XAPP623 (v2.1) February 28, 2005

Power Distribution System (PDS) Design: Using Bypass/Decoupling Capacitors

Author: Mark Alexander

Summary

This application note specifies how to build power distribution systems for Virtex™ devices. It also covers the basic principles of power distribution systems and bypass or decoupling capacitors. A step-by-step process is described where a power distribution system can be designed and verified. The final section discusses additional sources of power supply noise and provides resolutions.

Introduction

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

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

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

The first step in this process is to examine the utilization of the FPGA to get a rough idea of its transient current requirements. Next, a conservative decoupling network is designed to fit these requirements. The third step is to refine the network through simulation and modification of capacitor numbers and values. In the fourth step, the full design is built and in the fifth step it is measured. Measurements are made consisting of oscilloscope and possibly spectrum analyzer readings of power supply noise. Depending on the measured results, further iterations through the part selection and simulation steps could be necessary to optimize the PDS for the specific application. A sixth optional step is also given for cases where a perfectly optimized PDS is needed.

Basic Decoupling Network Principles

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

The purpose of a PDS is to provide power to the devices in a system. Each device in a system not only has its own power requirements for its operation, but also its own requirement for the

© 2003-2005 Xilinx, Inc. All rights reserved. All Xilinx trademarks, registered trademarks, patents, and further disclaimers are as listed at http://www.xilinx.com/legal.htm. All other trademarks and registered trademarks are the property of their respective owners. All specifications are subject to change without notice.

NOTICE OF DISCLAIMER: Xilinx is providing this design, code, or information "as is." By providing the design, code, or information as one possible implementation of this feature, application, or standard, Xilinx makes no representation that this implementation is free from any claims of infringement. You are responsible for obtaining any rights you may require for your implementation. Xilinx expressly disclaims any warranty whatsoever with respect to the adequacy of the implementation, including but not limited to any warranties or representations that this implementation is free from claims of infringement and any implied warranties of merchantability or fitness for a particular purpose.

cleanliness of that power. Most digital devices, including all Xilinx FPGAs, have a requirement that on all supplies, V_{CC} must not fluctuate more than 5% above or 5% below the nominal V_{CC} value. In this document V_{CC} is used generically to refer to all FPGA power supplies: V_{CCINT}, V_{CCO}, V_{CCAUX}, and V_{REF}. Multi-gigabit transceiver (MGT) analog supplies ($AV_{CCAUXTX}$, $AV_{CCAUXRX}$, VTTX, VTRX) are not covered here. For specific instructions on these supplies, see the *RocketIO™ Transceiver User Guide* (**Reference #1**).

This requirement specifies a maximum amount of noise present on the power supply, often referred to as "ripple voltage." If the device requirements state that V_{CC} must be within ±5% of the nominal voltage, that means peak to peak voltage ripple must be no more than 10% of the nominal V_{CC}. This assumes that nominal V_{CC} is exactly the nominal value given in the datasheet. If this is not the case, then V_{RIPPLE} must be adjusted to a value correspondingly less than 10%.

The power consumed by a digital device varies over time, and this variance occurs on all frequency scales. Low frequency variance of power consumption is usually the result of devices or large portions of devices being enabled or disabled. This can occur on time scales from milliseconds to days. High frequency variance of power consumption is the result of individual switching events inside a device, and this happens on the scale of the clock frequency and the first few harmonics of the clock frequency.

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

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

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

The second major component of the PDS is the bypass or decoupling capacitors. In this application note, the words "bypass" and "decoupling" are used interchangeably. Their function is to act as local energy storage for the device. They cannot provide DC power, as only a small amount of energy is stored in them (the voltage regulator is present to provide DC power). The function of this local energy storage is to respond very quickly to changing current demands. The capacitors are effective at maintaining power supply voltage at frequencies from hundreds of kilohertz to hundreds of megahertz in the milliseconds to nanoseconds range. Decoupling capacitors are of no use for all events occurring above or below this range. For example, if current demand in the device increases in a few picoseconds, the voltage at the device sags by some amount until the capacitors can supply extra charge to the device. If current demand in the device changes and maintains this new level for a number of milliseconds, the voltage regulator circuit, operating in parallel with the bypass capacitors, effectively takes over for them, changing its output to supply this new current.

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

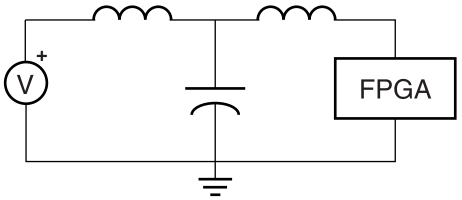

Figure 1: **Simplified PDS Circuit**

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

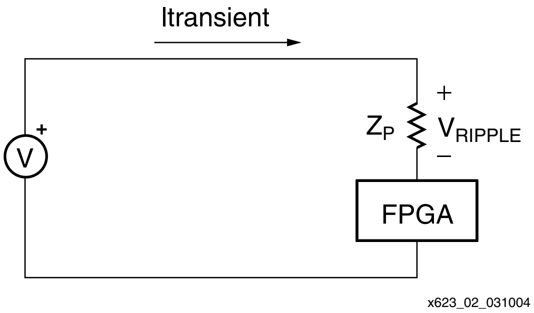

Figure 2: **Further Simplified PDS Circuit**

What is the Role of Inductance?

