Zynq-7000 All Programmable SoC PCB Design and Pin Planning Guide

UG933 (v1.2.1) October 1[illegible] 2012

XILINX®

Revision History

The following table show [illegible] revision hi [illegible] for this do [illegible]

Date	Ver [illegible]	Re [illegible]
[illegible]	[illegible]	Initial Xilinx re [illegible]
06/06/12	1.1	Corrected format issue.
08/29/12	1.2	Updated Table 3-1 and Table 3-2 for additional devices/packages. Added 680 µF capacitor specification to Table 3-3.
10/11/12	1.2.1	Corrected document number (changed UG993 to UG933).

Table of Contents

Chapter 1: Intro ction

Chapter 2: PCB Technology Basics

Chapter 3: Po Distribut em

Chapter 4: SelectIO Signaling

Chapter 5: Processing System (PS) Power and Signaling

Appendix A: Additional Resources

Chapter 1

Introduction

About This Guide

This guide provides information on PCB design and pin planning [illegible]-7000™ All Programmable SoC (AP SoC), with a focus on strategies for making decisions at the PCB and interface level.

This Zynq-7000 All Programmable SoC PCB Design [illegible] Planning Guide, part of an overall set of documentation on the Zynq-7000 A[illegible] available on the Xilinx web[illegible] www.xilinx.com/zynq.

Additional Support Resources

[illegible] mation reg[illegible] materials, trac[illegible] niques for high speed signals, refer to chapters four and five of UG483, *7-Series FPGAs PCB and Pin Planning Guide*. A comprehensive list of all additional resources is provided in Appendix A, Additional Resources.

PCB Technology Basics

Introducti[illegible]

Printed circuit boards (PCBs) are electrical systems, with electri[illegible] complicated as the discrete components and devices mounted to them[illegible] PCB [illegible]s complete control over many aspects of the PCB; however, curre[illegible]nology place[illegible]straints a[illegible] limits on the geometries and resulting electrical pr[illegible]. The following informatio[illegible] provided as a guide to the freedoms, limitations, a[illegible]niques for PCB designs usi[illegible] Zynq-7000 AP SoC devices.

This chapter contains the following sections:

- PCB Structures
- Transmissio[illegible]s
- Retur[illegible]

PCB Structures

PCB technology has not changed significantly in the last few decades. An insulator substrate material (usually FR4, an epoxy/glass composite) with copper plating on both sides has portions of copper etched away to form conductive paths. Layers of plated and etched substrates are glued together in a stack with additional insulator substrates between the etched substrates. Holes are drilled through the stack. Conductive plating is applied to these holes, selectively forming conductive connections between the etched copper of different layers.

While there are advancements in PCB technology, such as material properties, the number of stacked layers used, geometries, and drilling techniques (allowing holes that penetrate only a portion of the stackup), the basic structures of PCBs have not changed. The structures formed through the PCB technology are abstracted to a set of physical/electrical structures: traces, planes (or planelets), vias, and pads.

Traces

A trace is a physical strip of metal (usually copper) making an electrical connection between two or more points on an X-Y coordinate of a PCB. The trace carries signals between these points.

[illegible]

A plane is an u[illegible]rupted area of metal covering the entire PCB layer. A planelet, a variation of a pla[illegible] an uninterrupted area of metal covering only a portion of a PCB layer. Typically, a numbe[illegible] planelets exist in one PCB layer. Planes and planelets distribute power to a number of points on a PCB. They are very important in the transmission of signals along traces because they are the return current transmission medium

Vias

A via is a piece of metal making an electrical conne[illegible]tween two or more points [illegible] Z space of a PCB. Vias carry signals or power bet[illegible]ers of a PCB. In current plated-through-hole (PTH) technology, a via is fo[illegible]lating the inner surface o[illegible] drilled through the PCB. In current microvia tech[illegible]so known as High Densit[illegible] Interconnect or HDI[illegible]via is formed with a laser [illegible] the substrate material [illegible] deforming the co[illegible]tive plating. These microvia[illegible]etrate more than one [illegible] layers, however[illegible]can be st[illegible]tair-stepp[illegible] traversing [illegible] full b[illegible] thickness.

Pads and Antipads

Because PTH vias are conductive over the whole length of the via, a method is needed to selectively make electrical connections to traces, planes, and planelets of the various layers of a PCB. This is the function of pads and antipads. Pads are small areas of copper in prescribed shapes. Antipads are small areas in prescribed shapes where copper is removed. Pads are used both with vias and as exposed outer-layer copper for mounting of surface-mount components. Antipads are used mainly with vias. For traces, pads are used to make the electrical connection between the via and the trace or plane shape on a given layer. For a via to make a solid connection to a trace on a PCB layer, a pad must be present for mechanical stability. The size of the pad must meet drill tolerance/registration restrictions. Antipads are used in planes. Because plane and planelet copper is otherwise uninterrupted, any via traveling through the copper makes an electrical connection to it. Where vias are not intended to make an electrical connection to the planes or planelets passed through, an antipad removes copper in the area of the layer where the via penetrates.

Lands

For the purposes of soldering surface mount components, pads on outer layers are typically referred to as lands or solder lands. Making electrical connections to these lands usually requires vias. Due to manufacturing constraints of PTH technology, it is rarely possible to place a via inside the area of the land. Instead, this technology uses a short section of trace [illegible] surface pad. The minimum length of the connecting trace is determined by minimu[illegible] specifications from the PCB manufacturer. Microvia technology is not constrained, a[illegible] can be placed directly in the area of a solder land.

Dimensions

The major factors defining the dimensions of the PCB are PCB man[illegible]ing limits, AP SoC package geometries, and system compliance. Other factor[illegible] Manufacturing (DFM) and reliability impose further lim[illegible]t because [illegible]re application specific, they are not documented in thi[illegible]guide.

The dimensions of the AP SoC package, in combi[illegible]th PCB manufacturing lim[illegible] define most of the geometric aspects of the PCB[illegible]s described in this sectio[illegible] directly and indirectly. This significantly constrai[illegible] designer. The package ba[illegible] (1.0 mm for FF packa[illegible]s) defines the land pad lay[illegible]inimum surface feature [illegible] current PCB techno[illegible]y define the via arrangeme[illegible]a under the device. Mi[illegible] via diameters an[illegible]p-out ar[illegible]d those vi[illegible] by the PCB m[illegible]nufac[illegible] These diamete[illegible]t the am[illegible]e availabl[illegible]g of sig[illegible] in and out[illegible] array un[illegible]e device. The[illegible]e maximum [illegible]se breakout tra[illegible]B manufacturing li[illegible] the minimum trace width and minimum spacing.

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

- The number of signal layers is defined by the number of I/O signal traces routed in and out of an AP SoC 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 AP SoC and to provide references and isolation for signal layers.

Most PCBs for large AP SoCs 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 AP SoC package ball pitch and PCB via manufacturing constraints) is another factor in trace impedance. A designer often has little control over trace impedance in area of the via array beneath the AP SoC. When traces escape the via array, their width can change to the width of the target impedance (usually 50Ω single-ended).

Decoup[illegible]r placement and discrete termination resistor placement are other areas of trad[illegible]timization. DFM constraints often define a keep-out area around the perimeter of the [illegible]SoC (device footprint) where no discrete components can be placed. The purpose of th[illegible]eep-out area is to allow room for assembly and rework where necessary. For this [illegible]son, the area just outside the keep-out area is one where components compete for placement. It is up to the PCB designer to determine the high priority components. Decoupling capacitor placement constraints ar[illegible]apter 3, Power Distribution System. Termination resistor placeme[illegible]straints [illegible] determine[illegible] through signal integrity simulation, using IBIS or SPIC[illegible]els.

Transmission Lines

The combination o[illegible]ignal trace and a referenc[illegible]s a transmission line. A[illegible] signals in a PCB [illegible]m travel t[illegible]transmissi[illegible]ingle-ended I/O interf[illegible] both the signa[illegible] and the [illegible]lane are n[illegible]mit [illegible] from o[illegible] place to a[illegible] the PCB. [illegible]tial I/O inter[illegible]line is forme[illegible] [illegible]ion of two trac[illegible] 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 signal trace and the reference plane.

The dielectric constant of the material in the vicinity of the trace and reference plane is a property of the PCB laminate materials, and in the case of surface traces, a property of the air or fluid surrounding the board. PCB laminate is typically a variant of FR4, though it can also be an exotic material.

While the dielectric constant of the laminate varies from board to board, it is fairly constant within one board. Therefore, the relative impedance of transmission lines in a PCB is defined most strongly by the trace geometries and tolerances. Impedance variance can occur based on the presence or absence of glass in a local portion of the laminate weave, but this rarely poses issues except in high-speed (>6 Gb/s) interfaces.

Return Currents

An often neglected aspect of transmission lines and their signal integrity is return current. It is incorrect to assume that a signal trace by itself forms a transmission line. Currents flowing in a signal trace have an equal and opposite complimentary current flowing in the reference [illegible] 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 [illegible] formed by the trace and reference plane. While interruption of reference plane continuity beneath a trace is not as dramatic in effect as severing the signal trace, the performance of the transmission line and any devices sharing the reference plane is affected.

It is important to pay attention to reference plane continuity and return current paths. Interruptions of reference plane continuity, such as holes, slots, or isolation splits, cause significant impedance discontinuities in the signal traces. They can also be a significant source of crosstalk and contributor to Power Distribution System (PDS) noise. The importance of return current paths cannot be underestimated.

Chapter 3

Power Distribution System

Introduction

This chapter documents the power distribution system (PDS) for [illegible] SoC devices, including decoupling capacitor selection, placement, and [illegible] geometry. [illegible] simple decoupling method is provided for each device. Basic [illegible] design principles are covered [illegible] well as simulation and analysis methods. This chapter [illegible] the following sections:

- PCB Decoupling Capacitors
- Basic PDS Principles
- Simulation Methods
- PDS Measurements
- Troubleshooting

PCB Decoupling Capacitors

Recommended PCB Capacitors per Device

A simple PCB-decoupling network for the Zynq-7000 AP SoC devices is listed in Table 3-1 and Table 3-2. The optimized quantities of PCB decoupling capacitors assumes that the voltage regulators have stable output voltages and meet the regulator manufacturer's minimum output capacitance requirement.

Decoupling methods other than those presented in these tables can be used, but the decoupling network should be designed to meet or exceed the performance of the simple decoupling networks presented here. The impedance of the alternate network must be less than or equal to that of the recommended network across frequencies from 100 KHz to 100 MHz.

Because device capacitance requirements vary with CLB and I/O utilization, PCB decoupling guidelines are provided on a per-device basis. V_{CCINT}, V_{CCAUX}, V_{CCAUX_IO}, V_{CCBRAM}, and PS supply capacitors are listed as the quantity per device, while V_{CCO} capacitors are listed as

the quantity per I/O bank. Device performance at full utilization is equivalent across all devices when using these recommended networks.

Table 3-1 and Table 3-2 do not provide the decoupling networks required for the GTX transceiver power supplies. For this information, refer to UG476, *7 Series FPGAs GTX/GTH Transceivers User Guide*.

Requi[illegible]pacitor Quantities

Table 3-1 lists th[illegible]3 decoupling capacitor guidelines per V_{CC} supply rail for Zynq-7000 AP SoC devices.

Table 3-1: **Required PCB Capacitor Quantities per Device (PL)**

Package	Device	V_{CCINT}					V_{CCBRAM}				V_{CCAUX}			[illegible]			[illegible]Bank			Bank 0
		680 µF	330 µF	100 µF	4.7 µF	0.47 µF	100 µF	47 µF	4.7 µF	0.47 µF	47 µF	4.7 µF	0.47 µF	[illegible]	4.7 µF	0.47 µF	[illegible]	[illegible]	0.47 µF	47 µF
CLG400	7Z010	0	0	1	2	3	0	1	1	2	1	1	[illegible]	[illegible]	NA	NA	1	3	5	[illegible]
CLG400	7Z020	0	1	0	5	7	1	0	2	3	1	1	[illegible]	[illegible]	NA	NA	1	3	5	[illegible]
CLG484	7Z020	0	1	0	5	7	1	0	2	3	1	2	[illegible]	[illegible]	NA	NA	1	3	5	[illegible]
FBG484	7Z030	0	1	1	0	0	1	0	3	6	1	[illegible]	[illegible]	[illegible]	0	0	1	0	0	[illegible]
FBG676	7Z030	0	1	1	[illegible]	0	1	0	3	6	1	[illegible]	[illegible]	[illegible]	0	0	1	0	0	[illegible]
FBG676	7Z045	2	0	0	[illegible]	0	2	[illegible]	[illegible]	12	1	0	[illegible]	[illegible]	[illegible]	0	1	0	0	[illegible]
FFG676	7Z030	0	1	[illegible]	[illegible]	0	1	[illegible]	[illegible]	[illegible]	1	0	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0	[illegible]
[illegible]	7Z045	2	[illegible]	[illegible]	0	0	2	[illegible]	[illegible]	[illegible]	1	0	0	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0	1
[illegible]	[illegible]	[illegible]	[illegible]	0	0	0	2	[illegible]	[illegible]	12	2	0	0	1	[illegible]	[illegible]	1	0	0	1

Notes:

1. PCB Capacitor specifications are listed in Table 3-3.
2. Total includes all capacitors for all supplies, except for the MGT supplies MGTAVCC and MGTAVTT, which are covered in UG476, 7 Series FPGAs GTX/GTH Transceivers User Guide.

Table 3-2 lists the decoupling requirements for the Processing System (PS).

