT} is defined as half of the supply voltage V_{CCO} . While the canonical standard calls for this, better performance is achieved when only parallel termination is used at both transceivers. The designer may elect to use either the canonical standard termination or a modified termination (or no termination at all). It is up to the designer to verify through simulation and measurement that the signal integrity at the receiver is adequate.

GTL and GTL+ specify and require parallel termination at both transceivers. The designer **may not** elect to omit either of these terminations because the driver is open drain, meaning that the driver is not capable of driving a High signal and depends on the parallel termination to generate High logic levels.

Bidirectional Multi-Point Topologies

In more complex topologies, any transceiver in a multi-point bus can transmit to all other transceivers. Usually these topologies can only run at very slow clock rates because they only support very slow signal rise times (10 ns to 50 ns). While useful in some situations, the drawbacks usually outweigh the benefits. The constraints involved in designing these topologies with good signal integrity are beyond the scope of this document.

Multi-Gigabit Serial Signaling

This chapter is intentionally left blank for further updates.

Power Distribution System

This chapter covers the power distribution system for Virtex-4 FPGAs, including all details of decoupling capacitor selection, use of voltage regulators and PCB geometries.

This chapter contains the following sections:

- “Power Distribution Design Overview”
- “Basic Decoupling Network Principles”
- “PDS Design and Verification”
- “Other Concerns and Causes”
- “Calculation of Via Inductance”
- “SPICE Simulation Examples”
- “EDA Tools for PDS Design and Simulation”

Power Distribution Design Overview

FPGA designers are faced with a unique task when it comes to designing power distribution systems (PDS). Most other large, dense ICs (such as large microprocessors) come with very specific bypass capacitor requirements. Since these devices are only designed to implement specific tasks in their hard silicon, their power supply demands are fixed and only fluctuate within a certain range. FPGAs do not share this property. Since FPGAs can implement a practically infinite number of applications at undetermined frequencies and in multiple clock domains, it can be very complicated to predict what their transient current demands will be.

Since exact transient current behavior cannot be known for a new FPGA design, the only choice when designing the first version of an FPGA PDS is to go with a conservative worst-case design.

Transient current demands in digital devices are the cause of ground bounce, the bane of high-speed digital designs. In low-noise or high-power situations, the power supply decoupling network must be tailored very closely to these transient current needs, otherwise ground bounce and power supply noise will exceed the limits of the device. The transient currents in an FPGA are different from design to design. This chapter provides a comprehensive method for designing a bypassing network to suit the individual needs of a specific FPGA design.

The first step in this process is to examine the utilization of the FPGA to get a rough idea of its transient current requirements. Next, a conservative decoupling network is designed to fit these requirements. The third step is to refine the network through simulation and modification of capacitor numbers and values. In the fourth step, the full design is built and in the fifth step it is measured. Measurements are made consisting of oscilloscope and possibly spectrum analyzer readings of power supply noise. Depending on the measured

results, further iterations through the part selection and simulation steps could be necessary to optimize the PDS for the specific application. A sixth optional step is also given for cases where a perfectly optimized PDS is needed.

Basic Decoupling Network Principles

Before starting into the PDS design flow, it is important to understand the basic electrical principles involved. This section discusses the purpose of the PDS and the properties of its components. It also describes important aspects of discrete capacitor placement and mounting as well as PCB geometry and stackup recommendations.

Each device in a system not only has its own wattage requirements for its operation, but also its own requirement for the cleanliness of that power. Most digital devices, including all Virtex-4 FPGAs, have a requirement that on all supplies, V_{CC} must not fluctuate more than 5% above or 5% below the nominal V_{CC} value. In this document V_{CC} is used generically to refer to all FPGA power supplies: V_{CCINT} , V_{CCO} , V_{CCAUX} , and V_{REF} . Multi-gigabit transceiver (MGT) and system monitor analog supplies ($AV_{CCAUXTX}$, $AV_{CCAUXRX}$, V_{TTX} , V_{TRX} , AV_{DD} , AV_{SS}) are not covered here. For specific instructions on these supplies, see the *Virtex-4 RocketIO Multi-Gigabit Transceiver User Guide* (Reference #1) and the *Virtex-4 FPGA User Guide* (Reference #2).

This requirement specifies a maximum amount of noise present on the power supply, often referred to as “ripple voltage.” If the device requirements state that V_{CC} must be within $\pm 5\%$ of the nominal voltage, that means peak to peak voltage ripple must be no more than 10% of the nominal V_{CC} . This assumes that nominal V_{CC} is exactly the nominal value given in the data sheet. If this is not the case, 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. Low frequency variance of power consumption is usually the result of devices or large portions of devices being enabled or disabled. This can occur on time scales from milliseconds to days. High frequency variance of power consumption is the result of individual switching events inside a device, and this happens on the scale of the clock frequency and the first few harmonics of the clock frequency.

Since 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 in power supply voltage as possible.


When the current draw in a device changes, the power distribution system cannot respond to that change instantaneously. For the short time before the PDS responds, the voltage at the device changes. This is where power supply noise appears. There are two main causes for this lag in the PDS corresponding to the two major components of the PDS.

The first major component of the PDS is the voltage regulator. It observes its output voltage and adjusts the amount of current being supplied to keep the voltage constant. Most common voltage regulators make this adjustment on the order of milliseconds to microseconds. They are effective at maintaining output voltage for events at all frequencies from DC to a few hundred kilohertz (depending on the regulator). For all transient events that occur at frequencies above this range, there is a time lag before the voltage regulator can respond to the new level of demand. For example, if current demand in the device increases in a matter of nanoseconds, the voltage at the device sags by some amount until the voltage regulator can adjust to the new, higher level of current it must provide. This lag might take from microseconds to milliseconds, during which time the voltage sags.

The second major component of the PDS is the bypass or decoupling capacitors. In this document, the words “bypass” and “decoupling” are used interchangeably. Their function

is to act as local energy storage for the device. They cannot provide DC power, as only a small amount of energy is stored in them (the voltage regulator is present to provide DC power). The function of this local energy storage is to respond very quickly to changing current demands. The capacitors are effective at maintaining power supply voltage at frequencies from hundreds of kilohertz to hundreds of megahertz in the milliseconds to nanoseconds range. Decoupling capacitors are of no use for all 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 changes and maintains this new level for a number of milliseconds, the voltage regulator circuit, operating in parallel with the bypass capacitors, effectively takes over for them, changing its output to supply this new current.


Figure 4-1 shows the major components of the PDS: the power supply, the decoupling capacitors, and the active device being powered (in this case, an FPGA).



ug072_4_01_080104

Figure 4-1: **Simplified PDS Circuit**

Figure 4-2 shows a further simplified PDS circuit, showing all reactive components decomposed to a frequency-dependent resistor.



ug072_4_02_080104

Figure 4-2: **Further Simplified PDS Circuit**

What is the Role of Inductance?

There is a property of the capacitors and of the PCB current paths that retards changes in current flow. This is the reason why capacitors cannot respond instantaneously to transient currents, nor to changes that occur at frequencies higher than their effective range. This property is called inductance.