Table 3-2: **Required PCB Capacitor Quantities per Device (PS)**

Package	Device	V_{CCPINT}			V_{CCPAUX}			V_{CCO_DDR}			V_{CCO_MIO0}			V_{CCO_MIO1}			V_{CCPLL}(1)
		100 µF	4.7 µF	0.47 µF	100 µF	4.7 µF	0.47 µF	100 µF	4.7 µF	0.47 µF	100 µF	4.7 µF	0.47 µF	100 µF	4.7 µF	0.47 µF	10 µF
CLG400	7Z010	1	3	5	1	2	3	2	3	7	1	1	2	1	2	3	1
CLG400	7Z020	1	3	5	1	2	3	2	3	7	1	1	2	1	2	3	1
CLG484	7Z020	1	3	5	1	2	3	2	3	7	1	1	2	1	2	3	1
FBG484	7Z030	1	3	5	1	2	3	2	3	7	1	1	2	1	2	3	1
FBG676	7Z030	1	3	5	1	2	3	2	3	7	1	1	2	1	2	3	1
FBG676	7Z045	1	3	5	1	2	3	2	3	7	1	1	2	1	2	3	1

Table 3-2: **Required PCB Capacitor Quantities per Device (PS) *(Cont'd)***

Package	Device	V_{CCPINT}			V_{CCPAUX}			V_{CCO_DDR}			V_{CCO_MIO0}			V_{CCO_MIO1}			V_{CCPLL}(1)
		100 µF	4.7 µF	0.47 µF	100 µF	4.7 µF	0.47 µF	100 µF	4.7 µF	0.47 µF	100 µF	4.7 µF	0.47 µF	100 µF	4.7 µF	0.47 µF	10 µF
FFG676	7Z030	1	3	5	1	2	3	2	3	7	1	1	2	1	2	3	1
FFG676	7Z045	1	3	5	1	2	3	2	3	7	1	1	2	1	2	3	1
[illegible]	7Z04[illegible]	[illegible]	[illegible]	5	1	2	3	2	3	7	1	1	2	1	2	3	1

Notes:

1. V_{CCPLL} can be tied to V_{CCPAU} [illegible] g a ferrite bead filter (see Figure 5-2, page 52).

Capacitor Specifications

The electrical characteristics of the capacitors in Table 3-1 a [illegible] ecified in Table 3-3, and are followed by guidelines on acceptable [illegible] itutions. [illegible] alent seri [illegible] resistance (ESR) ranges specified for these capacitor [illegible] e over-ridden. [illegible] wever, thi [illegible] requires analysis of the resulting power distributio [illegible] impedance to ensure tha [illegible] resonant impedance spikes result.

Table 3-3: **PCB Capacitor Specifications**

Ideal Value	Value Range(1)	Body Size(2)	Type	ESL Maximum	ESR [illegible]	Voltage Rating(4)	Suggeste [illegible] Part Numb [illegible]
680 µF	C > 680 µF	V-C [illegible]	[illegible] -Terminal Tantalum	[illegible]	15 mΩ < [illegible]	[illegible] .5V	[illegible]
[illegible]	C > 330 [illegible]	[illegible]	2-Terminal Tantalum	[illegible]	15 mΩ < ESR < [illegible]	[illegible]	[illegible] V337M2R5ATE025
330 µF	C > 330 µF	D-Case	2-Terminal Niobium Oxide		15 mΩ < ESR < 40 mΩ	2.5V	NOSD337M002#0035
100 µF	C > 100 µF	1210	2-Terminal Ceramic X7R or X5R	5 nH	10 mΩ < ESR < 60 mΩ	2.5V	GRM32ER60J107ME20L
47 µF	C > 47 µF	1812	2-Terminal Ceramic X7R or X5R	5 nH	50 mΩ < ESR < 90 mΩ	6.3V	GRM43ER61A476KE19L
4.7 µF	C > 4.7 µF	0805	2-Terminal Ceramic X7R or X5R	2 nH	10 mΩ < ESR < 60 mΩ	6.3V	
0.47 µF	C > 0.47 µF	0402	2-Terminal Ceramic X7R or X5R	1.5 nH	5 mΩ < ESR < 15 mΩ	6.3V	

Notes:

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

PCB Bulk Capacitors

The purpose of the bulk capacitors is to cover the low-frequency range between where the voltage regulator stops working and where the on-package ceramic capacitors start working. As specified in Table 3-1 and Table 3-2, all AP SoC supplies require bulk capacitors.

Th[illegible]lum and niobium oxide capacitors specified in Table 3-3 were selected for their values [illegible]d ESR values. They are also ROHS compliant. If another manufacturer's tantalum, nio[illegible]xide, or ceramic capacitors are used, ensure they meet the specifications o[illegible]e 3-3 and are properly evaluated via simulation, s-parameter extraction, or ben[illegible]esting.

PCB High-Frequency Capacitors

Table 3-3 shows the requirements for the 0.47 μF and 4.7[illegible]capacito[illegible]402/0805 packages. Substitutions can be made for some chara[illegible]cs, but not oth[illegible]. For detai[illegible] refer to the notes in Table 3-3.

Capacitor Consolidation Rules

Sometimes a numbe[illegible]f I/O banks are powered [illegible]me voltage (e.g., 1.8V) a[illegible] recommended gui[illegible]es call for multiple bulk ca[illegible]is is also the case for V[illegible] V_{CCAUX}, V_{CCAUX}[illegible]d V_{CCBRA}[illegible]rger AP So[illegible]se many small[illegible]capa[illegible] can be consoli[illegible]into fewe[illegible]lue) capac[illegible] characteri[illegible]e consolid[illegible]itors (ESR and[illegible]e electrical [illegible]of the parallel co[illegible]ion of the recommen[illegible]tors.

For most consolidations of V_{CCO}, V_{CCINT}, V_{CCAUX}, V_{CCAUX_IO}, and V_{CCBRAM} capacitors, large tantalum capacitors with sufficiently low ESL and ESR are readily available.

Example

This example is of a Zynq-7000 AP SoC device with a single memory interface spanning three I/O banks, all powered from the same V_{CCO} voltage. The required PCB capacitor table (Table 3-1) calls for one 100 μF capacitor per bank. These three capacitors can be consolidated into one capacitor since three 100 μF capacitors can be covered by one 330 μF capacitor. The following is then true:

- The ESL of the specified combination is one-third that of the 100 μF capacitor. Three capacitors, each at 3 nH equals one capacitor at 1 nH.
- The ESR of the specified combination is one-third that of the 100 μF capacitor. Three capacitors, each at 10-60 mΩ equals one capacitor at 3–20 mΩ.

Three 100 μF capacitors each with 3 nH ESL and 50 mΩ ESR are replaced by one 330 μF capacitor with a 1 nH ESL and a 25 mΩ ESR.

PCB Capacitor Placement and Mounting Techniques

PCB Bulk Capacitors

Bulk capacitors can be large and sometimes are difficult to place very close to the AP SoC. Fortunately, this is not a problem because the low-frequency energy covered by bulk [illegible]nsitive to capacitor location. Bulk capacitors can be placed almost anywhere [illegible]B, but the best placement is as close as possible to the AP SoC. Capacitor mou[illegible] should follow normal PCB layout practices, tending toward short and wide shapes con[illegible]ng to power planes with multiple vias.

0805 Ceramic Capacitor

The 4.7 µF 0805 capacitor covers the middle frequency ra[illegible]ome impact on its performance. The capacitor should be placed as [illegible]as possible [illegible]AP SoC. A[illegible] placement within two inches of the device's outer e[illegible]cceptable. The capacitor mounting (solder lands, traces, and vias) should b[illegible]zed for low inductance. Vi[illegible] should be butted directly against the pads. Vias c[illegible]ated at the ends of the pa[illegible] Figure 3-1B), but are more optimally located at t[illegible]f the pads (see Figure 3-1[illegible] placement at the sides of the pads decreases the [illegible]s overall parasitic inducta[illegible] increasing the mutu[illegible]nductive coupling of one [illegible]her. Dual vias can be pla[illegible] both sides of the [illegible] (see Figur[illegible]D) for even[illegible]itic inductance, but wit[illegible] diminishing retu[illegible]

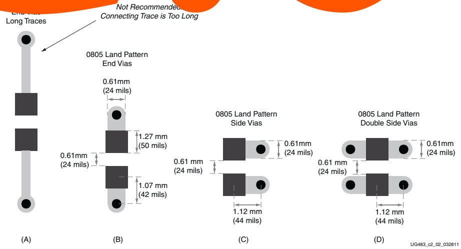

Figure 3-1: **Example 0805 Capacitor Land and Mounting Geometries**

0402 Ceramic Capacitor

The 0.47 µF 0402 capacitor covers the high frequency range. These capacitors should be placed as close to the AP SoC as possible, within 0.5 inches of the outer edge. Follow the same land pattern rules as with the 0805 ceramic capacitors.

Basic PD[illegible]rinciples

The purpose of th[illegible]S and the properties of its components are discussed in this section. The important aspe[illegible] of capacitor placement, capacitor mounting, PCB geometry, and PCB stackup recommendations are also described.

Noise Limits

In the same way that devices in a system have a re[illegible]nt for the amount of curre[illegible] consumed by the power system, there is also a re[illegible] for the cleanliness of the [illegible] This cleanliness requirement specifies a maximu[illegible] of noise present on the [illegible] supply, often referred to as ripple voltage (V_{RIPP}[illegible]igital devices, including [illegible] Zynq-7000 AP SoC [illegible]ces, require that V_{CC} sup[illegible]ctuate more than ±5% [illegible] nominal V_{CC} value[illegible]s means th[illegible]e peak-to-p[illegible]must be no more than 1[illegible] the nominal V_{C}[illegible]is docum[illegible]m V_{CC} is u[illegible] for the fol[illegible]ng AP[illegible] power suppl[illegible]$_{INT}$, V_{CCO}, [illegible]$_{AUX\_IO}$, V_{CC}[illegible]s. This [illegible]al V_{CC} is e[illegible]nominal value [illegible] sheet. If not, [illegible]$_{RIPPLE}$ must be adjusted to [illegible]ue correspondingly less [illegible]%.

The power consumed by a digital device varies over time and this variance occurs on all frequency scales, creating a need for a wide-band PDS to maintain voltage stability.

- Low-frequency variance of power consumption is usually the result of devices or large portions of devices being enabled or disabled. This variance occurs in time frames from milliseconds to days.
- High-frequency variance of power consumption is the result of individual switching events inside a device. This occurs on the scale of the clock frequency and the first few harmonics of the clock frequency up to about 5 GHz.

Because the voltage level of V_{CC} for a device is fixed, changing power demands are manifested as changing current demand. The PDS must accommodate these variances of current draw with as little change as possible in the power-supply voltage.

When the current draw in a device changes, the PDS cannot respond to that change instantaneously. As a consequence, the voltage at the device changes for a brief period before the PDS responds. Two main causes for this PDS lag correspond to the two major PDS components: the voltage regulator and decoupling capacitors.

The first major component of the PDS is the voltage regulator. The voltage regulator observes its output voltage and adjusts the amount of current it is supplying to keep the output voltage constant. Most common voltage regulators make this adjustment in milliseconds to microseconds. Voltage regulators effectively maintain the output voltage for events at all frequencies from DC to a few hundred kHz, depending on the regulator (some are effective at regulating in the low MHz). For transient events that occur at [illegible]ove this range, there is a time lag before the voltage regulator responds to the new [illegible]and level.

For example, if [illegible]evice's current demand increases in a few hundred picoseconds, the voltage at the de[illegible] sags by some amount until the voltage regulator can adjust to the new, higher level o[illegible]quired current. This lag can last from microseconds to milliseconds. A second component is needed to substitute for the regulator during this time, preventing the voltage from sagging.

This second major PDS component is the decoupling [illegible]tor (also kno[illegible] a bypass capacitor). The decoupling capacitor works as the d[illegible]local energy storage. The capacitor cannot provide DC power because it sto[illegible] a small amount of energy (voltage regulator provides DC power). This local [illegible]torage should respond ve[illegible] quickly to changing current demands. The capac[illegible]tively maintain power-su[illegible] voltage at frequencies from hundreds of kHz to [illegible]f MHz (in the millisecon[illegible] nanoseconds range[illegible]iscrete decoupling capaci[illegible] useful for events occur[illegible] above or below th[illegible]nge.

For example, [illegible]nt deman[illegible]vice increase[illegible]e voltage [illegible]some amou[illegible]e capacitors ca[illegible]ge to the device. [illegible]demand in the device [illegible]tains this new level for [illegible]seconds, the voltage-regulator circuit, operating in parallel with the decoupling capacitors, replaces the capacitors by changing its output to supply this new level of current.

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

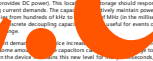

Figure 3-2: **Simplified PDS Circuit**

Figure 3-3 shows a simplified PDS circuit with all reactive components represented by a frequency-dependent resistor.

Figure 3-3: **Further Simplified PDS Circuit**

Role of Inductance

Inductance is the property of the capacitors and t[illegible]urrent paths that slows do[illegible] changes in current flow. Inductance is the reason [illegible]acitors cannot respond instantaneously to transient currents or to chang[illegible]ccur at frequencies highe[illegible] their effective range.

Inductance can be [illegible]ught of as th[illegible] momentum [illegible]harge moving through [illegible] conductor repre[illegible] some an[illegible]urrent. If t[illegible]rent change[illegible]e cha[illegible] moves at a di[illegible]rate. Bec[illegible]ntum (store[illegible]s associat[illegible] with thi[illegible]me amoun[illegible]d energy is req[illegible]or speed up the [illegible]ow. The greater the indu[illegible]ce, the greater the resista[illegible] change, and the longer the time required for the current level to change. A voltage develops across the inductance as this change occurs.

The PDS, made up of a regulator and multiple stages of decoupling capacitors, accommodates the device current demand and responds to current transients as quickly as necessary to maintain the voltage within the specified limits. When these current demands are not met, the voltage across the device's power supply changes. This is observed as noise. Inductance in the current path of the capacitors should be minimized, because it retards the ability of decoupling capacitors to quickly respond to changing current demands.

Inductances occur between the AP SoC device and capacitors and between the capacitors and the voltage regulator (see Figure 3-2). These inductances occur as parasitics in the capacitors and in all PCB current paths. It is important that each of these parasitics be minimized.