Inductance can be thought of as momentum of charge. Where charge is moving at some rate through a conductor, this implies some amount of current. If the level of current is to change, the charge must move at a different rate. Because there is momentum (stored magnetic field energy) associated with this charge, it takes some amount of time for the charge to slow down or speed up. The greater the inductance, the greater the resistance to change, and the longer it takes for the current level to change.

The purpose of the PDS is to accommodate whatever current demands the device(s) could have, and respond to changes in that demand as quickly as possible. When these demands are not met, the voltage across the device's power supply changes. This is observed as noise. Since inductance retards the abilities of bypass capacitors to quickly respond to changing current demands, it should be minimized.

Figure 4-1 shows inductances between the FPGA device and capacitors, and between the capacitors and the voltage regulator. These inductances arise as parasitics of the capacitors themselves and of all current paths in the PCB. It is important that each of these be minimized.

Capacitor Parasitic Inductance


Of a capacitor's various properties, the capacitance value is often considered the most important. However in the domain of PCB PDS design, the property of parasitic inductance (ESL or Equivalent Series Inductance) is of the same or greater importance.

The one factor that influences parasitic inductance more than any other is the dimensions of the package. Very simply, physically small capacitors tend to have lower parasitic inductance than physically large capacitors. Just as a short wire has less inductance than a long wire, a short capacitor has less inductance than a long capacitor. Likewise, as a fat or wide wire has less inductance than a narrow wire, so too does a fat capacitor have less inductance than a narrow capacitor.

For these reasons, when choosing decoupling capacitors, the smallest package should be chosen for a given value. Similarly, for a given package size (essentially a fixed inductance value), the highest capacitance value available in that package should be chosen.

Surface-mount chip capacitors are the smallest capacitors available, making them a good choice for discrete bypass capacitors. For values from 2.2 μF down to very small values such as 0.001 μF , X7R or X5R type capacitors are usually used. These have low parasitic inductance, and an acceptable temperature characteristic. For larger values, such as 1000 μF , tantalum capacitors are used. These have low parasitic inductance and a relatively high equivalent series resistance (ESR), giving them a low-quality factor and consequently a very wide range of effective frequencies. They also provide a comparatively high capacitance value in a small package size, thus reducing board real-estate costs. In cases where tantalum capacitors are not available, low-inductance electrolytic capacitors can be used. Other new technologies with similar characteristics are also available.

A real capacitor has characteristics not only of capacitance but also inductance and resistance. Figure 4-3 shows the parasitic model of a real capacitor. A real capacitor should be treated as an RLC circuit.




ug072_4_03_080104

Figure 4-3: Parasitics of a Real, Non-Ideal Capacitor

Figure 4-4 shows the impedance characteristic of a real capacitor. Overlaid on this plot are the curves corresponding to the capacitor's capacitance and parasitic inductance (ESL).

These two curves combine to form the total impedance characteristic of the RLC circuit formed by the parasitics of the capacitor.



ug072_4_04_08010

Figure 4-4: **Contribution of Parasitics to Total Impedance Characteristics**

As capacitive value is increased, the capacitive curve moves down and to the left. As parasitic inductance is decreased, the inductive curve moves down and to the right. Since parasitic inductance for capacitors in a given package is essentially fixed, the inductance curve remains fixed. As different capacitor values are selected in that same package, the capacitive curve moves up and down relative to the fixed inductance curve. The only way to decrease the total impedance of a capacitor if the package is fixed is to increase the value of the capacitor. The only way to move the parasitic inductance curve down (and consequently lower the total impedance characteristic), is to connect additional capacitors in parallel.

Inductance from PCB Current Paths


The parasitic inductance of current paths in the PCB have two distinct sources: the capacitor mounting, and the power and ground planes of the PCB.

Mounting Inductance

In this context, the mounting refers to the capacitor's solder land on the PCB, the trace (if any) between the land and via, and the via itself.

The vias, traces, and pads of a capacitor mounting contribute anywhere from 300 pH to 4 nH of inductance depending on the specific geometry. Since the inductance of a current path is proportional to the area of the loop the current traverses, it is important to minimize the size of this loop. 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 4-5.

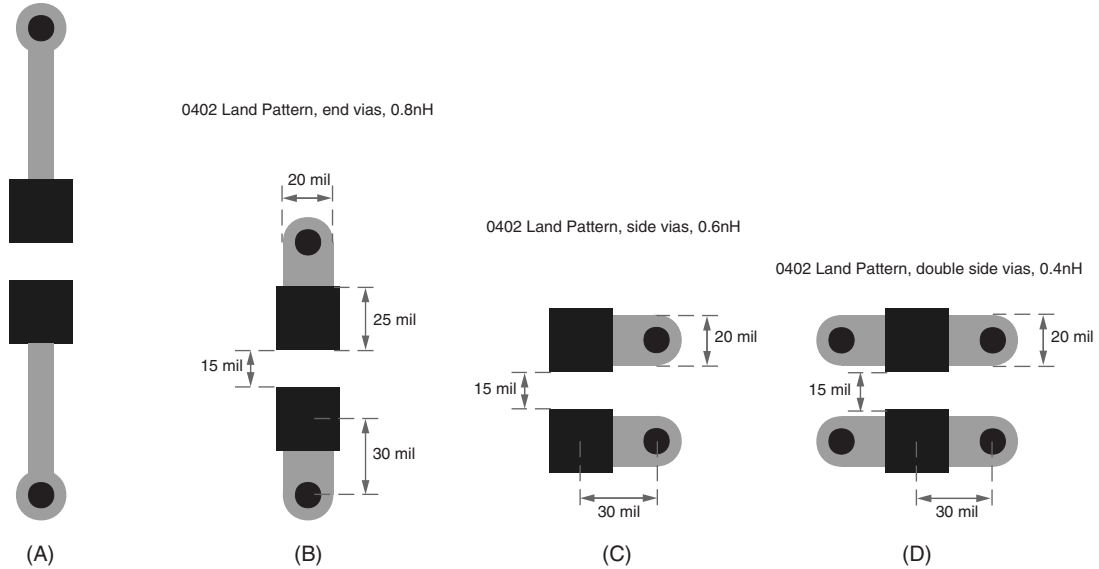


ug072_4_05_080104

Figure 4-5: Cutaway View of PCB with Capacitor Mounting

By shortening the connecting traces, the area of this loop is minimized and the inductance is reduced. Similarly, by reducing the via length through which the current flows, loop area is minimized and inductance is reduced.

0402 Land Pattern, end vias, long traces, 4nH - BAD



ug072_4_06_080104

Figure 4-6: Example Capacitor Land and Mounting Geometries

The existence and/or length of a connecting trace has a big impact on parasitic inductance of the mounting. Wherever possible, there should be no connecting trace (Figure 4-6a) - the via should butt up against the land itself (Figure 4-6b). Additionally, the connecting trace should be made as wide as possible. Further improvements can be made to the mounting by placing vias to the side of capacitor lands (Figure 4-6c), or doubling the number of vias (Figure 4-6d). Currently, very few PCB manufacturing processes allow via-in-pad geometries, but this is another good option. The technique of using multiple vias per land is important when using ultra-low inductance capacitors, such as reverse aspect ratio capacitors (AVX's LICC).

Many times in an effort to squeeze more parts into a small area, PCB layout engineers opt to share vias among multiple capacitors. **This technique should not be used under any circumstances.** The capacitor mounting (lands, traces, and vias) typically contributes about the same amount or more inductance than the capacitor's own parasitic inductance. If a second capacitor is connected into the vias of an existing capacitor, it only improves the PDS by a very small amount. It is better to reduce the total number of capacitors and maintain a one-to-one ratio of lands to vias.

Plane Inductance

The power and ground planes of a PCB have some amount of inductance associated with them. The geometry of these planes determines their inductance.

Since power and ground planes are by definition a planar structure, current does not just flow through them in one direction. It tends to spread out as it travels from one point to another, in accordance with a property similar to skin effect. For this reason, inductance of planes can be described as "spreading inductance," and is specified in units of henries per square. The square is dimensionless, as it is the shape of a section of plane, not the size, that determines its inductance.

Spreading inductance acts like any other inductance — to resist changes to the amount of current in a conductor. In this case, the conductor is the power plane or planes. This

quantity should be reduced as much as possible, since it retards the ability of capacitors to respond to transient currents in the device. Since the X-Y shape of the plane is typically something the designer has little control over, the only controllable factor is the spreading inductance value. This is primarily determined by the thickness of the dielectric separating a power plane from its associated ground plane.

In high-frequency power distribution systems of the type discussed here, power and ground planes work in pairs. Their inductances do not exist independently of each other. The spacing (and the dielectric constant of the material) between power and ground planes determines the spreading inductance of the pair. The closer the spacing (the thinner the dielectric), the lower the spreading inductance. [Table 4-1](#) gives approximate values of spreading inductance for different thicknesses of FR4 dielectric ([Reference #3](#)).

Table 4-1: Capacitance and Spreading Inductance Values for Various Thicknesses of FR4 Power-Ground Plane Sandwiches

Dielectric Thickness (mil, microns)	Inductance (pH/square)	Capacitance (pF/in ² , pF/cm ²)
4, 102	130	225, 35
2, 51	65	450, 70
1, 25	32	900, 140

Since closer spacing results in decreased spreading inductance, it is best, wherever possible, to place V_{CC} planes directly adjacent to GND planes in the 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 demands it.

Besides offering a low-inductance current path, power-ground sandwiches also offer some high-frequency decoupling capacitance. As plane area increases and as the separation between power and ground planes decreases, the value of this capacitance increases. At the same time, since the parasitic inductance of this capacitance is decreasing, its effective frequency band center frequency increases. Capacitance per square inch is also given in [Table 4-1](#).

This capacitance alone is usually not enough to give power-ground sandwiches a compelling advantage. However, when viewed as a bonus on top of low spreading inductance, it is an advantage most designers gladly take.

PCB Stackup and Layer Order

The placement of V_{CC} and Ground planes in the PCB stackup (determined by layer order) has a significant impact on the parasitic inductances of power current paths. For this reason, PCB designers need to consider layer order in the early stages of the design cycle, putting high-priority supplies in the top half of the stackup and low-priority supplies in the bottom half of the stackup.

Power supplies with high transient current should have their associated V_{CC} planes close to the top surface (FPGA side) of the PCB stackup to decrease the distance in the vertical direction that currents travel through V_{CC} and GND vias before reaching the associated V_{CC} and GND planes. As mentioned in the previous section, every V_{CC} plane should have a GND plane adjacent to it in the stackup to reduce spreading inductance. Since high-frequency currents couple tightly due to skin effect, the GND plane adjacent to a given V_{CC}

plane tends to carry the majority of the current complementary to that in the V_{CC} plane. For this reason, adjacent V_{CC} and GND planes are considered as a pair.

Not all V_{CC} and GND plane pairs can reside in the top half of the PCB stackup, because manufacturing constraints typically require the PCB stackup to be symmetrical about the center with respect to dielectric thicknesses and etched copper areas. The PCB designer must determine which V_{CC} and GND plane pairs have high priority or carry high-frequency energy, and which pairs have low priority or carry lower frequency energy.

Capacitor Effective Frequency

Every capacitor has a narrow frequency band where it is most effective as a decoupling capacitor. Outside this band, it does have some contribution to the PDS but in general it is much smaller. The frequency bands of some capacitors are wider than others. The ESR of the capacitor determines the quality factor (Q) of the capacitor, which determines the width of the effective frequency band. Tantalum capacitors generally have a very wide effective band, while X7R and X5R chip capacitors, with their lower ESR, generally have a very narrow effective band.

The effective frequency band corresponds to the capacitor's resonant frequency. While an ideal capacitor only has a capacitive characteristic, real non-ideal capacitors also have a parasitic inductance ESL and a parasitic resistance ESR. These parasitics act in series to form an RLC circuit (Figure 4-3). The resonant frequency associated with that RLC circuit is the resonant frequency of the capacitor.

To determine the resonant frequency of an RLC circuit, Equation 4-1 is used:

$$F = \frac{1}{2\pi\sqrt{LC}} \quad \text{Equation 4-1}$$

Alternatively, a frequency sweep SPICE simulation of the circuit could be performed, and the frequency where the minimum impedance value occurs would be the resonant frequency.

It is important to distinguish between the capacitor's self-resonant frequency and the effective resonant frequency of the mounted capacitor when it is part of the system. This is simply the difference between taking into account only the capacitor's parasitic inductance, and taking into account its parasitic inductance as well as that of the vias, planes, and connecting traces lying between it and the FPGA. The self-resonant frequency of the capacitor F_{RSELF} (the value reported in a capacitor data sheet), is considerably higher than its effective mounted resonant frequency in the system, F_{RIS} . Since the mounted capacitor's performance is what is important, it is the mounted resonant frequency that is used when evaluating a capacitor as part of a larger PDS.

The main contributors to mounted parasitic inductance are the capacitor's own parasitic inductance, the inductance of PCB lands and connecting traces, the inductance of vias, and power plane inductance. Vias traverse a full board stackup on their way to the device when capacitors are mounted on the underside of the board. These vias contribute something in the range of 300 pH to 1,500 pH on a board with a finished thickness of 60 mils; vias in thicker boards have higher inductance. Because there are two of these paths in series with each capacitor, twice this value should be added to the capacitor's parasitic inductance. This quantity, the parasitic inductance of the capacitor mounting, is designated L_{MOUNT} . To determine the total parasitic inductance of the capacitor in-system, L_{IS} , the capacitor's parasitic inductance L_{SELF} is added to the parasitic inductance of the mounting, L_{MOUNT} :

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

Example

X7R Ceramic Chip capacitor (AVX capacitor data used here)

$$C = 0.01 \mu\text{F}$$

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

$$F_{\text{RSELF}} = 53 \text{ MHz}$$

$$L_{\text{MOUNT}} = 0.8 \text{ nH}$$

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

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

$$L_{\text{IS}} = 1.7 \text{ nH}$$

Plugging in the values from the example:

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

$$F_{\text{RIS}} = \frac{1}{2\pi\sqrt{(1.7 \times 10^{-12} \text{ H}) \cdot (1 \times 10^{-8} \text{ F})}} = 3.8 \times 10^7 \text{ Hz}$$

F_{RIS} : Mounted Capacitor Resonant Frequency: 38 MHz

Since a decoupling capacitor is only effective at a narrow band of frequencies around its resonant frequency, it is important that the resonant frequency be taken into account when choosing a collection of capacitors to build up a decoupling network.