Capacitor Parasitic Inductance

The capacitance value is often considered to be a capacitors's most important characteristic. In power system applications, the parasitic inductance (ESL) has the same or

greater importance. Capacitor package dimensions (body size) determine the amount of parasitic inductance. Physically small capacitors usually have lower parasitic inductance than physically large capacitors.

Requirements for choosing decoupling capacitors:

- For a specific capacitance value, choose the smallest package available.

  - or -

- For a specifi[illegible] size (essentially a fixed inductance value), choose the highest capacitance [illegible] available in that package.

Surface-mount chip capacitors are the smallest capacitors available and are a good choice for discrete decoupling capacitors:

- For values from 100 µF to very small values such [illegible] µF, ceramic [illegible] X5R typ[illegible] capacitors are usually used. These capacitors [illegible] parasitic inductance and [illegible] ESR, with an acceptable temperature characte[illegible]
- For larger values, such as 47 µF to 1000 µF, [illegible] capacitors are usually used[illegible] capacitors have a low parasitic inductance [illegible] ESR, giving them a low [illegible] factor and consequently a very wide range [illegible] frequencies.

If tantalum [illegible] are not [illegible] or cannot [illegible] ESR, low-inductance electrolytic [illegible] can be [illegible] they [illegible] values[illegible] Other [illegible] with [illegible] characteristics are [illegible], POSCAP, and [illegible] SMT).

A real capacitor of any type then not only has capacitance characteristics but also inductance and resistance characteristics. Figure 3-4 shows the parasitic model of a real capacitor. A real capacitor should be treated as an RLC circuit (a circuit consisting of a resistor (R), an inductor (L), and a capacitor (C), connected in series).

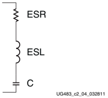

Figure 3-4: **Parasitics of a Real, Non-Ideal Capacitor**

Figure 3-5 shows a real capacitor's impedance characteristic. Overlaid on this plot are curves corresponding to the capacitor's capacitance and parasitic inductance (ESL). These two curves combine to form the RLC circuit's total impedance characteristic, softened or sharpened by the capacitor's ESR.

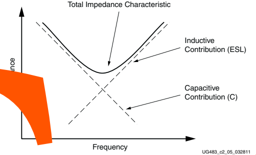

Figure 3-5: **Contribution of Parasitics to Total Impedanc[illegible]**

As capacitive value is increased, the capacitive curve [illegible] down and left. As parasitic inductance is decreased, the inductive curve move[illegible] and right. Because parasiti[illegible] inductance for capacitors in a specific package is [illegible] inductance curve for cap[illegible] in a specific package remains fixed.

As different capacito[illegible] values are selected in the [illegible]ge, the capacitive curve [illegible] up and down again[illegible]he fixed inductance curve, [illegible] Figure 3-6.

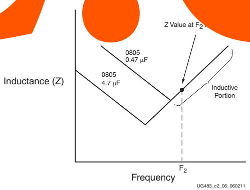

Figure 3-6: **Effective Frequency Example**

The low-frequency capacitor impedance can be reduced by increasing the value of the capacitor; the high-frequency impedance can be reduced by decreasing the inductance of the capacitor. While it might be possible to specify a higher capacitance value in the fixed package, it is not possible to lower the inductance of the capacitor (in the fixed package) without putting more capacitors in parallel. Using multiple capacitors in parallel divides the

parasitic inductance, and at the same time, multiplies the capacitance value. This lowers both the high and low frequency impedance at the same time.

PCB Current Path Inductance

The parasitic inductance of current paths in the PCB have three distinct sources:

- Capa[illegible]ing
- PCB power [illegible]round planes
- AP SoC mount[illegible]

Capacitor Mounting Inductance

Capacitor mounting refers to the capacitor's solder lan[illegible] the PCB, t[illegible] (if any) between the land and via, and the via.

The vias, traces, and capacitor mounting pads of [illegible]nal capacitor contribute inductance between 300 pH to 4 nH depending [illegible]ecific geometry.

Because the current path's inductance is proport[illegible] loop area the current tra[illegible] it is important to m[illegible]ize this loop size. The loo[illegible]f the path through one plane, up through [illegible] via, throu[illegible] connecting [illegible] land, through the capa[illegible] through the oth[illegible]d and co[illegible]race, dow[illegible]ther via [illegible]nto the other plane[illegible]wn in Figu[illegible]

Figure 3-7: **Example Cutaway View of PCB with Capacitor Mounting**

A connecting trace length has a large impact on the mounting's parasitic inductance and if used, should be as short and wide as possible. When possible, a connecting trace should not be used and the via should butt up against the land. Placing vias to the side of the capacitor lands or doubling the number of vias, further reduces the mounting's parasitic inductance.

Some PCB manufacturing processes allow via-in-pad geometries, an option for reducing parasit[illegible] Using multiple vias per land is important with ultra-low inductance capacitors, s[illegible]everse aspect ratio capacitors that place wide terminals on the sides of the capacitor bo[illegible]stead of the ends.

PCB layout engine[illegible] often try to squeeze more parts into a small area by sharing vias among multiple capacitors. This technique should not be used under any circumstances. PDS improvement is very small when a second capacitor is co[illegible]xisting capacitor's vias. For a larger improvement, reduce the tot[illegible]mber [illegible]rs and maintain a one-to-one ratio of lands to vias.

The capacitor mounting (lands, traces, and vias) ty[illegible]ontributes about the same amount or more inductance than the capacitor's [illegible]sitic self-inductance.

Plane Inductance

Some inductance [illegible]sociated w[illegible]e PCB powe[illegible]d planes. The geometry [illegible] these planes de[illegible]nes their [illegible].

Curren[illegible] as it flows [illegible] point to anoth[illegible]y similar to skin [illegible] power and ground p[illegible]s. Inductance in planes ca[illegible]escribed as spreading inductance and is specified in units of henries per square. The square is dimensionless; the shape of a section of a plane, not the size, determines the amount of inductance.

Spreading inductance acts like any other inductance and resists changes to the amount of current in a power plane (the conductor). The inductance retards the capacitor's ability to respond to a device's transient currents and should be reduced as much as possible. Because the designer's control over the X-Y shape of the plane can be limited, the only controllable factor is the spreading inductance value. This is determined by the thickness of the dielectric separating a power plane from its associated ground plane.

For high-frequency power distribution systems, power and ground planes work in pairs, with their inductances coexisting dependently with each other. The spacing between the power and ground planes determines the pair's spreading inductance. The closer the spacing (the thinner the dielectric), the lower the spreading inductance. Approximate values of spreading inductance for different thicknesses of FR4 dielectric are shown in Table 3-4.

Table 3-4: **Capacitance and Spreading Inductance Values for Different Thicknesses of FR4 Power-Ground Plane Sandwiches**

Dielectric Thickness		Inductance	Capacitance	
(micron)	**(mil)**	**(pH/square)**	**(pF/in2)**	**(pF/in2)**
102	4	130	225	35
[illegible]	2	65	450	70
25	1	32	900	140

Decreased sprea[illegible] inductance corresponds to closer spacing of V_{CC} and GND planes. When possible, pl[illegible] the V_{CC} planes directly adjacent to the GND planes in the PCB stackup. Facing V_{CC} and GN[illegible] planes are sometimes referred to as sandwiches. While the use of V_{CC} – GND sandwiches was not necessary in the past for previous tec[illegible] (lead frames, wire bond packages), the speeds involved and the sheer am[illegible]ired for fast, dense devices often demand it.

However, because of the presence of substrate dec[illegible] capacitors in Zynq-7000 AP [illegible] devices, there is a limit to the amount of fast tran[illegible]rent demanded from PCB decoupling capacitors. This means that there is li[illegible]fit from dielectric thickne[illegible] below 50μ (2 mil). Dielectric thickness of 50μ or [illegible]een V_{CC} and GND layers i[illegible] sufficient for Zynq-7000 AP SoC devices.

Besides offering a[illegible]-inductanc[illegible]ent path, p[illegible] sandwiches also offer [illegible] high-frequency[illegible]upling ca[illegible] As the pla[illegible]es and as t[illegible] separa[illegible] between po[illegible] ground p[illegible]ases, the va[illegible]creases. [illegible]uare inch is [illegible] Table 3-4. Howe[illegible]or capacitance arising from these PCB power-ground plane pairs is generally inconsequential, given the substrate decoupling capacitors present in Zynq-7000 AP SoC devices.

AP SoC Mounting Inductance

The PCB solder lands and vias that connect the AP SoC power pins (V_{CC} and GND) contribute an amount of parasitic inductance to the overall power circuit. For existing PCB technology, the solder land geometry and the dogbone geometry are mostly fixed, and parasitic inductance of these geometries does not vary. Via parasitic inductance is a function of the via length and the proximity of the opposing current paths to one another.

The relevant via length is the portion of the via that carries transient current between the AP SoC solder land and the associated V_{CC} or GND plane. Any remaining via (between the power plane and the PCB backside) does not affect the parasitic inductance of the via (the shorter the via between the solder lands and the power plane, the smaller the parasitic inductance). Parasitic via inductance in the AP SoC mounting is reduced by keeping the relevant V_{CC} and GND planes as close to the AP SoC as possible (close to the top of the PCB stackup).

Device pinout arrangement determines the proximity of opposing current paths to one another. Inductance is associated with any two opposing currents (for example, current

flowing in a V_{CC} and GND via pair). A high degree of mutual inductive coupling between the two opposing paths reduces the loop's total inductance. Therefore, when given a choice, V_{CC} and GND vias should be as close together as possible.

The via field under an AP SoC has many V_{CC} and GND vias, and the total inductance is a function of the proximity of one via to another:

- [illegible] plies (V_{CCINT} and V_{CCAUX}), opposing current is between the V_{CC} and GND pi[illegible]
- For I/O V_{CC} s[illegible]ies (V_{CCO}), opposing current is between any I/O and its return current path, whether [illegible]ied by a V_{CCO} or GND pin.

To reduce parasitic inductance:

- Core V_{CC} pins such as V_{CCINT} and V_{CCAUX} are placed i[illegible]eckerb[illegible]ngement in the pinout.
- V_{CCO} and GND pins are distributed among the [illegible].

Every I/O pin in the Zynq-7000 AP SoC device pi[illegible]jacent to a return-curren[illegible]

AP SoC pinout arrangement determines the PCB [illegible]ement. The PCB designer [illegible] control the proximi[illegible] opposing current paths b[illegible]ol over the trade-offs be[illegible] the capacitor's m[illegible]ng induc[illegible]d AP SoC [illegible]ductance:

- Both mo[illegible]nductanc[illegible]ed by placi[illegible] the PCB [illegible]lf and plac[illegible]pacitors on the [illegible]ng the [illegible]pacitor's via length).
- If power planes are placed in the PCB stackup's bottom half, the capacitors must be mounted on the PCB backside. In this case, AP SoC mounting vias are already long, and making the capacitor vias long (by coming down from the top surface) is a bad practice. A better practice is to take advantage of the short distance between the underside of the PCB and the power plane of interest, mounting capacitors on the underside.

PCB Stackup and Layer Order

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

- High-priority supplies should be placed closer to the AP SoC (in the PCB stackup's top half)
- Low-priority supplies should be placed farther from the AP SoC (in the PCB stackup's bottom half)

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

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

Capacitor Effective Frequency

Every capacitor has a narrow frequency band wh[illegible]ost effective as a decoup[illegible] capacitor. This band is centered at the capacitor[illegible]nant frequency F_{RSELF}. Th[illegible] effective frequency [illegible]nds of some capacitors ar[illegible]n others. A capacitor's E[illegible] determines the ca[illegible]tor's quality (Q) factor, and[illegible]r can determine the wi[illegible] the effective fre[illegible]y band:

- Tantalu[illegible]tors gene[illegible] very *wide* ef[illegible]
- [illegible] chip capacitors wit[illegible]er ESR, generally have [illegible]*rrow* effective frequency band.

An ideal capacitor only has a capacitive characteristic, whereas *real* non-ideal capacitors also have a parasitic inductance (ESL) and a parasitic resistance (ESR). These parasitics work in series to form an RLC circuit (Figure 3-4). The RLC circuit's resonant frequency is the capacitor's self-resonant frequency.

To determine the RLC circuit's resonant frequency, use Equation 3-1:

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

Equation 3-1

Another method of determining the self-resonant frequency is to find the minimum point in the impedance curve of the equivalent RLC circuit. The impedance curve can be computed or generated in SPICE using a frequency sweep. See the Simulation Methods section for other ways to compute an impedance curve.

It is important to distinguish between the capacitor's self-resonant frequency and the mounted capacitor's effective resonant frequency when the capacitor is part of the system, F_{RIS}. This corresponds to the resonant frequency of the capacitor with its parasitic inductance, plus the inductance of the vias, planes, and connecting traces between the capacitor and the AP SoC.

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

Mounted parasitic inductance is a combination of the capacitor's own parasitic inductance [illegible]uctance of: PCB lands, connecting traces, vias, and power planes. Vias traverse a [illegible]he device when capacitors are mounted on the PCB backside. For a board with a [illegible]d thickness of 1.524 mm (60 mils), these vias contribute approximately 300 pH to 1,500 [illegible]the capacitor’s mounting parasitic inductance, L_{MOUNT}) depending on the spacing betw[illegible] vias. Wider-spaced vias and vias in thicker boards have higher inductance.