Capacitor Anti-Resonance

One common problem associated with capacitors in an FPGA PDS is anti-resonant spikes in the PDS aggregate impedance. These spikes can be caused by bad combinations of energy storage devices in the PDS (such as discrete capacitors, parasitic inductances, power and ground planes). If the inter-plane capacitance of the power and ground planes has an especially low Z with a high-quality factor, the crossover point between the high-frequency discrete capacitors and this plane capacitance might exhibit a high-impedance anti-resonance peak. If the FPGA has high transient current demand at this frequency (acting as a stimulus), a large noise voltage results. The PDS can be improved only by bringing down the impedance of the anti-resonant spike. To mitigate this problem, either the characteristics of the high-frequency discrete capacitors or the characteristics of the V_{CC} and Ground planes must be changed.

Capacitor Placement

Capacitors need to be close to the device to perform the decoupling function. There are two basic reasons for this requirement.

First, increased spacing between the device and decoupling capacitor increases the distance travelled by the current in the power and ground planes, and hence, the inductance of the current path between the device and the capacitor. Since the inductance of this path (the loop followed by current as it goes 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, decreasing its inductance is a matter of decreasing the loop area. Shortening the distance between the device and the decoupling capacitor(s) reduces the inductance resulting in a less impeded transient current flow.

Because of the dimensions of PCBs, this reason tends to be less important with regard to placement than the second reason.

The second reason deals with the phase relationship between the FPGA noise source and the mounted capacitor. Their phase relationship determines the capacitor's effectiveness. For a capacitor to be effective in providing transient current at a certain frequency (for instance, the optimum frequency for that capacitor), it must be within a fraction of the wavelength associated with that frequency. The placement of the capacitor 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 relevant factor.

Noise from the FPGA falls into certain frequency bands, and different sizes of decoupling capacitors take care of different frequency bands. For this reason, capacitor placement is determined based on the effective frequency of each capacitor.

When the FPGA initiates a change in its current demand, it causes a small local disturbance in the voltage of the PDS (a point in the power and ground planes). For a decoupling capacitor to counteract this, the capacitor has to first see a voltage difference. There is a finite time delay between the start of the disturbance at the FPGA power pins and the start of the capacitor's view of the disturbance. This time delay is equal to the distance from FPGA power pins to capacitor, divided by the propagation speed of current through FR4 dielectric (the substrate of the PCB where the power planes are embedded). There is another delay of the same duration for the compensation current from the capacitor to reach the FPGA.

Therefore, for any transient current demand in the FPGA, there is a round-trip delay to the capacitor before any relief is seen at the FPGA. For placement distances greater than one quarter of a wavelength of some frequency, the energy transferred to the FPGA is negligible.

For decreasing distances less than a quarter wavelength, the energy transferred to the FPGA increases to 100% at zero distance. Efficient energy transfer from the capacitor to the FPGA requires placement of the capacitor at a fraction of a quarter wavelength of the FPGA power pins. This fraction should be small because the capacitor is also effective at frequencies slightly above its resonant frequency, where the corresponding wavelength is shorter.

In practical applications, one tenth of a quarter wavelength is a good target. This leads to placing a capacitor within one fortieth of a wavelength of the power pins it is decoupling. The wavelength corresponds to F_{RIS} , the capacitor's mounted resonant frequency.

Example

0.001 μ F X7R Ceramic Chip capacitor, 0402 package

$L_{IS} = 1.6$ nH

$$F_{RIS} = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{1.6 \times 10^{-9} \times 0.001 \times 10^{-6}}} = 125.8 \text{ MHz}$$

Equation 4-2 calculates T_{RIS} , the mounted period of resonance, from F_{RIS} .

$$T_{RIS} = \frac{1}{F_{RIS}} = \frac{1}{125.8 \times 10^6} = 7.95 \text{ ns} \quad \text{Equation 4-2}$$

Equation 4-3 computes the wavelength based on T_{RIS} and the propagation velocity in FR4 dielectric.

$$\lambda = \text{Wavelength} = \frac{T_{RIS}}{V_{PROP}} \quad \text{Equation 4-3}$$

$$\text{where } V_{PROP} = 166 \times 10^{-12} \frac{s}{inch}$$

$$\lambda = \frac{T_{RIS}}{V_{PROP}} = \frac{7.95 \times 10^{-9}}{166 \times 10^{-12}} = 47.9 \text{ inches}$$

$$R_{PLACE} = \frac{\lambda}{40} \quad \text{Equation 4-4}$$

$$R_{PLACE} = \frac{\lambda}{40} = \frac{47.9 \text{ inches}}{40} = 1.20 \text{ inches}$$

In this example, the effective frequency, equal to the resonant frequency, can be determined by Equation 1. This effective frequency is determined to be 125.8 MHz. The reciprocal of this is taken to give the resonant period, 7.95 ns using Equation 2. Using the propagation speed of current in FR4 (approximately 166 ps per inch), the wavelength associated with this capacitor is computed to be approximately 48 inches using Equation 3. As computed in Equation 4, one fortieth of this is 1.2 inches. Therefore the target placement radius (R_{PLACE}) for capacitors of this size is within 1.2 inches (3.0 cm) of the power and ground pins they are decoupling.

All other capacitor sizes follow in the same manner. A radius of 1.2 inches is not terribly difficult to achieve in current PCB technology. It does not require placing capacitors directly underneath the device on the opposite side of the PCB. It is acceptable for capacitors to be mounted around the periphery of the device, provided the target radius is maintained. The 0.001 μF capacitors are among the smallest in the decoupling network, so placement radii less than an inch are unnecessary. For larger capacitors, the target placement radius expands quickly as the resonant frequency goes down. A 4.7 μF capacitor, for example, can be placed anywhere on the board, as its target radius of 98 inches is much bigger than most PCBs (corresponding to the resonant frequency of 1.56 MHz).

Example Capacitor Layout

Figure 4-7 is an example of the bottom-side PCB artwork showing the capacitor layout. Black fill and hatch represents plated copper, red represents vias, blue represents silkscreen labels, and purple represents package outlines. The FPGA footprint can be seen as the regular array of red via dots in the upper portion of the figure at the center. The absence of vias in a cross pattern at the center of the device indicates that solder lands on the top surface had their associated vias escape out toward the corners.

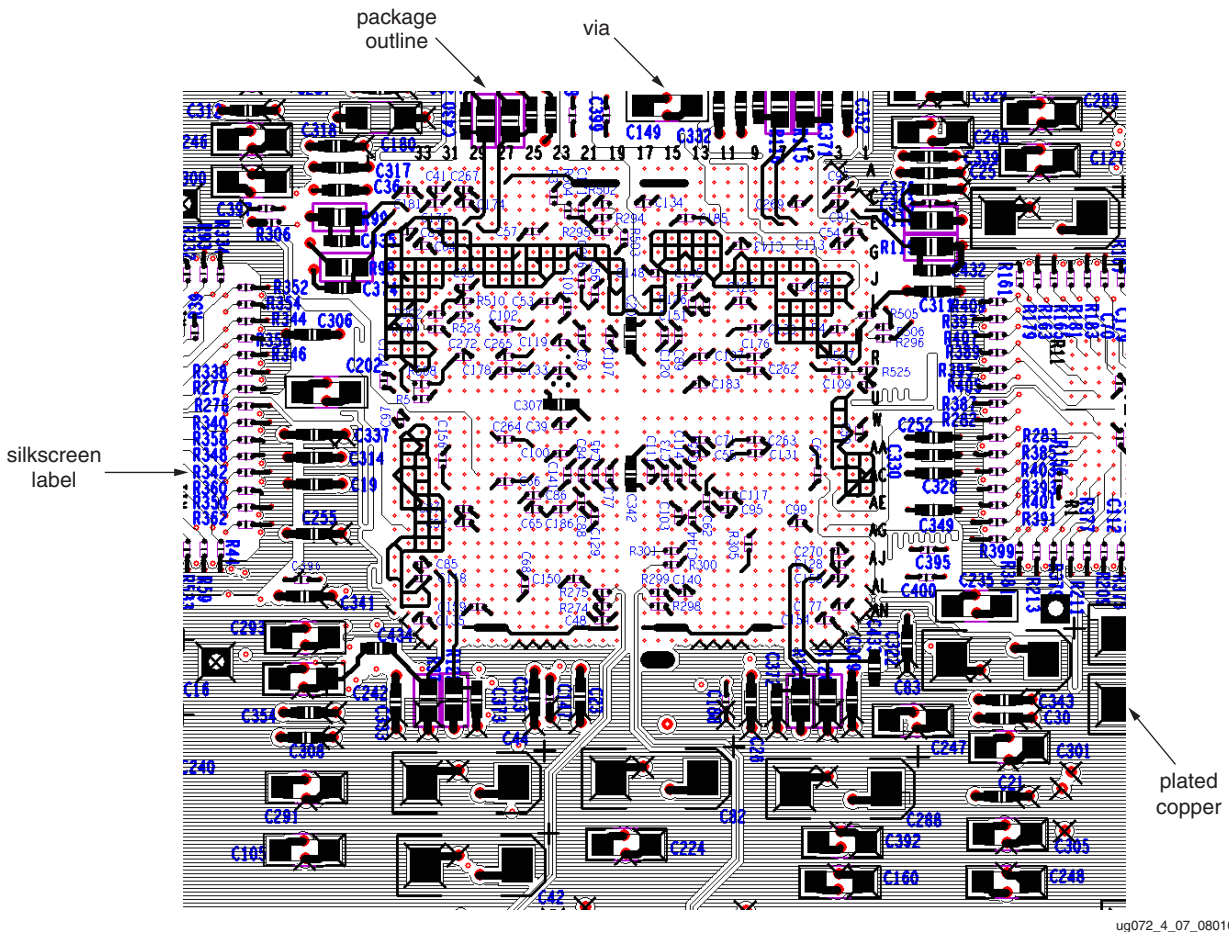


Figure 4-7: Example PCB Layout Showing Capacitor Placement on the Bottom Surface

In this example, many of the high-frequency 0402 decoupling capacitors are placed within the footprint of the FPGA on the opposite side of the board (C150, C117). There are also a handful of 0603 decoupling capacitors and termination resistors (C307, R274). Larger capacitors are placed outside the footprint of the FPGA, moving farther away from the FPGA with increasing size (C247, C288).

Traces connecting capacitor lands to vias are kept as short as possible. Also for large-package capacitors with large separation between solder lands (C42, C224), vias are inserted in between the solder lands to reduce parasitic inductance of the mounting.

It is not necessary to place high-frequency capacitors within the footprint of the FPGA. It is perfectly acceptable to place all capacitors around the periphery of the device, provided all V_{CC} planes have a Ground plane adjacent to them, separated by a dielectric less than 4 mils in thickness. Also, in cases where V_{CC} and Ground plane pairs are in the top half of the stackup (closer to the device), it is advantageous to place capacitors on the top surface of the board, around the periphery of the device.

In cases where large numbers of external termination resistors are used, placement of the termination resistors takes priority over the decoupling capacitors. Moving away from the device in concentric rings, termination resistors should be closest to the device, followed by the smallest-value decoupling capacitors, then followed by larger-value decoupling capacitors.

PDS Design and Verification

Having discussed the basic operating principles of power distribution systems, this section introduces a step-by-step process for designing and verifying a PDS.

Step 1: Determining Critical Parameters of the FPGA

In designing the first iteration of the decoupling capacitor network, the basic objective is to have one capacitor per V_{CC} pin used on the device. Therefore, the *effective* number of V_{CC} pins for each supply must be determined.

Very few designs use 100% of the various resources in the FPGA. The FPGA package and the PDS inside it are very carefully sized to meet the needs of a fully utilized die without being overly conservative. The number of V_{CC} and GND pins on a package for a given device is determined based on the needs of a 100% utilized FPGA. The determining factor is not DC power handling abilities — it is transient current impedance. Decoupling capacitor requirements track very closely since they are based on the same factor. For this reason, the number of V_{CC} pins on each supply is used as an indicator of the number of capacitors needed on that supply. All supplies must be considered: V_{CCINT} , V_{CCAUX} , V_{CCO} , and V_{REF} .

It is only necessary to provide one capacitor per V_{CC} pin if all pins are used. There is no need to decouple V_{REF} pins if they are not used as V_{REF} . Conversely, V_{CCAUX} and V_{CCINT} pins must always be fully decoupled, i.e., they must always have one capacitor per pin. V_{CCO} can be pro-rated according to I/O utilization.

Step 2: Designing the Generic Bypassing Network

A number of Xilinx test boards and customer designs were analyzed to discern some trends of successful PDS designs. In 80% to 100% utilized designs with power supply noise on the order of half the maximum allowed power supply noise ($V_{RIPPLE}/2$), the PDS generally has approximately one capacitor per V_{CC} pin on a per-supply basis. The generic bypassing network is designed with this range of capacitors in mind. The pro-rated number of V_{CCO} pins is used.

Given the number of discrete capacitors needed as determined above, a distribution of capacitor values adding up to that total number must be determined. To cover a broad range of frequencies, a broad range of capacitor values must be used. The proportion of high-frequency capacitors to low-frequency capacitors is an important factor.

The objective of a parallel combination of a number of values of capacitors is to keep a low and flat power supply impedance over frequencies from the 500 kHz range to the 500 MHz range. Both large value (low frequency) and small value (high frequency) capacitors are needed. Small value capacitors tend to have less of an impact on the total impedance profile, so a greater number of small value capacitors are needed to yield the same impedance level impact as a small number of large value capacitors.

To keep the impedance profile smooth and free of anti-resonance spikes, a capacitor is generally needed at least in every decade of the capacitor value range. The typical ceramic capacitor range generally spans values from 0.001 μF to 4.7 μF . The exact value of these capacitors is not critical. What is critical is having some capacitor value in every order of magnitude over this range. More values are better, as a flatter impedance profile is yielded.