To determine the capacitor’s total parasitic inductance in the [illegible] capacitor's parasitic inductance, L_{SELF}, is added to the mounting’s pa[illegible] induc[illegible]OUNT:

$$L_{IS} = L_{SELF} + L_{MOUNT}$$ Equatio[illegible]

For example, using X7R Ceramic Chip capacitor i[illegible]dy size:

C = 0.01 µF (selected by user)

L_{SELF} = 0.9 nH ([illegible]acitor data sheet paramete[illegible]

F_{RSELF} = 53 [illegible] capacito[illegible]t paramet[illegible]

L[illegible] (based o[illegible]nting geomet[illegible]

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

$$L_{IS} = L_{SELF} + L_{MOUNT} = 0.9\,nH + 0.8\,nH$$
$$L_{IS} = 1.7\,nH$$ Equation 3-3

The values from the example are used to determine the mounted capacitor resonant frequency (F_{RIS}). Using Equation 3-1:

$$F_{RIS} = \frac{1}{2\pi\sqrt{L_{IS}C}}$$ Equation 3-4

$$F_{RIS} = \frac{1}{2\pi\sqrt{(1.7\times10^{-9}H)\cdot(0.01\times10^{-6}F)}} = 38\times10^{6}Hz$$ Equation 3-5

F_{RSELF} is 53 MHz, but F_{RIS} is lower at 38 MHz. The addition of mounting inductances shifts the effective-frequency band down.

A decoupling capacitor is most effective at the narrow-frequency band around its resonant frequency, and thus, the resonant frequency must be reviewed when choosing a capacitor collection to build up a decoupling network. This being said, capacitors can be effective at frequencies considerably higher and lower than their resonant frequency. Recall that capacitors of differing values in the same package share the same inductance curve. As

shown in Figure 3-6, for any given frequency along the inductive portion of the curve, the capacitors are equally effective.

Capacitor Anti-Resonance

One problem associated with combinations of capacitors in a PDS of an AP SoC is [illegible]kes in the PDS aggregate impedance. The cause for these spikes is a bad combinati[illegible]gy storage elements in the PDS (intrinsic capacitances, discrete capacitors, pa[illegible]nductances, and power and ground planes).

Anti-resonance ca[illegible]ise between any two consecutive stages of a power distribution system, such as bet[illegible]en the high-frequency PCB capacitors and the PCB plane capacitance. The inter-plane capacitance of the power and ground planes genera[illegible] a high-Q factor. If the high-frequency PCB capacitors also are high-Q, the c[illegible]een the high-frequency discrete capacitors and the plane capac[illegible] might ex[illegible] high-impedance anti-resonance peak. If the AP SoC h[illegible]gh transient cu[illegible]nt deman[illegible] this frequency (as a stimulus), a large noise voltag[illegible]cur.

To correct this type of problem, the characteristic[illegible]gh-frequency discrete cap[illegible] or the characteristics of the V_{CC} and ground pla[illegible]e changed, or AP SoC ac[illegible] shifted to a different frequency away from the r[illegible]

Capacitor Pl[illegible]ment [illegible]ound

To perfor[illegible]upling fu[illegible]acitors shoul[illegible]e being [illegible]

Increased spacing between the AP SoC and decoupling capacitor increases the current flow distance in the power and ground planes, and it often increases the current path's inductance between the device and the capacitor.

The inductance of this current path (the loop followed by current as it travels from the V_{CC} side of the capacitor to the V_{CC} pin[s] of the AP SoC, and from the GND pin[s] of the AP SoC to the GND side of the capacitor[s]), is proportional to the loop area. Inductance is decreased by decreasing the loop area.

Shortening the distance between the device and the decoupling capacitor reduces the inductance, resulting in a less impeded transient current flow. Because of typical PCB dimensions, this lateral plane travel tends to be less important than the phase relationship between the AP SoC noise source and the mounted capacitor.

The phase relationship between the AP SoC's noise source and the mounted capacitor determines the capacitor's effectiveness. For a capacitor to be effective in providing transient current at a certain frequency (for example, the capacitor's resonant frequency), the phase relationship, based on the distance travelled by the current from the AP SoC to the capacitor, must be within a fraction of the corresponding period.

The capacitor’s placement determines the length of the transmission line interconnect (in this case, the power and ground plane pair) between the capacitor and AP SoC. The propagation delay of this interconnect is the key factor.

AP SoC noise falls into certain frequency bands, and different sizes of decoupling capacitors take care of different frequency bands. Thus, capacitor placement requirements are [illegible] by each capacitor’s effective frequency.

When the A[illegible]tiates a current demand change, it causes a small local disturbance in the PDS voltag[illegible]int in the power and ground planes). Before it can counteract this, the decoupling capa[illegible] must first sense a voltage difference.

A finite time delay (Equation 3-6) occurs between the start of the disturbance at the AP SoC power pins and the point when the capacitor senses the distur[illegible]

$$\textit{Time Delay} = \frac{\textit{Distance from the FPGA power pi[illegible] capacitor}}{\textit{Signal propagation speed thro[illegible] dielectric}} \qquad \textit{Equatio[illegible]}$$

The dielectric is the substrate of the PCB where [illegible] planes are embedded.

Another delay of the[illegible]me duration occurs whe[illegible]ensation current from th[illegible] capacitor flows to [illegible]AP SoC. For any transient c[illegible]nd in the AP SoC, a rou[illegible] delay occurs bef[illegible]ny relief [illegible] the AP So[illegible]

- Neglig[illegible]gy is trans[illegible]e AP SoC wit[illegible] greater than [illegible] a demand fr[illegible]s wavelength.
- Energy transferred to the AP SoC increases from 0% at one-quarter of a wavelength to 100% at zero distance.
- Energy is transferred efficiently from the capacitor to the AP SoC when capacitor placement is at a fraction of a quarter wavelength of the AP SoC power pins. This fraction should be small because the capacitor is also effective at some frequencies (shorter wavelengths) above its resonant frequency.

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

When using large numbers of external termination resistors or passive power filtering for transceivers, priority should be given to these over the decoupling capacitors. Moving away from the device in concentric rings, the termination resistors and transceiver supply filtering should be closest to the device, followed by the smallest-value decoupling capacitors, then the larger-value decoupling capacitors.

V_{REF} Stabilization Capacitors

In V_{REF} supply stabilization, one capacitor per pin is placed as close as possible to the V_{REF} pin. The capacitors used are in the 0.022 µF – 0.47 µF range. The V_{REF} capacitor's primary function is to reduce the V_{REF} node impedance, which in turn reduces crosstalk coupling. Since no low-frequency energy is needed, larger capacitors are not necessary.

[illegible] en Internal V_{REF} is not used. Internal V_{REF} is a feature in the Zynq-7000 AP SoC devi[illegible] in the reference voltage rail is generated internally, which in turn allows the V_{REF} [illegible] o be used as regular I/O pins.

Power Supply [illegible]onsolidation

Powering 1.8V V_{CCO} and V_{CCAUX} from a common PCB plane i[illegible]000 AP SoC designs. However, careful consideration must be given t[illegible]er supply [illegible]n particula[illegible] any noise on the V_{CCO} rail should not violate the rec[illegible]ded operating [illegible]ndition ra[illegible] for the V_{CCAUX} supply. See DS187, *Zynq-7000 EPP [illegible]0 and XC7Z020): DC and A[illegible] Switching Characteristics*

Unconnected V_{CCO} Pins

In some cases, on[illegible] more I/O banks in an AP S[illegible]ed (for example, when [illegible] SoC has far more [illegible] pins than [illegible]n requires[illegible]s, it might b[illegible]esirab[illegible] leave the ban[illegible]ciated V[illegible] connected, [illegible] layout constrai[illegible]ding of po[illegible]ound planes [illegible]s obstacles to [illegible]ring and exiting the p[illegible] array, more copper are[illegible]e for other planelets in the otherwise used plane layer).

Leaving the V_{CCO} pins of unused I/O banks floating reduces the level of ESD protection on these pins and the I/O pins in the bank. However, ESD events at the unconnected solder balls in the inner rows of the pinout array are unlikely and not considered a high risk.

Simulation Methods

Simulation methods, ranging from very simple to very complex, exist to predict the PDS characteristics. An accurate simulation result is difficult to achieve without using a fairly sophisticated simulator and taking a significant amount of time.

Basic lumped RLC simulation is one of the simplest simulation methods. Though it does not account for the distributed behavior of a PDS, it is a useful tool for selecting and verifying that combinations of decoupling capacitor values will not lead to large anti-resonances. Lumped RLC simulation is a good method for establishing equivalence of decoupling networks, such as evaluating an alternative to the capacitors of Table 3-3.

Lumped RLC simulation is performed either in a version of SPICE or other circuit simulator, or by using a mathematical tool like MathCAD or Microsoft Excel. Istvan Novak publishes a free Excel spreadsheet for lumped RLC simulation (among other useful tools for PDS simulation) on his website under **Tool Download**:

http://www.electrical-integrity.com

[illegible] few EDA tool vendors for PDS design and simulation. These tools span a wide rang [illegible] nistication levels.

Table 3-5: **EDA T [illegible] or PDS Design and Simulation**

Tool	Vendor	Website URL
ADS	Agilent	http://www.agilen [illegible]
SIwave, HFSS	Ansoft	http://w [illegible] soft.c [illegible]
Specctraquest Power Integrity	Cadence	htt [illegible] w.cadence.com
Speed 2000, PowerSI, PowerDC	Sigrity	[illegible] ww.sigrity.com
Hyperlynx PI	Mentor	[illegible] ww.mentor.com

PDS Meas [illegible] ements

Measurement [illegible] e used t [illegible] whether a [illegible] e measur [illegible] a unique ta [illegible] ny specialized [illegible] en developed. [illegible] describes the noise [illegible] itude and noise spectr [illegible] urements.

Noise Magnitude Measurement

Noise measurement must be performed with a high-bandwidth oscilloscope (minimum 3 GHz oscilloscope and 1.5 GHz probe or direct coaxial connection) on a design running realistic test patterns. The measurement is taken at the device's power pins or at an unused I/O driven High or Low (referred to as a *spyhole measurement*).

V_{CCINT} and V_{CCAUX} can only be measured at the PCB backside vias. V_{CCO} can also be measured this way, but more accurate results are obtained by measuring static (fixed logic level) signals at unused I/Os in the bank of interest.

When making the noise measurement on the PCB backside, the via parasitics in the path between the measuring point and AP SoC must be considered. Any voltage drop occurring in this path is not accounted for in the oscilloscope measurement.

PCB backside via measurements also have a potential problem: decoupling capacitors are often mounted directly underneath the device, meaning the capacitor lands connect directly to the V_{CC} and GND vias with surface traces. These capacitors confuse the measurement by acting like a short circuit for the high-frequency AC current. To make sure

the measurements are not shorted by the capacitors, remove the capacitor at the measurement site (keep all others to reflect the real system behavior).

When measuring V_{CCO} noise, the measurement can be taken at an I/O pin configured as a driver to logic 1 or logic 0. In most cases, the same I/O standard should be used for this “spyhole” as for the other signals in the bank. Measuring a static logic 0 shows the crosstalk [illegible] routing, package routing) induced on the victim. Measuring a static logic 1 shows [illegible] crosstalk components as well as the noise present on the V_{CCO} net for the I/O bank. [illegible] subtracting (coherently in time) the noise measured on static logic 0 from the noise measured on static logic 1, the noise on V_{CCO} at the die can be viewed. For an accurate result, the static logic 0 and static logic 1 noise must be measured at the same I/O location. This means storing the time-domain waveform information from both logic states and performing the subtraction operation on the two waveforms in a post-process math computation tool such as MATLAB or Excel.

Oscilloscope Measurement Methods

There are two basic ways of using the oscilloscope [illegible] power system noise, each [illegible] different purpose. The first surveys all possible [illegible]ts, while the second is used [illegible] focusing on individual noise sources.

- Place the oscilloscope in infinite persistence [illegible] acquire all noise over a long [illegible] period (many seconds or minutes). If the design [illegible] many different modes [illegible] using different sources in [illegible] amounts, [illegible] conditions [illegible] modes [illegible] should [illegible]ration while the oscilloscope is [illegible]urement.
- [illegible] oscilloscope in averaging mode and trigger on [illegible] aggressor event. This can show the amount of noise correlated with the aggressor event (any events asynchronous to the aggressor are removed through averaging).

Power system noise measurements should be made at a few different AP SoC locations to ensure that any local noise phenomena are captured.

Figure 3-8 shows an averaged noise measurement taken at the V_{CCO} pins of a sample design. In this case, the trigger was the clock for an I/O bus interface sending a 1-0-1-0 pattern at 250 Mb/s.

Figure 3-8: **Averaged Mea[illegible]ement of V_{CCO} Supply with [illegible] Sending Patterns at 250[illegible]**

Figure 3-9 show[illegible] infinite p[illegible] noise me[illegible] the same desi[illegible] with [illegible] wider variety [illegible]ctivity. B[illegible] infinite pers[illegible]nes all no[illegible] events o[illegible] period, bo[illegible]d and non-co[illegible]ary aggressor, [illegible] system excursions are [illegible].

Figure 3 [illegible] **Infinite Persistence Meas[illegible] ame Supply**

The measureme[illegible] own in Fi[illegible] nd Figure [illegible] the peak-to[illegible] ak noi[illegible] the peak-to-[illegible] se is outs[illegible] cified accep[illegible] sheet val[illegible] V[illegible] + 5[illegible] upling ne[illegible] dequate or a [illegible] PCB layout.

3.0.1 Noise Spectrum Measurements

Having the necessary information to improve the decoupling network requires additional measurements. To determine the frequencies where the noise resides, noise power spectrum measurement is necessary. A spectrum analyzer or a high-bandwidth oscilloscope coupled with FFT math functionality can accomplish this.

The FFT math function can be built into the oscilloscope, however, many of these functions do not have resolution sufficient to give a clear picture of the noise spectrum. Alternatively, a long sequence of time-domain data can be captured from an oscilloscope and converted to frequency domain using MATLAB or other post-processing software supporting FFT. This method has the advantage of showing as much resolution as the user is willing to process. If neither math capacity is available, the noise frequency content can be approximated by visually examining the time-domain waveform and estimating the individual periodicities present in the noise.

A spectrum analyzer is a frequency-domain instrument, showing the frequency content of a voltage signal at its inputs. Using a spectrum analyzer, the user sees the exact frequencies where the PDS is inadequate.