A ratio of capacitors giving a relatively flat impedance is one where the quantity of capacitors is roughly doubled for every decade of decrease in size. In other words, if the

bottom three values in the network were 1.0 μF , 0.1 μF and 0.01 μF , the network might have two 1.0 μF capacitor, four 0.1 μF capacitors, and eight 0.01 μF capacitors.

In addition, low-frequency capacitance in the form of tantalum, OS-CON, or electrolytic capacitors is needed. These large capacitors typically have a higher ESR than ceramic chip capacitors, making them effective over a wider range of frequencies. This also makes the capacitors less likely to contribute to anti-resonance spikes. For this reason, it is not necessary to maintain the rule of one value per decade. Generally, one value in the 470 μF to 1000 μF range is sufficient.

A set of percentages is helpful for calculating these ratios based on the total number of capacitors. Table 4-2 and Table 4-3 give these percentages for the two classes of supplies: V_{CCINT} , V_{CCAUX} , V_{CCO} , and V_{REF} .

Table 4-2: Capacitor Value Ratios for a Balanced Decoupling Network (V_{CCINT} , V_{CCAUX} , V_{CCO})

Capacitor Value	Quantity Percentage	Capacitor Type
470 μF to 1000 μF	4%	Tantalum
1.0 to 4.7 μF	14%	X7R 0805
0.1 to 0.47 μF	27%	X7R 0603
0.01 to 0.047 μF	55%	X7R 0402

Table 4-3: Capacitor Value Ratios for a Balanced Decoupling Network (V_{REF})

Capacitor Value	Quantity Percentage	Capacitor Type
0.1 μF to 0.47 μF	50%	X7R or X5R 0603
0.01 μF to 0.047 μF	50%	X7R or X5R 0402

For every power supply except V_{REF} these ratios should be roughly maintained. For V_{REF} supplies, the values should be distributed in a 50/50 ratio of 0.1 μF to 0.47 μF capacitors and 0.01 μF to 0.047 μF capacitors. Since the primary function of V_{REF} decoupling capacitors is to reduce the impedance of V_{REF} nodes thus reducing crosstalk coupling, very little low-frequency energy is needed. Therefore, only capacitors in the 0.01 μF - 0.47 μF range are necessary.

Step 3: Simulation

During simulation, the generic decoupling network is verified and in some cases refined. The designer can experiment with different values of capacitors or different packages to achieve an optimum power supply impedance profile for the constraints of the system. A number of levels of PDS design tools available from various EDA vendors are listed in [“EDA Tools for PDS Design and Simulation,” page 54](#).

The simulation circuit is essentially a parallel combination of the decoupling capacitors with associated parasitics. The simulator calculates the aggregate impedance over the pertinent range of frequencies. The equivalent circuit can be created and analyzed in SPICE (see [“SPICE Simulation Examples,” page 52](#) for an example SPICE deck) or in one of the tools listed in [“EDA Tools for PDS Design and Simulation.”](#) A more limited but still effective approach is to plot the impedance profile in a spreadsheet tool (for example, Microsoft Excel).

Note that a lumped RLC simulation of this type does not reflect the distributed RLC properties of the V_{CC} and Ground planes of the PCB stackup. The effects of these planar structures usually begin to manifest in the 500 MHz range, and are dependent on the geometries of the planes (for example, length and width). These are difficult to predict without the use of a distributed model, such as what is offered by a tool like Speed2000, SIwave, Spectraquest Power Integrity, or a full-mesh RLC SPICE simulation. For this reason, it is unwise to draw any conclusions from the results of a lumped RLC simulation above 500 MHz.

In using any of these tools to simulate a bypassing network, it is important to have accurate parasitic values. Obtaining accurate self-parasitic data from the capacitor vendor or from in-house testing is important. The mounting parasitics lying in the path between the bypass capacitor and the FPGA also need to be taken into account. These parasitics combined in series give the mounted capacitor parasitic resistance and inductance. The section on [“Mounting Inductance,” page 33](#) covers the details of mounting modeling. [“Calculation of Via Inductance,” page 52](#) lists equations for via parasitic inductance. A more accurate inductance number for a particular geometry can be obtained using a field solver such as Ansoft’s HFSS. For the following simulation, a value of 0.8 nH to 0.9 nH in mounting inductance was added to each capacitor’s parasitic self-inductance to come up with L_{IS} . This parameter reflects the inductance of small capacitor mountings in a board on the order of 60 mils thick. Thicker board stackups have a higher associated via inductance.

[Figure 4-8](#) shows a simple impedance plot from a simulation of the parallel combination of these capacitors, taking into account their parasitics and the approximate parasitics of the PCB. An equivalent SPICE netlist is included in [“SPICE Simulation Examples,” page 52](#). [Table 4-4](#) lists the capacitor quantities, values, and parasitic values used in the simulation. The RLC characteristics of the V_{CC} and GND planes of the PCB are not taken into account.




Figure 4-8: PDS Impedance Versus Frequency Plot

Table 4-4: Values Used in Impedance Plot of Figure 4-8

Quantity	Symbol	Package	Capacitive Values (μF)	Parasitic Inductance (nH)	Parasitic Resistance (ohms)
2	❖	E	680	2.8	0.57
7	◐	0805	2.2	2.0	0.02
13	◆	0603	0.22	1.8	0.06
26	■	0402	0.022	1.5	0.20

This collection of capacitors is a good start. The impedance is below 0.033Ω from 500 KHz to 150 MHz, and increases to 0.11Ω at 500 MHz. Over this range there are no significant anti-resonance spikes. These capacitors are used in the board design.

Step 4: Building the Design

At this stage, the PCB is laid out with the final capacitor networks verified in simulation. The board is built. See earlier sections on capacitor placement and land geometries for detailed layout information.

Step 5: Measuring Performance

In the performance measurement step, measurements are made to determine whether the PDS is adequate for the devices it is serving. Determining whether or not a bypassing network is adequate for a given design is relatively simple. The measurement is performed

with a high-bandwidth oscilloscope (1 GHz oscilloscope and 1 GHz probe at minimum), on a design running realistic test patterns.

Noise Magnitude Measurement

The measurement is taken either directly at the power pins of the device, or across a pair of unused I/O, one driven High and one driven Low. V_{CCINT} and V_{CCAUX} can only be measured at the PCB backside vias. V_{CCO} can also be measured this way, but more accurate results are obtained measuring fixed signals at unused I/Os in the same bank.


When making the noise measurement at the back-side of the board, it is necessary to take into account the parasitics of the vias in the path between the measuring point and FPGA, as any voltage drop occurring in this path is not accounted for in the oscilloscope measurement. Backside via measurements also have a potential pitfall. Many times, decoupling capacitors are mounted directly underneath the device, meaning that the capacitor lands may connect to these V_{CC} and GND vias directly with surface traces. These capacitors can confound the measurement, as they act like a short circuit for high-frequency AC current. To make sure such capacitors do not short the measurement, capacitors at the measurement site must be removed.