Excessive noise at a certain frequency indicates a frequency where the PDS impedance is too high for the device's transient current demands. Using this information, the designer can modify the PDS to accommodate the transient current at the specific frequency. This is accomplished by either adding capacitors with effective frequencies close to the noise frequency or otherwise lowering the PDS impedance at the critical frequency.

T[illegible]ectrum measurement should be taken in the same manner as the peak-to-peak noise [illegible] directly underneath the device, or at a static I/O driven High or Low. A spectrum an[illegible]kes its measurements using a 50Ω cable instead of an active probe.

- A good meth[illegible]ttaches the measurement cable through a coaxial connector tapped into the power [illegible]d ground planes close to the device. This is *not* available in most cases.
- Another method attaches the measurement cable at th[illegible]ing capacitor in the vicinity of the device that has been [illegible]ved. The c[illegible]nter conductor and shield are soldered directly to the [illegible]tor lands. Alternatively, a pr[illegible] station with 50Ω RF probes can be used to tou[illegible]ecoupling capacitor lands[illegible]

To protect the spectrum analyzer's sensitive front[illegible]itry, add a DC blocking ca[illegible] or attenuator in line. This isolates the spectrum [illegible]om the device supply vo[illegible]

Figure 3-10 is an ex[illegible]ple of a noise spectrum m[illegible] of the V_{CCO} power-sup[illegible] noise, with multi[illegible]O sendin[illegible]ns at 100 [illegible]

Figur[illegible] **Screens**[illegible]**rum Analyz**[illegible] **of V**[illegible]

[illegible] Decoupling [illegible]ork Design

If a highly optimized PDS is needed, measurements and simulations of a prototype system can inform the PDS design. Using knowledge of the noise spectrum generated by the prototype system along with knowledge of the system's power system impedance, the unique transient current of the design can be determined and accommodated.

To measure the noise spectrum of the design under operating conditions, use either a spectrum analyzer or an oscilloscope with FFT. The power system impedance can be determined either through direct measurement or simulation, or a combination of these two as there are often many variables and unknowns.

Both the noise spectrum and the impedance are functions of frequency. By examining the quotient of these per frequency point, transient current as a function of frequency is computed (Equation 3-7):

$$I(f) = \frac{V(f)\ \text{From Spectrum Analyzer}}{Z(f)\ \text{From Network Analyzer}} \qquad \text{Equation 3-7}$$

Using the data sheet's maximum voltage ripple value, the impedance value needed at all frequencies can be determined. This yields a target impedance as a function of frequency. A

specially designed capacitor network can accommodate the specific design's transient current.

Troubleshooting

In some [illegible]oper design work is done up-front, but noise problems still exist. This next section d[illegible]s possible issues and suggested resolution methods.

Possibility 1: Ex[illegible]sive Noise from Other Devices on the PCB

Sometimes ground and/or power planes are shared among man[illegible]nd noise from an inadequately decoupled device affects the PDS at other [illegible]uses of this noise are:

- RAM interfaces with inherently high-transient c[illegible]emands resulting either fr[illegible] temporary periodic contention or high-curren[illegible]
- Large ASICs

When unacceptable [illegible]ounts of noise are measu[illegible]at these devices, the loc[illegible] and the compone[illegible]coupling [illegible]works should [illegible]

Possibility [illegible]sitic Ind[illegible]f Planes, [illegible]aces

[illegible] the decoupling netwo[illegible]apacitance is adequate, [illegible]re is too much inductance in the path from the capacitors to the AP SoC.

Possible causes are:

- Wrong decoupling capacitor connecting-trace geometry or solder-land geometry
- The path from the capacitors to the AP SoC is too long

- and/or -

- A current path in the power vias traverses an exceptionally thick PCB stackup

For inadequate connecting trace geometry and capacitor land geometry, review the loop inductance of the current path. If the vias for a decoupling capacitor are spaced a few millimeters from the capacitor solder lands on the board, the current loop area is greater than necessary.

To reduce the current loop area, vias should be placed directly against capacitor solder lands. *Never* connect vias to the lands with a section of trace.

Other improvements of geometry are via-in-pad (via under the solder land), not shown, and via-beside-pad (vias straddle the lands instead of being placed at the ends of the lands). Double vias also improve connecting trace geometry and capacitor land geometry.

Exceptionally thick boards (> 3.2 mm or 127 mils) have vias with higher parasitic inductance.

[illegible] tic inductance, move critical V_{CC}/GND plane sandwiches close to the top surface [illegible] e AP SoC is located, and place the capacitors on the top surface where the AP SoC is l[illegible] d.

Possibility 3: I/[illegible] ignals in PCB are Stronger Than Necessary

If noise in the V_{CCO} PDS is still too high after refining the PDS [illegible] e slew rate and/or drive strength can be reduced. This applies to bo[illegible] tputs [illegible] SoC and inputs to the AP SoC. In severe cases, excessive overs[illegible] n inputs to th[illegible] SoC can reverse-bias the IOB clamp diodes, injecting curren[illegible] e V_{CCO} PDS.

If large amounts of noise are present on V_{CCO}, the [illegible] ength of these interfaces [illegible] be decreased, or different termination should be [illegible] input or output paths).

Possibility 4: I/O [illegible] gnal Return Current Tr[illegible] Sub-Optimal Paths

I/O signal retur[illegible] ents can [illegible] excessive [illegible] DS. For eve[illegible] gnal transmitted [illegible] ice into t[illegible] eventually [illegible] e is an equ[illegible] nt flowing [illegible] CB into the dev[illegible] system. If a [illegible] mpedance return current path is not available, a less optimal, higher impedance path is used. When I/O signal return currents flow over a less optimal path, voltage changes are induced in the PDS, and the signal can be corrupted by crosstalk. This can be improved by ensuring every signal has a closely spaced and fully intact return path.

Methods to correct a sub-optimal return current path:

- Restrict signals to fewer routing layers with verified continuous return current paths.
- Provide low-impedance paths for AC currents to travel between reference planes (high-frequency decoupling capacitors at PCB locations where layer transitions occur).

Se[illegible]IO Signaling

Introducti[illegible]

The Zynq-7000 AP SoC resources are the general-purpose I/O [illegible] ettings. With numerous I/O standards and hundreds of variants within [illegible] standa[illegible] SelectIO resources offer a flexible array of choices for designin[illegible] nterfaces.

This chapter provides some strategies for choosin[illegible] ndard, topography, and termination, and offers guidance on simulation a[illegible] urement for more detailed decision making and verification. In many cases, [illegible] vel aspects of the system [illegible] device choices or standards support) define the [illegible] ces to be used. In cases such constraints are[illegible] t defined, it is up to the s[illegible] gner to choose I/O inter[illegible] standards and opti[illegible] e them acc[illegible] ing to the p[illegible] system.

This chapter c[illegible] s the foll[illegible] ns:

- [illegible] ypes
- 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 GND common to the driver and receiver. In single-ended interfaces, a signal's assertion (whether it is High or Low) is based on its voltage level relative to a fixed voltage threshold that is referenced to GND. When the voltage of the signal is higher than the V_{IH} threshold, the state is considered High. When the voltage of the signal is lower than

the V_{IL} threshold, the state is considered Low. TTL is one common example of a single-ended I/O standard.

To reach higher interface speeds and increase noise margin, some single-ended I/O standards rely on a precise dedicated local reference voltage other than GND. HSTL and SSTL are examples of I/O standards that rely on a V_{REF} to resolve logic levels. V_{REF} can be [illegible] as a fixed comparator input.

Higher-per[illegible] interfaces typically make use of differential signaling – a convention that transmits t[illegible]mplementary signals referenced to one another. In differential interfaces, a sign[illegible]ssertion (whether it is High or Low) is based on the relative voltage levels of the two co[illegible]lementary 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 [illegible] the P and N signals have similar swing, and have a common-mode vo[illegible] above [illegible]ough this is not always the case). LVDS is one common example o[illegible]erential I/O s[illegible]ard.

SDR versus DDR Interfaces

The difference between Single Data Rate (SDR) a[illegible]e Data Rate (DDR) interfa[illegible] to do with the relationship of the data signals of a[illegible]e clock signal of that bus. [illegible] systems, data is onl[illegible]gistered at the input flip-[illegible]eceiving device on eithe[illegible] rising *or* the fallin[illegible]ge of the [illegible] One full clo[illegible]quivalent to one bit ti[illegible] DDR systems, d[illegible]registere[illegible]put flip-flo[illegible] device[illegible]oth th[illegible] rising *and* f[illegible]ges of the [illegible]full clock pe[illegible]bit times. Th[illegible]and DDR ha[illegible] to do with whet[illegible]ard carrying the signals is single-ended or differential. A single-ended interface can be SDR or DDR, and a differential interface can also be SDR or DDR.

Single-Ended Signaling

A variety of single-ended I/O standards are available in the Zynq-7000 AP SoC I/O. For a complete list of supported I/O standards and detailed information about each one, refer to the "SelectIO Resources" chapter of UG471, *7 Series FPGAs SelectIO Resources User Guide*. Tables at the end of this chapter summarize for each supported I/O standard which ones support DRIVE and SLEW attributes, bidirectional buffers, and the DCI options. It also describes which I/O standards are supported in the high-performance (HP) and high-range (HR) I/O banks.

Modes and Attributes

Some I/O standards can be used only in unidirectional mode, while some can be used in bidirectional mode or unidirectional mode.

Some I/O standards have attributes to control drive strength and slew rate, as well as the presence of weak pull-up or pull-down and weak-keeper circuits (not intended for use as parallel termination). Drive strength and slew rate can be used to tune an interface for adequate speed while not overdriving the signals. Weak pull-ups, weak pull-downs, and weak keepers can be used to ensure a known or steady level on a floating or 3-stated signal. The "SelectIO Resources" chapter of UG471, *7 Series FPGAs SelectIO Resources User Guide* [illegible] which standards support these attributes. Refer to this user guide for more information.

LVCMOS, when [illegible] 6 mA DRIVE and FAST slew, has an approximate output impedance close to 50Ω, allowing it to be used as a crude approximation of a controlled-impedance driver. The impedance match of the weak driver to the transmission line is only approximate and varies with voltage and temperature. LVDCI and HSLVDCI, true controlled-impedance drivers, are adaptive, maintain a much closer impedance mat[illegible] constant over voltage and temperature.

Input Thresholds

The input circuitry of the single-ended standards [illegible] two categories: those wit[illegible] input thresholds and those with input threshold[illegible] V$_{REF}$ voltage. The use o[illegible] has three advantages:

- It allows for ti[illegible] control [illegible] threshold [illegible]
- It remove[illegible]dence o[illegible] for the thres[illegible]
- [illegible] input threshol[illegible] closer together, w[illegible] need for a large voltage swing of the signal at the input receiver

Two 1.8V I/O standards that illustrate this are LVCMOS18 and SSTL18 Class 1. The thresholds for 1.8V LVCMOS are set at 0.63V and 1.17V (necessitating that the signal at the receiver swing a full 540 mV at minimum to make a logic transition). The thresholds for SSTL18 Class 1 are set at $V_{REF} - 0.125V$ and $V_{REF} + 0.125V$, or for a nominal V_{REF} of 0.9V, set at 0.775V and 1.025V (necessitating that the signal at the receiver only swing 250 mV at minimum to make a logic transition). This smaller required swing allows for higher frequency of operation in the overall link. A smaller swing at the driver means reduced DC power is required with less transient current. A historical drawback to the use of V_{REF} was that the semi-dedicated V_{REF} pins of the bank could not be used as I/Os whenever an I/O standard was used in a bank that required the V_{REF} supply. However, with the 7 series devices, the reference voltage can either be provided using the semi-dedicated V_{REF} pins, or optionally generated internally using the Internal V_{REF} feature. See UG471, *7 Series FPGAs SelectIO User Guide* for more details on Internal V_{REF}. For more information on V_{REF} decoupling and decoupling of all other supplies, see Chapter 3, Power Distribution System.

Topographies and Termination

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

[illegible] standards can be used in countless topographies depending on the [illegible] system. SelectIO drivers and receivers adhering to a standard (SSTL, LVCMOS, etc[illegible] can be used according to the letter of the standard (published by a standards body [illegible]as EIA/TIA or JEDEC) or they can be mixed and matched with drivers or receivers from an[illegible]r standard or hybrid I/O. An I/O standard specification might define something as limit[illegible]as the V_{IL} and V_{IH} of the receiver, or it might define every aspect of the interface, including driver impedance and slew rate, PCB trace length and topography, value and position of passive termination, the maximum input cap[illegible]iving device, and even the maximum number of receivers.

It is up to the designer to apply the standard in que[illegible] the system in which it is working. There are many decisions to make with r[illegible] topographies and termin[illegible] which determine the signal integrity of the interf[illegible]of utmost importance tha[illegible] signal integrity of each interface be verified thro[illegible]simulation and measure[illegible]

Termination general[illegible]efers to impedance-matc[illegible]edance-compensating [illegible] that are used to m[illegible]ain signal i[illegible]rity in an int[illegible] many types of elemen[illegible] be used as *term*[illegible]*s* (such a[illegible]s, capacito[illegible]is discussion[illegible]limite[illegible] resistive term[illegible]. In gene[illegible]r and diode[illegible]are more [illegible]

Unidirectional Topographies and Termination

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

Unidirectional Point-to-Point Topographies

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

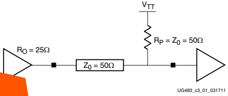

Figure 4-1: [illegible]el-Terminated Unidirectional, Point-to-Point Topography

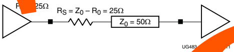

Figure 4-2: **Series-Terminated Unidirectional, Poi[illegible]int Topography**

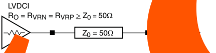

Figure 4-3: [illegible]ntrolled [illegible]Driver Unid[illegible]to-Point [illegible]phy

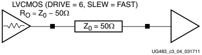

Figure 4-4: **"Weak Driver" Unidirectional, Point-to-Point Topography**

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

Assuming transmission lines with 50Ω characteristic impedance and a driver output impedance (R_O) of 25Ω, a 25Ω series termination (Figure 4-2) or a 50Ω parallel termination (Figure 4-1) is appropriate. Controlled-impedance drivers, whether implemented with DCI or with weak LVCMOS drivers, should be sized to have an output impedance (R_O) of 50Ω. This corresponds to VRN and VRP resistors equal to 50Ω for DCI. Weak LVCMOS drivers of 6 mA to 8 mA drive strength have an output impedance approximately equal to 50Ω (Figure 4-3).

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 possible to use a Thevenin parallel termination. Thevenin parallel termination consists of a voltage divider with a parallel equivalent resistance (R_{PEQ}) equal to the characteristic impedance of the transmission line (50Ω in most cases). The divided voltage point is [illegible] at V_{TT}. Figure 4-5 illustrates a Thevenin parallel termination powered from 2.5V V_{CCO} [illegible] of two 100Ω resistors, resulting in a V_{TT} of 1.25V and a parallel equivalent res[illegible] (R_{PEQ}) of 50Ω.

Parallel terminati[illegible]an be less desirable than series termination or controlled-impedance drivers because it [illegible]ipates more power. This trade-off must be weighed against other trade-offs to determine the optimum termination topography for an interface.

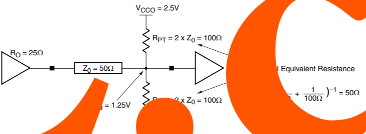

[illegible]igure 4-5: **[illegible]arallel Terminat[illegible]**

Table 4-1 lists example I/O interface types that can be used with the unidirectional point-to-point topography.

Table 4-1: **Example I/O Interface Type for Unidirectional Point-to-Point Topographies**

LVTTL
LVCMOS
LVDCI
SSTL Class I
HSTL Class I

LVTTL and LVCMOS do not specify any canonical termination method. Series termination at the driver or parallel termination at the receiver are both appropriate considerations.

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

Every I/O standard can have different requirements for termination techniques. In some cases the specification for the I/O standard can rigidly define the termination topology. Other standards might not have any hard requirements, but rather might simply provide

examples of termination topologies. An example of a standard with specific termination requirements is HSTL. HSTL Class I is a unidirectional I/O standard that recommends a parallel termination at the receiver. In the case of HSTL Class I, the termination voltage V_{TT} is defined as half of the supply voltage V_{CC}. The designer can ultimately elect either not to use termination at all or to use a different termination, such as series termination at the driver. There are a number of reasons why this selection might be advantageous in a given [illegible] to the designer to verify through simulation and measurement that the signal [illegible] receiver is adequate.

The SSTL [illegible] to not have rigid requirements for termination topology. Rather, the JEDEC specifications [illegible] example termination techniques that tend to be the commonly used topographies. The “[illegible] Resources” chapter of UG471, *7 Series FPGAs SelectIO Resources User Guide* provides example termination techniques for each of the I/O standards, including the SSTL standards, for the purpose of providing a good starting point for [illegible] to HSTL, it is ultimately up to the designer to verify through simulation and [illegible] integrity at the receiver is adequate.

Unidirectional Multi-Drop Topographies

In more complex topographies, a single driver [illegible] multiple receivers. The [illegible] represent loads that must be fed by individual [illegible] line stubs. From a [illegible] integrity standpoint, the best topography to use [illegible] is a single long [illegible] line with the driver [illegible] one end and parallel [illegible] other, with receivers connected to the [illegible] trace by [illegible] stubs in [illegible] of topography is [illegible] referred to as [illegible] *multi-[illegible]*phy.

[illegible] aspects [illegible]graphy. The first [illegible] a single parallel termination at the far end of the transmission line. Series termination at the driver or a controlled impedance driver must *never* be used. Parallel termination is the only applicable termination type for this topography. The second critical aspect is the length of the connecting stubs at each receiver. These must remain short: no more than a fraction of a signal rise time in length. With a typical signal rise time of 600 ps, a stub no longer than 700 ps/4 = 150 ps, or 0.9 inches (22.86 mm) should be used. As the stubs become longer, they present a larger impedance discontinuity to the signal travelling down the transmission line, and can support significant reflections. These impedance discontinuities corrupt the signal. With increasing numbers of loads and increasing length of stubs, the signal is corrupted to the point where it is no longer usable.

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

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

Figure 4-6 illustrates a Thevenin parallel termination powered from V_{CCO}, made up of two 100Ω resistors, resulting in a V_{TT} of $V_{CCO}/2$ and a parallel equivalent resistance of 50Ω. This figure shows a topography with one driver (an LVCMOS driver) and four receivers. The driver is on the left side, the receivers are spaced at interim points across the 50Ω transmission line, and the Thevenin parallel termination of two 100Ω resistors is on the right side.

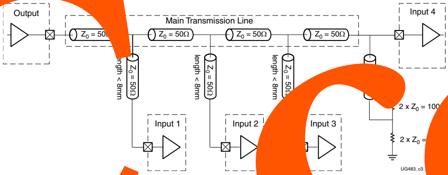

Figure 4-6: **Basic Multi-Drop [illegible]**

The main tran[illegible]n line sh[illegible]t as short [illegible] the [illegible]0 inches [illegible] more are [illegible]or most I/[illegible] as long as p[illegible] is maintained [illegible]urces are avoid[illegible] lengths of interim s[illegible]the main transmission line need not be equal. Their relative lengths can be arbitrary. Receivers at different points along the main transmission line receive the signal with varying amounts of delay, but all signal rise times are similar.

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

Table 4-2 lists example I/O interface types that can be used with the unidirectional multi-drop topography.

Table 4-2: **Example I/O Interface Types for Unidirectional Multi-Drop I/O Topographies**

LVTTL
LVCMOS
HSTL
SSTL

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

Bidirectional Topography and Termination

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

Bidirectional [illegible]-Point Topographies

The simplest [illegible] topography is point to point. That is, there are two transceivers connected by a [illegible] line. Because bidirectional interfaces need to operate equally well in both directions, symmetry of the topography is desirable. While asymmetrical topographies can be designed with reasonably good signal [illegible] way to ensure good signal integrity is to keep the topography [illegible] termination used on one side of the link should also be used on [illegible] side of the [illegible] Series termination (Figure 4-8) is rarely appropriate for [illegible] interfaces as incoming signals are attenuated by the series resistor of the [illegible] transceiver. Parallel [illegible] (Figure 4-7) almost always achieves better signal [illegible] both receivers.

Controlled-impedance drivers, whether crudely [illegible] in the form of a weak [illegible] driver or adaptively controlled in the form LVDCI [illegible], also can have good [illegible] shown in Figure [illegible] Figure 4-10, and Figure [illegible] with a low-drive [illegible] LVCMOS driver). [illegible] use IBIS [illegible] to [illegible] optimal termination [illegible] value, V_{TT} [illegible] and [illegible] these [illegible].

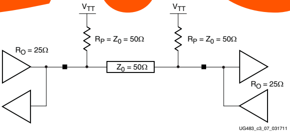

Figure 4-7: **Parallel Terminated Bidirectional Point-to-Point Topography**

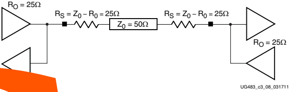

Figure 4-8: [illegible] **Terminated Bidirectional Point-to-Point Topography: Not Recommended**

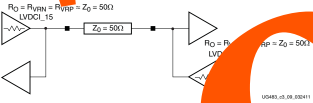

Figure 4-9: **DCI Controlled Impedance Bidirectio** [illegible] **-Point Topography**

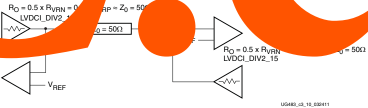

Figure 4-10: **HSLVDCI Controlled Impedance Driver Bidirectional Point-to-Point Topography**

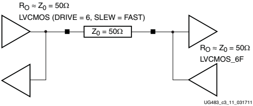

Figure 4-11: **"Weak Driver" Bidirectional Point-to-Point Topography**

In general, parallel resistive termination (R_P) has a value equal to the characteristic impedance Z_0 of the transmission line it is terminating. Some interfaces, such as DDR2

memory interfaces, use 75Ω termination resistors instead of 50Ω in an effort to open the data eye. In this case, the trade-off is eye height against a small amount of signal reflection from the impedance discontinuity. Controlled-impedance drivers are typically tuned such that the driver output impedance (R_O) is equal to the characteristic impedance (Z_0) of the transmission line it is terminating.

[illegible] transmission lines with 50Ω characteristic impedance and a driver output [illegible] 50Ω parallel terminations are appropriate (Figure 4-7). Controlled-[illegible]ce drivers, whether implemented with DCI or with weak LVCMOS drivers, should [illegible]ed to have an output impedance (R_O) of 50Ω. An example of the use of a controlled-in[illegible]dance driver would be the LVDCI_15 I/O standard. By using 50Ω external precision [illegible]stors placed on the VRN and VRP pins for that bank, the resulting controlled output impedance for that bank would be 50Ω. If 100Ω resistors were already required on the VRN and VRP pins (for the purpose of creati[illegible]uivalent split termination circuit equal to 50Ω), and a 50Ω controlled [illegible]dance dr[illegible]required in the same bank, this could be accomplished by using th[illegible]/2" versions o[illegible] drivers, s[illegible] as LVDCI_DIV2_15 (Figure 4-9 and Figure 4-10). We[illegible]MOS drivers of 6 mA to 8 m[illegible] drive strength have an output impedance approxi[illegible]qual to 50Ω (Figure 4-11).

Parallel terminations have the best performance [illegible] (the voltage source conne[illegible] the parallel termination resistor) is equal to half [illegible]aling voltage, since this [illegible] typically the center[illegible]tage of the data eye. For [illegible] (V_{CCO} = 2.5V), V_{TT} is id[illegible] 1.25V. In cases wh[illegible]this voltag[illegible]t available, [illegible]le to use a Thevenin pa[illegible] termination. Th[illegible] parallel [illegible]n consists [illegible]ivider with [illegible]arallel resistance e[illegible]he charac[illegible]edance of t[illegible] in most [illegible] voltage po[illegible]gned to be at V_{TT}. [illegible]trates a Thevenin parallel termination powered from 2.5V V_{CCO}, made up of two 100Ω resistors, resulting in a V_{TT} of 1.25V and a parallel equivalent resistance (R_{PEQ}) of 50Ω.

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

Figure 4-12: **Thevenin Parallel Termination (Bidirectional Point-to-Point Topography)**

Table 4-3 lists example I/O interface types that can be used with the bidirectional point-to-point topography.

Table 4-3: **Example I/O Interface Types for Bidirectional Point-to-Point I/O Topographies**

LVTTL
LVCMOS
LVDCI
HSLVDCI
SSTL15
SSTL15 DCI
SSTL18 CLASS II
SSTL18 CLASS II DCI
HSTL CLASS II
HSTL CLASS II DCI

LVTTL and LVCMOS do not specify any canonical termination method. Series termination [illegible] not recommended for bidirectional interfaces. Parallel termination and weak drivers, however, are both appropriate.

LVDCI and HSLVDCI both implicitly [illegible] controlled [illegible] driver termination.

HSTL Class II specifies parallel [illegible] at both transmitter [illegible] voltage V_{TT} is defined [illegible] the supply [illegible] V_{CCO}. The designer [illegible] not to use termination at all or to use a different termination. It is up to the designer to verify through simulation and measurement that the signal integrity at the receiver is adequate.

The JEDEC specifications for SSTL provide examples of both series termination and parallel termination. The termination voltage V_{TT} is defined as half of the supply voltage V_{CCO}. While the specification document provides examples depicting series termination at the drivers, it is important to note that the purpose of this is to attempt to match the impedance of the driver with that of the transmission line. Because the Zynq-7000 AP SoC SSTL drivers target to have output impedances close to 40–50Ω, better signal integrity can be achieved without any external source-series termination. When possible, it is a better starting point to consider the use of the 3-state DCI I/O standards ("T_DCI"), which provide internal parallel termination resistors that are only present when the output buffer is in 3-state. It is up to the designer to carefully choose the I/O standard(s) at the 7 series device, drive strengths, and on-die termination (ODT) options at the other device(s) in the interface (usually DRAM ICs) and termination topography though careful simulation and measurement. See UG471, *7 Series FPGAs SelectIO User Guide* for more details on the available I/O standards and options.

Bidirectional Multi-Point Topographies

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

Chapter 5

Pr[illegible]ing System (PS) Power and Signaling

Power

Zynq-7000 AP SoC devices are divided into several po[illegible]omains. Figu[illegible] shows a[illegible] overview of those domains.

Figure 5-1: **Power Domains**

Main Power Supplies

V_{CCPINT} – PS Internal Logic Supply

V_{CCPINT} is a 1.0V nominal supply that powers all of the PS internal logic circuits.

V_{CCPAUX} – PS Auxiliary Logic Supply

V_{CCPAUX} is a 1.8V nominal supply that powers the PS I/O banks.

V_{CCPLL} – PS PLL Supply

V_{CCPLL} is a 1.8V nominal supply that provides power to the three PS PLLs. It can be powered separately or tied to the V_{CCPAUX} supply. If powered by V_{CCPAUX}, V_{CCPLL} must be filtered through a 120Ω @ 100 MHz, size 0602 ferrite bead and a 10 μF, size 0602 decoupling capacitor to minimize PLL jitter. Figure 5-2 shows an example of the filtering circuit for [illegible]mended ferrite bead is Murata BLM18SG121TN1. Recommended capacitor is Murata [illegible]G106ME47.

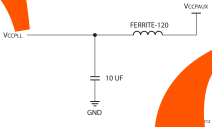

Figure 5-2: **Connecting V[illegible]**

PS DDR Pow[illegible]Supplie[illegible]

V[illegible]DR I/O Su[illegible]

V_{CCO_DDR} is a 1.2V–1.8V nominal supply that supplies the DDR I/O bank input and output drivers. This supply sources the DDR output drivers, input receivers and termination circuitry. Its requirements are defined by the type of interface (DDR2, DDR3 or LPDDR2), memory speed, and the data bus width. Table 5-1 shows the supply voltages for the different memory types.