When measuring V_{CCO} noise, the measurement can be taken at a pair of I/O pins configured as strong drivers to logic 1 and logic 0. This technique, when performed correctly, can also show die-level noise.

Measuring the driven logic 1 against the driven logic 0 shows the degree of rail collapse at the die. Measuring a driven logic 0 against PCB ground shows the amount of ground bounce the die is experiencing relative to the PCB PDS. Since all grounds are common at the die and package levels of the device (excepting AGND of MGTs, V_{REFN} and AV_{SS} of the system monitor, and ADCs), a ground bounce measurement taken at an unused I/O pin shows the ground bounce of all supplies. Rail collapse measurements, on the other hand, only apply to V_{CCO} .


To make these measurements, the oscilloscope should be in infinite persistence mode, to acquire noise over a long time period (many seconds or minutes). If the design operates in a number of different modes, utilizing different resources in different amounts, these various conditions and modes should be in operation while the oscilloscope is acquiring the noise measurement. Noise measurements should be made at a few different V_{CC}/GND pairs on the FPGA to eliminate the effects of a local noise phenomena.

Figure 4-9 shows an instantaneous noise measurement taken at the V_{CCINT} pins of a sample design. Figure 4-10 shows an infinite persistence noise measurement of the same design. Since the infinite persistence measurement catches ALL noise events over a long period, it obviously yields more relevant results.



ug072_4_09_080104

Figure 4-9: Instantaneous Measurement of V_{CCO} Supply, with Multiple I/O Sending Patterns at 100 MHz



ug072_4_10_080104

Figure 4-10: Infinite Persistence Measurement of Same Supply

This measurement represents the peak-to-peak noise. If it is greater than or equal to the maximum V_{CC} ripple voltage specified in the data sheet (10% of V_{CC}), then the bypassing network is not adequate. The maximum voltage ripple allowed for this particular supply, with a nominal value of 1.5V DC, is 10% of this, or 150mV. The scope shots show noise in the range of 60 mV. From this measurement, it is clear that the decoupling network is adequate.

If, however, the measurement showed noise greater than 10% of V_{CC} , the PDS would be inadequate. To have a working, robust design, changes should be made to the PDS. A greater number of capacitors, different capacitance values, or different numbers of the various decoupling capacitor values will bring the noise down.

Having the necessary information to improve the decoupling network requires additional measurements. Specifically, measurement of the noise power spectrum is necessary to determine the frequencies where the noise resides. There are many ways to do this. A spectrum analyzer works well as does an oscilloscope with FFT math functionality. Alternatively, a long sequence of time-domain data can be captured from an oscilloscope and converted to frequency domain using MATLAB or other software supporting FFT. It is also possible to get a basic feel for the frequency content of the noise by simply looking at the time-domain waveform and measuring the individual periodicities present in the noise.


Noise Spectrum Measurements

A spectrum analyzer is a frequency-domain instrument. It shows the frequency content of a voltage signal at its inputs. When used to measure an inadequate PDS, the user can see the exact frequencies where the PDS is inadequate. Excessive noise at a certain frequency indicates a frequency where the PDS impedance is too high for the transient current demands of the device. Armed with this information, the designer can modify the PDS to accommodate the transient current at the specific frequency. This is accomplished either by adding capacitors with resonant frequencies close to the frequency of the noise or by lowering the PDS impedance at the critical frequency through other means.

The noise spectrum measurement should be taken in the same place as the peak-to-peak noise measurement — directly underneath the device, or at a pair of unused I/O driven High and Low. A spectrum analyzer takes its measurements through a 50 Ω cable, rather than through an active probe like the oscilloscope. One of the best ways to attach the cable for measurements is through an SMA connector tapped into the power and ground planes in the vicinity of the device. In most cases this is not available. Another way to attach the cable for measurement of noise in the power planes is to remove a decoupling capacitor in the vicinity of the device, and solder the center conductor and shield of the cable directly to the capacitor lands. Alternatively, a probe station can be used.

In most cases, distinct bands of noise at fixed frequencies are seen. These correspond to the clock frequency and its harmonics. The height of each band represents its relative power. The majority of the energy is usually contained in tight bands around 3 or 4 of the harmonics, with power falling off as frequency increases.

Figure 4-11 shows an example of a noise spectrum measurement. It is a screenshot of a spectrum analyzer measurement of power supply noise on V_{CC0} , with multiple I/O sending patterns at 150 MHz.



ug072_4_11_080104

Figure 4-11: Screenshot of Spectrum Analyzer Measurement of V_{CC0}

The noise bands correspond to frequencies where the FPGA has a demand for current but is not receiving it from the PDS. This could be because there is not enough capacitance, or because there is enough capacitance but the parasitic inductance of the path separating the capacitors from the FPGA is too great. Whatever the cause, the impedance of the power supply at this frequency is too high. Conversely, at frequencies where there is very little or no noise, the impedance may be lower than it needs to be. To solve these problems, the bypassing network must be modified. New capacitor values, or different quantities of the original values should be chosen.

Step 6: Optimum Bypassing Network Design (Optional)

In cases where a highly optimized PDS is needed, further measurements can be taken to guide the design of a carefully tailored decoupling network. A network analyzer can be used to measure the impedance profile of a prototype PDS, giving an output similar to what was discussed in the simulation section. The network analyzer sweeps a stimulus across a range of frequencies and measures the impedance of the PDS at each frequency. Its output is impedance as a function of frequency.

Since the spectrum analyzer gives an output of voltage as a function of frequency, these two measurements can be used together to determine transient current as a function of frequency.

$$I(f) = \frac{V(f) \text{ From Spectrum Analyzer}}{Z(f) \text{ From Network Analyzer}}$$

Armed with an understanding of the design's transient current requirements, the designer can make better PDS choices. With maximum voltage ripple value from the data sheet, the value of impedance needed at all frequencies can be determined. This yields a target impedance as a function of frequency. Given this, a network of capacitors can be designed to accommodate the transient current of the specific design.

This six-step process lays out a closed-loop method for designing and verifying a power distribution system. Its use ensures an adequate PDS for any design.

Other Concerns and Causes

If this step-by-step method does not yield a design meeting the required noise specifications, then other aspects of the system should be analyzed for possible changes.

Possibility 1: Excessive Noise from Other Devices on Board

When ground and/or power planes are shared among many devices, as is often the case, noise from an inadequately decoupled device can affect the PDS at other devices. RAM interfaces with inherently high transient current demands due to temporary periodic contention and high-current drivers are a common cause; large microprocessors are another. If unacceptable amounts of noise are measured locally at these devices, an analysis should be done on the local PDS and decoupling networks of the component.

Possibility 2: Parasitic Inductance of Planes, Vias, or Connecting Traces

In this case there is enough capacitance in the bypassing network, but too much inductance in the path from the capacitors to the FPGA. This could be due to a bad choice of connecting trace or solder land geometry, too long a path from capacitors to the FPGA, and/or a current path in power vias that traverse an exceptionally thick stackup.