Table 5-1: **PS DDR Interface I/O Supply**

DDR interface	DDR2	DDR3	LPDDR2
Voltage	1.8V	1.5V	1.2V

PS_DDR_VREF0, PS_DDR_VREF1 – PS DDR Reference Voltage

PS_DDR_VREF0 and PS_DDR_VREF1 provide a voltage reference for the PS_DDR_DQ and PS_DDR_DQS differential input receivers. They need to be tied to a termination voltage (V_{tt}) equal to $V_{CCO\_DDR}/2$. For example, for DDR3, V_{CCO_DDR} is set to 1.5V, then V_{REF} shall be set to 0.75V. A resistor divider can be used to generate PS_DDR_VREF0 and PS_DDR_VREF1. A 0.01 uF capacitor shall be added for decoupling.

PS_DDR_VRN, PS_DDR_VRP – PS DDR Termination Voltage

PS_DDR_VRN and PS_DDR_VRP provide a reference for dynamic controlled impedance (DCI) calibration. For memory types that require termination (DDR2, DDR3) VRP must be pulled Low to GND and VRN needs to be pulled High to V_{CCO_DDR}. The resistor value on VRP and VRN should be twice the memory's trace and termination impedance. For example, for a [illegible] with 40Ω termination and board impedance, an 80Ω resistor must be used to pull-up/[illegible]nd VRN.

Table 5-2 shows [illegible]equired values for the DCI VRN/VRP pull-down/pull-up resistors.

Table 5-2: **DCI VRP [illegible]d VRN Values**

	LPDDR2	DDR2	DDR3
VRP/VRN	40Ω	100Ω	[illegible]
	(type I DCI trace impedance 40Ω)	(type III DCI trace impedance 50Ω)	([illegible] III DCI trace [illegible]edance 40Ω)

Unused DDR Memory

When no PS DDR memory is used, V_{CCO_DDR} sho[illegible] to V_{CCPAUX}. PS_DDR_VR[illegible] and PS_DDR_VRN/P[illegible] be left floating, though i[illegible]ended to tie them to gro[illegible] this instance.

PS MIO [illegible] Suppl[illegible]

V_{CCO_MIO0} – PS MIO Bank 0 I/O Supply

V_{CCO_MIO0} is a 1.8–3.3 volt supply. It powers I/O Bank 500 which contains PS_MIO[15:0], PS_CLK, and PS_POR_B I/Os.

V_{CCO_MIO1} – PS MIO Bank 1 I/O Supply

V_{CCO_MIO1} is a 1.8-3.3 volt supply. It powers I/O Bank 501 which contains PS_MIO[53:16], PS_VREF_MIO, and PS_SRST_B I/Os.

Configuring the V_{CCO_MIO0}, V_{CCO_MIO1} Voltage Mode

The PS I/O banks can operate in two different voltage modes, low (1.8V) mode and high (2.5V – 3.3V) mode. Before powering on, the banks must be configured for the correct mode otherwise, damage might occur. The I/O bank voltage is set by pulling pins MIO[7] and MIO[8] either High or Low. Table 5-3 shows the voltage mode configuration (VMODE) for MIO Bank 0 and Bank 1.

MIO[7] and MIO[8] are dual use pins that are shared with the high-speed QSPI/NAND/SRAM interface signals. Special care needs to be taken to avoid signal integrity issues.

CAUTION! *If the MIO bank voltage is incorrectly set, the I/O behaves unpredictably and damage might occur. For example, avoid setting the MIO voltage to 3.3V while using HSTL18.*

[illegible] Vo[illegible] [illegible]figuration

I/O Bank	[illegible]ply Name	Programming Pin	1.8V Mode	2.5V, 3.3V Modes
MIO Bank 0	[illegible]_MIO0	MIO[7]	20 KΩ resistor to V_{CCO_MIO0}	20 KΩ resistor to ground
MIO Bank 1	V_{CCO_MIO1}	MIO[8]	20 KΩ resistor to V_{CCO_MIO}[illegible]	20 KΩ resistor to ground

PS_MIO_VREF – RGMII Reference Voltage

PS_MIO_VREF provides a reference voltage for the [illegible] differential input receivers. [illegible] needs to be tied to a voltage equal to one half V[illegible] For example, when using [illegible] HSTL18 RGMII interface, V_{CCO_MIO1} is set to 1.8V [illegible] MIO_VREF shall be set to [illegible] A resistor divider can be used to generate PS_M[illegible] .01 uF capacitor shall be [illegible] for decoupling.

Power Seq[illegible]cing

Ref[illegible] *[illegible]nq-7000 EP[illegible] 0 and XC7Z020) [illegible]ing [illegible]cteristics* for power supply sequencing requirements.

Power Supply Ramp Requirements

Refer to DS187, *Zynq-7000 EPP (XC7Z010 and XC7Z020) DC and AC Switching Characteristics* for for power supply ramp requirements.

PCB Decoupling Capacitors

Decoupling guidelines for the PS supplies can be found in Table 3-2, page 12.

PS Clock and Reset

PS_CLK – Processor Clock

PS_CLK shall be connected to a clock generator providing a 30-60 MHz clock. The clock must be a single-ended LVCMOS signal, using the same voltage level as the V_{CCO_MIO0} I/O

voltage for bank MIO0. Refer to DS187, *Zynq-7000 EPP (XC7Z010 and XC7Z020) DC and AC Switching Characteristics* for further PS_CLK requirements.

PS_POR_B – Power on Reset

The PS power-on reset (PS_POR_B) is an active-Low signal used to hold the PS in reset until [illegible]lies are stable and at their required voltage levels. It is suggested to generate [illegible] from the power supply *power-good* signal or a voltage supervisor chip. At the time PS_[illegible]B is released, the system clock (PS_CLK) must have been stable for 2,000 clocks. PS_[illegible]_B should be pulled High to V_{CCO_MIO0}. When asserting PS_POR_B, the pulse length must [illegible] longer than 1 µs.

PS_SRST_B – External System Reset

The PS system reset (PS_SRST_B) is an active-Low signal that [illegible]tly used for d[illegible]ng propos[illegible] PS_SRST_B must be High to begin the boot process. If PS_[illegible]s not used it can be pulled Hi[illegible] V_{CCO_MIO1}. When asserting PS_SRST_B, the pulse length [illegible] longer than 1 µs.

Boot Mode Pins

Pins MIO[2] to M[illegible]] are use[illegible]gure the b[illegible] bypass, and [illegible]O vol[illegible] All designs m[illegible]ude 20 K[illegible]r pull-dow[illegible] set the requir[illegible]

MIO[2] and MIO[8] are dual use pins that are shared with the high-speed QSPI/NAND/SRAM interface signals. Special care needs to be taken to avoid signal integrity issues. To avoid signal integrity issues, limit the stub length to the pull-up or pull-down resistors to < 10 mm.

When system design requires the modes to be changeable, it is recommended to not use a resistor tree to set the mode but instead connect one pull-up/down resistor to the mode pin and place a jumper on the other side of the resistor to select between pull-up or pull-down. See Figure 5-3 for an example.

Note: PROGRAM_B, INIT_B, and DONE should not be left floating. Refer to UG470, *7 Series FPGAs Configuration User Guide* for more information on how to treat these pins.

Note: The PL system JTAG interface, PL_JTAG, should have its signals TDI, TMS, and TCK pulled-up.

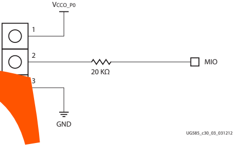

Figure 5-3: **Setting Mode Pins**

Dynamic Memory

Zynq-7000 AP SoC devices support DDR2, DDR3[illegible]DR2 (mobile DDR) dynam[illegible] memory. The memory is connected to dedicate[illegible]O Bank 502. This bank ha[illegible] dedicated I/O, termi[illegible]tion, and reference voltag[illegible]

DDR runs at very [illegible] speeds a[illegible]ial care ne[illegible]n in board layout to e[illegible] signal integrity[illegible]ollowing[illegible]ow the re[illegible] for DDR[illegible]mory designs for [illegible]00 AP So[illegible]

DDR Interface Signal Pins

Table 5-4 lists all dynamic memory interface signals in Bank 502.

Table 5-4: **DDR Interface Signal Pins**

Pin Name	Direction	Description
DDR_CK_P	O	Differential clock output positive
DDR_CK_N	O	Differential clock output negative
DDR_CKE	O	Clock enable
DDR_CS_B	O	Clock select
DDR_RAS_B	O	RAS row address select
DDR_CAS_B	O	CAS column address select
DDR_WE_B	O	Write enable
DDR_BA[2:0]	O	Bank address
DDR_A[14:0]	O	Address
DDR_ODT	O	Output dynamic termination
DDR_DRST_B	O	Reset

Table 5-4: **DDR Interface Signal Pins *(Cont'd)***

Pin Name	Direction	Description
DDR_DQ[31:0]	I/O	Data
DDR_DM[3:0]	O	Data mask
DDR_DQS_P[3:0]	I/O	Differential data strobe positive
[illegible]	I/O	Differential data strobe negative
DDR_VRP	I/O	Used to calibrate input termination
DDR_VRN	I/O	Used to calibrate input termination
DDR_VREF[1:0]	I/O	Reference voltage

Unused DDR pins should be connected as shown in Table 5-5.

Table 5-5: **DDR Unused Pins**

Unconnected Pins	Comments
O	Unconnected
IO	Unconnected, internal pull-up by s[illegible]

Dynamic Memory Implementation

Figure 5-4, Figure [illegible] and Figure 5-6 show exam[illegible]menting DDR memory [illegible] typical boards.

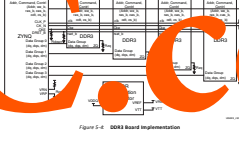

Figure 5-4: **DDR3 Board Implementation**

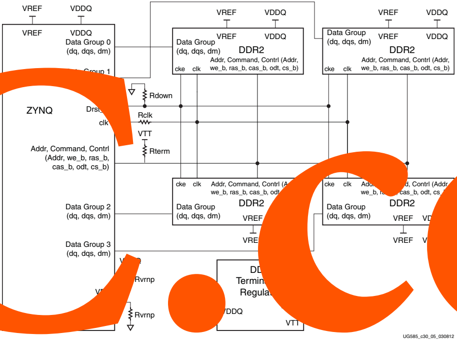

Figure 5-5: **DDR2 Board Implementation**

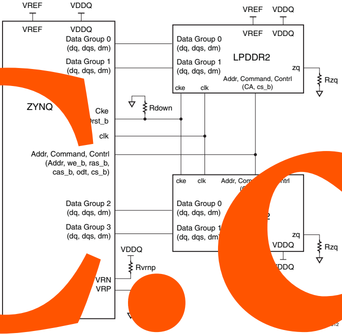

Figure 5-6: **LPDDR2 Board Implementation**

DDR Supply Voltages

Table 5-6 lists the different supply, reference and termination voltages required for LPDDR2/DDR2/DDR3 memory. These voltages are also required to power the DDR I/O bank, reference, and termination voltages.

Note: Xilinx recommends that each DDR memory IC should have an independent V_{REF} voltage source and divider instead of sharing a common one.

Table 5-6: **DDR Voltage**

Voltage	LPDDR2	DDR2	DDR3	Comments
V_{CCO_DDR} V_{DDQ}	1.2V	1.8V	1.5V	LPDDR2 devices, also require V_{DD1} (1.8V) and V_{DD2} (1.2V)
V_{TT}	$V_{DDQ}/2$	$V_{DDQ}/2$	$V_{DDQ}/2$	Use a DDR termination regulator or a resistor voltage divider to generate V_{TT} and V_{REF}
PS_DDR_V_{REF0} PD_DDR_V_{REF1} V_{REF}	$V_{DDQ}/2$	$V_{DDQ}/2$	$V_{DDQ}/2$	

DDR Termination

For better signal integrity, DDR2 and DDR3 clock, address, command and control signals need to be terminated. DDR_CKE and DDR_DRST_B do not require termination but need to be pulled to GND with a 4.7 KΩ resistor.

[illegible] not require termination.

Table 5-7 sh[illegible] DDR termination requirements.

Table 5-7: **DDR Terminat**[illegible]

Termination	LPDDR2	DDR2	DDR3	Comments
Rterm	N/A	50Ω	40Ω	
Rclk	N/A	100Ω	80Ω	
Rdown	4.7KΩ	4.7 KΩ	4.7 KΩ	There [illegible] DR_DRST_B in L[illegible]/DDR2 devi[illegible] sid[illegible]

Note: DDR3 memory also supports terminated DQS s[illegible]ough the TDQS_P and TDQS_[illegible] This feature is not supported on Zynq-7000 AP SoC [illegible]d those pins should be left f[illegible]

DDR Trace Len[illegible]th

All DDR memory [illegible]ces shoul[illegible]ed as clos[illegible]q-7000 AP SoC device[illegible] possible. Tabl[illegible]hows the [illegible] recommen[illegible]signals. [illegible] trace legt[illegible]so include [illegible]ge delay.

Table 5-8: **DDR Max Trace Length**

Signal Group	LPDDR2	DDR2	DDR3	Comments
Data Group	1.5″	2″	2″	Without guard-band
		5″	5″	With guard-band
Address, Command, Control	1.5″	3″	3″	Without guard-band
		6″	6″	With guard-band

In addition, DDR signals also require matched trace lengths. Table 5-9 shows the recommended trace length matching for DDR. Differential traces should be length matched to ±5 mil if possible.

Table 5-9: **DDR Trace Length Match**

Signal Group	LPDDR2	DDR2	DDR3	Comments
DQ/DM to DQS_P/N in data group	±5 ps	±10 ps	±5 ps	
Address/Control to CK_P/N	±10 ps	±25 ps	±10 ps	

DDR Trace Impedance

All DDR signals except DDR_DRST_B and DDR_CKE require controlled impedance. Table 5-10 shows the required trace impedance for DDR signals.

Table 5-10: **DDR Trace Impedance**

[illegible]	LPDDR2	DDR2	DDR3	Comments
Single-ended	40Ω	50Ω	40Ω	±10% tolerance
Differential	80Ω	100Ω	80Ω	±10% tolerance

DDR3 and LPDDR2 memory also require an additional resistor connected to the ZQ pin to calibrate the device's output impedance. Table 5-11 shows the required RZQ values.

Table 5-11: **DDR ZQ**

ZQ	LPDDR2	DDR2	DDR3
Rzq	240Ω	N/A	240Ω

DDR Routing Topology

Based on the chosen memory type, the number [illegible] devices and layout requirements, different routing topologies can be [illegible] DDR memory. Figure 5-7 [illegible] three different topologies.

Note: For PS [illegible] Qxx, ensure [illegible] es are kept to [illegible] ould always [illegible] omitted for [illegible] or other function[illegible] (PS_ADDR14)

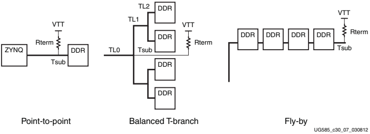

Figure 5-7: **DDR Routing Topologies**

In a balanced T-branch configuration, trace lengths TL1, TL2 and Tsub shall be kept as short as possible; Rterm shall be close to and balanced among the loads.

In a point-to-point configuration, Rterm shall be placed behind and close to the last load.

Table 5-12 shows the recommended routing topologies. DQ data pins within a byte lane are allowed to be swapped.

Table 5-12: **DDR Routing Topology**

Signal Group	LPDDR2	DDR2	DDR3	Number of DDR Devices
[illegible]	Point-to-point	Point-to-point	Point-to-point	
Clock	Point-to-point	Point-to-point	Point-to-point	1
	T-branch	T-branch	Fly-by/T-branch	2
	N/A	T-branch	Fly-by	4
Address, Command, Control	Point-to-point	Point-to-point	Point-to-point	1
	T-branch	T-branch	Fly-by/T-branch	2
	N/A	T-branch	Fly-by	4

SD/SDIO Peripheral Controller

The SD/SDIO peripheral controller controls communication with SDIO devices and SD memory cards connected to the PS.

SDIO Interface Guidelines

[illegible] (SDx_CLK) needs to be terminated with a [illegible]s resistor. The resistor shall be placed as close to the MIO pin as possible.

MIO IP Layout Guidelines

This section lists MIO interface-specific layout guidelines.

QSPI

The trace delay shall be between 0.7 ns – 0.8 ns. If the QSPI operating frequency is intended to be greater than FQSPICLK2 (as defined by DS187, *Zynq-7000 EPP (XC7Z010 and XC7Z020) DC and AC Switching Characteristics*), no loads may be placed on MIO[8] other than a pull-up or pull-down resistor.

Note: When operating at a QSPI frequency greater than FQSPICLK2, quad-SPI external loopback must be enabled. QSPIx_CLK should not have pull-up/pull-down if on a shared MODE pin. For FQSPICLK1 and FQSPICLK2 operating requirements, refer to DS187, *Zynq-7000 EPP (XC7Z010 and XC7Z020) DC and AC Switching Characteristics*.

Note: MIO[7] and MIO[8] are dual use pins that are shared with the high-speed QSPI/NAND/SRAM interface signals. Special care needs to be taken to avoid signal integrity issues.

USB0/1

If the MIO bank is configured to 3.3V, the maximum board delay shall be 0.8 ns. For other bank voltages, the maximum board delay shall be 1.3 ns.

SMC

If the MI[illegible] onfigured to 3.3V, the maximum board delay shall be 1 ns. Otherwise, the maximum [illegible] delay shall be 1.3 ns.

SPI 0/1 Master [illegible]de

If the MIO bank is configured to 3.3V, the maximum board delay [illegible] ns. Otherwise, the maximum board delay shall be 1.3 ns.

Ethernet MIO GEM0

RX Path: For HSTL 1.8V, the clock to data skew sha[illegible]ween 1.5 ns (min) and 2 ns [illegible]

TX Path: For HSTL 1.8V, the clock to data skew s[illegible]ween 1.9 ns (min) and 2.[illegible] (max).

Ethernet MIO [illegible]1 (MIO[illegible]39)

RX Path: F[illegible]8V, the clo[illegible]skew shall be [illegible]and 2 ns (max).

TX Path: For HSTL 1.8V, the clock to data skew shall be between 1.7 ns (min) and 2.2 ns (max).

Trace A (MIO 2-19)

TX Path: The clock to data skew shall be 2.8 ns to center align the clock to the data for all voltages.

Trace B (MIO 22- 27)

TX Path: The clock to data skew shall be 2.6 ns to center align the clock to the data for all voltages.

Temperature Sensing Diodes

Note: If unused, the DXP and DXN pins for the temperature sensing interface should be tied to GND.

Appendix A

Additional Resources

Xilinx Resources

Product Support and Documentation

www.xilinx.com/support

Xilinx Glossary

www.xilinx.com/company/terms.htm

Device User Gui[illegible]

http://www[illegible].com/sup[illegible]entation/[illegible]

[illegible]x.com/supp[illegible]mentation/7_seri[illegible]htm

Xilinx Design Tools: Installation and Licensing Guide **(UG798)**

http://www.xilinx.com/support/documentation/sw_manuals/xilinx14_1/iil.pdf

Xilinx Design Tools: Release Notes Guide **(UG631)**

http://www.xilinx.com/support/documentation/sw_manuals/xilinx14_1/irn.pdf

Xilinx Forums and Wiki Links

- http://forums.xilinx.com
- http://wiki.xilinx.com
- http://wiki.xilinx.com/zynq-linux
- http://wiki.xilinx.com/zynq-uboot

Xilinx git Websites

http://git.xilinx.com

Solution Centers

See the Xilinx Solution Centers for support on devices, software tools, and intellectual property at all stages of the design cycle. Topics include design assistance, advisories, and troubleshooting tips.

Reference

Zynq-7000 AP SoC Documents

Refer to the following Zynq-7000 AP SoC documents f[illegible]her referen[illegible]

- CS692, *Zynq-7000 All Programmable SoC Pr*[illegible]*rief*
- DS190, *Zynq-7000 All Programmable SoC* [illegible]
- DS187, *Zynq-7000 All Programmable SoC* [illegible]*t*
- UG585, *Zyn*[illegible]*00 All Programmable SoC* [illegible]*eference Manual*
- UG865, *Z*[illegible]*000 All P*[illegible]*able SoC* [illegible] *Pinout Specifi*[illegible]*ations*
- UG82[illegible]*-7000 All* [illegible]*ble SoC Sof*[illegible]
- [illegible]*ynq-7000 All Pr*[illegible]*able SoC XC7Z020* [illegible]

These documents can be found on the Zynq-7000 documentation page at:

http://www.xilinx.com/support/documentation/zynq-7000.htm

PL Documents – Device and Boards

To learn more about the PL resources, refer to the following 7 Series FPGA User Guides:

- DS821, *Xilinx LogiCORE IP 7 Series FPGAs Integrated Block for PCI Express Product Specification*
- UG471, *Xilinx 7 Series FPGAs SelectIO Resources User Guide*
- UG472, *Xilinx 7 Series FPGAs Clocking Resources User Guide*
- UG473, *Xilinx 7 Series FPGAs Memory Resources User Guide*
- UG474, *Xilinx 7 Series FPGAs Configurable Logic Block User Guide*
- UG476, *Xilinx 7 Series FPGAs GTX Transceiver User Guide*
- UG477, *Xilinx 7 Series FPGAs Integrated Block v1.3 for PCI Express User Guide*

- UG479, *Xilinx 7 Series FPGAs DSP48E1 User Guide*
- UG480, *Xilinx 7 Series FPGAs XADC User Guide*
- UG483, *Xilinx 7 Series FPGAs PCB and Pin Planning Guide*

These user guides and additional relevant information can be found in the 7 Series overall of documentation on the 7 Series FPGAs, which is available on the Xilinx website:

http://w[illegible].com/support/documentation/7_series.htm

Advanced eXtensible Interface (AXI) Documents

Refer to UG761, *AXI Reference Guide* for further reference on the AXI protocol.

Software Documents

***Zynq-7000 EPP Software Developers Guide* (UG821)**

http://www.xilinx.com/support/documentatio[illegible]000_user_guides.htm

***Zynq-7000 EPP Concepts, Tools, and Techniques Gu*[illegible]3)**

http://www.xil[illegible]om/suppo[illegible]umentatio[illegible]s/xilinx14_x/ug873_zyn[illegible]pdf

The so[illegible]or stand a[illegible]SBL are provide[illegible]nx IDE Design [illegible]dded Edition. The Lin[illegible]vers are provided via th[illegible]Open Source Wiki at:

http://wiki.xilinx.com

Xilinx Alliance Program partners provide system software solutions for IP, middleware, operation systems, etc. For the latest information refer to the Zynq-7000 landing page at:

http://www.xilinx.com/products/silicon-devices/epp/zynq-7000

git Information

- http://git-scm.com
- http://git-scm.com/documentation
- http://git-scm.com/download

Design Tool Documents

Xilinx Command Line Tools User Guide **(**[UG628](UG628)**)**

Xilinx ISE Manuals

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

Xilinx XPS/E[illegible]rted IP Website

http://www.xi[illegible].com/ise/embedded/edk_ip.htm

ChipScope Pro Software and Cores User Guide **(**UG029**)**

PlanAhead Tutorial: Debugging with ChipScope **(**UG677**)**

Xilinx Problem Solvers

http://www.xilinx.com/support/troubleshoot.h[illegible]

EDK Documentation

The entire EDK documentation set can be access[illegible]

http://www[illegible].com/sup[illegible]nentation/[illegible]1.htm

- [illegible]ystem Tools [illegible] Manual (UG11[illegible]

  http://www.xilinx.com/support/documentation/xilinx14_1/est_rm.pdf

- *OS and Libraries Document Collection* (UG643):

  http://www.xilinx.com/support/documentation/xilinx14_1/oslib_rm.pdf

- *Platform Specification Format Reference Manual* (UG642):

  http://www.xilinx.com/support/documentation/xilinx14_1/psf_rm.pdf

EDK Tutorials Website

http://www.xilinx.com/support/documentation/dt_edk_edk14-1_tutorials.htm

Platform Studio and EDK Website

http://www.xilinx.com/ise/embedded_design_prod/platform_studio.htm

XPS/EDK Supported IP Website

http://www.xilinx.com/ise/embedded/edk_ip.htm

Third-Party IP and Standards Documents

To learn about functional details related to vendor IP cores contained in Zynq-7000 devices or related international interface standards, refer the following documents:

Note: ARM documents can be found at: http://infocenter.arm.com/help/index.jsp

- *[illegible] Level 2 Cache Controller (L2C-310) TRM* (also called PL310)
  - *ARM [illegible]pecification Revision 2.0, 1999* (IHI 0011A)
  - *ARM Arch[illegible]ure Reference Manual* (Need to register with ARM)
  - *ARM Cortex[illegible] Series Programmer's Guide*
  - *ARM Cortex-A9 Technical Reference Manual, Revision r3p0*
  - *ARM Cortex-A9 MPCore Technical Reference Manu[illegible]ision r[illegible]*407F) – includes descriptions for accelerator coherenc[illegible] (ACP), CPU p[illegible] timers a[illegible] watchdog timers (AWDT), event bus, general[illegible]pt controller (GIC), and sn[illegible] control unit (SCU)
  - *ARM Cortex-A9 NEON Media Processing E[illegible]nical Reference Manual, R[illegible] r3p0*
  - *ARM Cortex-[illegible] Floating-Point Unit Techn[illegible]e Manual, Revision r3p0*
  - *ARM Cor[illegible] v1.0 Arc[illegible] Specificati[illegible]* descriptions f[illegible] ATB B[illegible] and A[illegible]ation
  - *[illegible]ght Program[illegible]e Architecture Sp[illegible]*
  - *ARM Debug Interface v5.1Architecture Specification*
  - *ARM Debug Interface v5.1 Architecture Specification Supplement*
  - *ARM CoreSight Components TRM* – includes descriptions for embedded cross trigger (ECT), embedded trace buffer (ETB), instrumentation trace macrocell (ITM), debug access port (DAP), and trace port interface unit (TPIU)
  - *ARM CoreSight PTM-A9 TRM*
  - *ARM CoreSight Trace Memory Controller Technical Reference Manual*
  - *ARM Generic Interrupt Controller v1.0 Architecture Specification* (IHI 0048B)
  - *ARM Generic Interrupt Controller PL390 Technical Reference Manual* (DDI0416B)
  - *ARM PrimeCell DMA Controller (PL330) Technical Reference Manual*
  - *ARM Application Note 239: Example programs for CoreLink DMA Controller DMA-330*
  - *ARM PrimeCell Static Memory Controller (PL350 series) Technical Reference Manual, Revision r2p1, 12 October 2007* (ARM DDI 0380G)
- *BOSCH, CAN Specification Version 2.0 PART A and PART B, 1991*

- *Cadence, Watchdog Timer (SWDT) Specification*
- *IEEE 802.3-2008 - IEEE Standard for Information technology-Specific requirements - Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, 2008*
- *Intel Corp., Enhanced Host Controller Interface Specification for Universal Serial Bus, [illegible]*
- *ISO 118[illegible]ard USB Association, USB 2.0 Specification*
- *SD Associatio[illegible]art A2 SD Host Controller Standard Specification Ver2.00 Final 070130*
- *SD Association,[illegible]rt E1 SDIO Specification Ver2.00 Final 070130*
- *SD Group, Part 1 Physical Layer Specification Ver2.00 Final 0605[illegible]*