In the case of inadequate connecting trace and capacitor land geometry, it is important to keep in mind the loop inductance of the current path. If the vias for a bypass capacitor are spaced a few millimeters from the capacitor solder lands on the board, the current loop area is greater than it needs to be (Figure 4-6a). Vias should be placed directly against capacitor solder lands (Figure 4-6b). Never connect vias to the lands with a section of trace (Figure 4-6a). Other improvements of geometry are via-in-pad (where the via is actually under the solder land), not shown, and via beside pad (where vias are not at the ends of the lands, but rather astride them), see Figure 4-6c. Double vias are a further improvement (Figure 4-6d).

If the inductance of the path in the planes is too great, there are two parameters that can be changed; the length of the electrical path, and the spreading inductance of the planes themselves.

The path length is determined by capacitor placement. Capacitors must be placed close to the power/ground pin pairs on the device being bypassed. This is especially important for the smallest capacitors in the network, since care has been taken to choose capacitors with low parasitic inductance. There is no use in connecting a low-inductance, high-frequency capacitor to a device through a high-inductance path. Larger capacitors inherently have a high parasitic self inductance allowing the proximity to the device to be less important.

The spreading inductance of the planes is controlled by the plane spacing and by the dielectric constant of the material between them. See section on “[Plane Inductance](#),” page 35.

When boards are exceptionally thick (greater than 90 mils or 2.3 mm), vias have higher parasitic inductance. In these cases, the following changes to the design should be considered. The first is to move the V_{CC}/GND plane sandwiches close to the top surface the FPGA is on. The second is to place the highest frequency capacitors on the top surface. Both changes 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 making refinements to the PDS, the I/O interface power can be scaled back. This goes for both outputs from the FPGA and inputs to the FPGA. In some cases, excessive overshoot on inputs to the FPGA can reverse-bias the clamp diodes in the IOBs. This can put large amounts of noise into V_{CCO} . If this condition is occurring, the drive strength of these interfaces should be decreased, or termination should be used (both on input and output paths).

Possibility 4: I/O Signal Return Current Travelling in Sub-Optimal Paths

Excessive noise in the PDS can be caused by I/O signal return currents. 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 back into the device's power/ground system. If there is no low-impedance return current path available, a less optimal, higher impedance path is used. When this occurs, voltage changes are induced in the PDS.

This situation can be improved by ensuring that every signal has a closely spaced and fully intact return path. Various strategies could be required including restricting signals to only a few of the available routing layers, and providing low-impedance paths for AC currents to travel between reference planes (decoupling capacitors at specific locations on the PCB).

Calculation of Via Inductance

Via inductance is a major contributor to the parasitic inductance of a capacitor mounting. The dimensions of a via largely determine its parasitic inductance. Equation 4-5, from Grover, is used to determine the self-inductance of a single filled via based on its length and diameter. Dimensions are in inches and nanohenries.

$$L = 5.08 \times h \times \left[\ln\left(\frac{4 \times h}{d}\right) - 0.75 \right] \quad \text{Equation 4-5}$$

Example

To calculate the inductance of a via going from the bottom surface of the board to the top surface of the board, use the board finished thickness for via length: the board finished thickness at 62 mils, the via diameter at 3 mils. There are 1000 mils in an inch.

$$h = 0.062 \text{ in}$$

$$d = 0.003 \text{ in}$$

$$L = 5.08 \times h \times \left[\ln\left(\frac{4 \times h}{d}\right) - 0.75 \right]$$

$$L = 5.08 \times 0.062 \times \left[\ln\left(\frac{4 \times 0.062}{0.003}\right) - 0.75 \right]$$

$$L = 5.08 \times 0.062 \times 3.67$$

$$L = 1.15 \text{ nH}$$

This result is the self-inductance of a single via. The self-inductance is only one part of the total inductance of the current loop the via is a part of. Since the mutual inductance of vias with opposing currents (power and ground) has its own effect on the total inductance, it should be taken into account when greater accuracy is desired. The mutual inductance of closely spaced complementary vias lowers the total inductance by a small amount.

SPICE Simulation Examples

This appendix demonstrates the method used to simulate decoupling capacitor networks in SPICE. HSPICE techniques are discussed here. Other variants of SPICE or dedicated PDS simulation software can also be used. The simulation referenced below is purely for illustrative purposes. Simulator details are beyond the scope of this discussion and are left to the readers' investigation. The HSPICE result is included in Figure 4-12. A schematic representation is included in Figure 4-13.

These capacitor networks represent the capacitance and parasitics of an 18-capacitor network. The general capacitor array impedance calculation follows these steps:

1. Formulate a netlist for the L-C-R network.
2. Understand where the input node and output node are located.
3. Apply an AC stimulus to the input port.
4. Run an AC analysis on the L-C-R network.
5. Measure the input current as well as the input AC voltage.
6. Formulate $Z = V/I$.
7. Plot the result using a log scale for ease of viewing.

In this approach, the AC stimulus is set to 1A. The AC Analysis directive sweeps an AC current waveform across a prescribed set of frequency points. The number of frequency points per decade is commented in the appended HSPICE netlist. With the AC current magnitude set to 1A, the impedance is calculated based on $Z = V/I$. Thus, V is the main calculated variable — the voltage at the capacitor array positive node.

Two other details complete the SPICE decks:

1. There is a DC bias resistor to ground.
2. There is a small input resistor connecting the AC source to the L-C-R network (this is optional).

Item 1 is necessary to decrease simulation time. It allows SPICE to quickly calculate an operating point for the circuit prior to AC analysis. This is accomplished by providing SPICE a DC path to the L-C-R network (to ground by way of the bias resistor). Item 2 is optional, but convenient. It provides a component to monitor the input current to the L-C-R network.

For viewing the simulated impedance result in HSPICE, the `.net` directive is executed in order that HSPICE calculates Z_{IN} for direct plotting.

HSPICE Netlist

The HSPICE netlist is available on the Xilinx website:
<http://www.xilinx.com/bvdocs/appnotes/xapp623.zip>

HSPICE Output

Figure 4-12 shows the HSPICE output: $Z_{IN}(MAG)$ using the AWAVES graphical viewer.




Figure 4-12: HSPICE Output

Schematic Circuit

Figure 4-13 shows a capacitor array with corresponding parasitic inductance and resistance.




Figure 4-13: Schematic Circuit

EDA Tools for PDS Design and Simulation

Table 4-5 lists the some vendors of EDA tools for PDS design and simulation.

Table 4-5: EDA Tools for PDS Design and Simulation

Tool	Vendor	Website URL
SIwave	Ansoft	http://www.ansoft.com
Spectraquest Power Integrity	Cadence	http://www.cadence.com
Speed 2000	Sigrity	http://www.sigrity.com
Star HSPICE	Synopsys	http://www.synopsys.com
UCADES3.exe	UltraCAD	http://www.ultracad.com

References

1. [UG076](#), *Virtex-4 RocketIO Multi-Gigabit Transceiver User Guide*.
2. [UG070](#), *Virtex-4 FPGA User Guide*.
3. Larry D. Smith, *Decoupling Capacitor Calculations For CMOS Circuits*, Proceedings EPEP Conference, November 1984.
4. Frederick W. Grover Ph.D., *Inductance Calculations: Working Formulas and Tables*, D. Van Nostrand Company, Inc., 250 Fourth Avenue, New York, NY, 1946.