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

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

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

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

Capacitor Parasitic Inductance

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

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

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

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

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

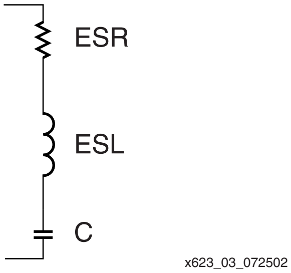

Figure 3: **Parasitics of a Real, Non-Ideal Capacitor**

Figure 4 shows the impedance characteristic of a real capacitor. Overlaid on this plot are the curves corresponding to the capacitor's capacitance and parasitic inductance (ESL). These two curves combine to form the total impedance characteristic of the RLC circuit formed by the parasitics of the capacitor.

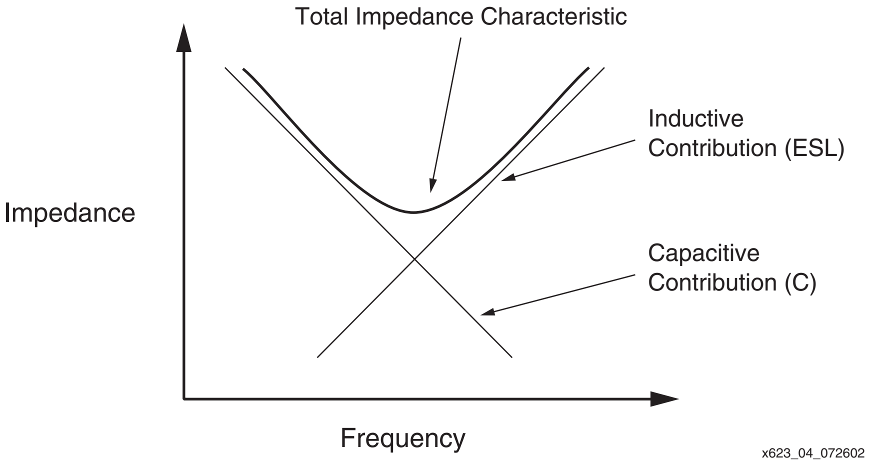

Figure 4: **Contribution of Parasitics to Total Impedance Characteristics**

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

Inductance from PCB Current Paths

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

Mounting Inductance

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

The vias, traces, and pads of a capacitor mounting contribute anywhere from 300 pH to 4 nH of inductance depending on the specific geometry. Since the inductance of a current path is proportional to the area of the loop the current traverses, it is important to minimize the size of this loop. The loop consists of the path through one power plane, up through one via, through the connecting trace to the land, through the capacitor, through the other land and connecting trace, down through the other via, and into the other plane, as shown in Figure 5.

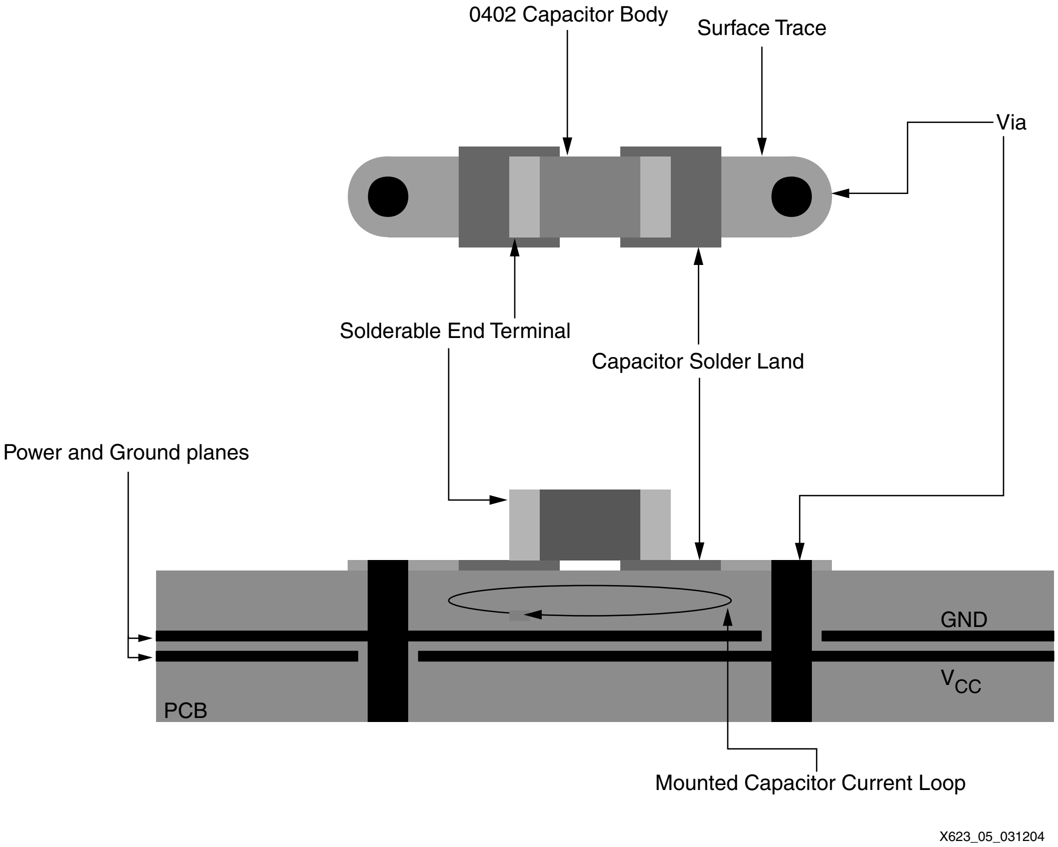

Figure 5: **Cutaway View of PCB with Capacitor Mounting**

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

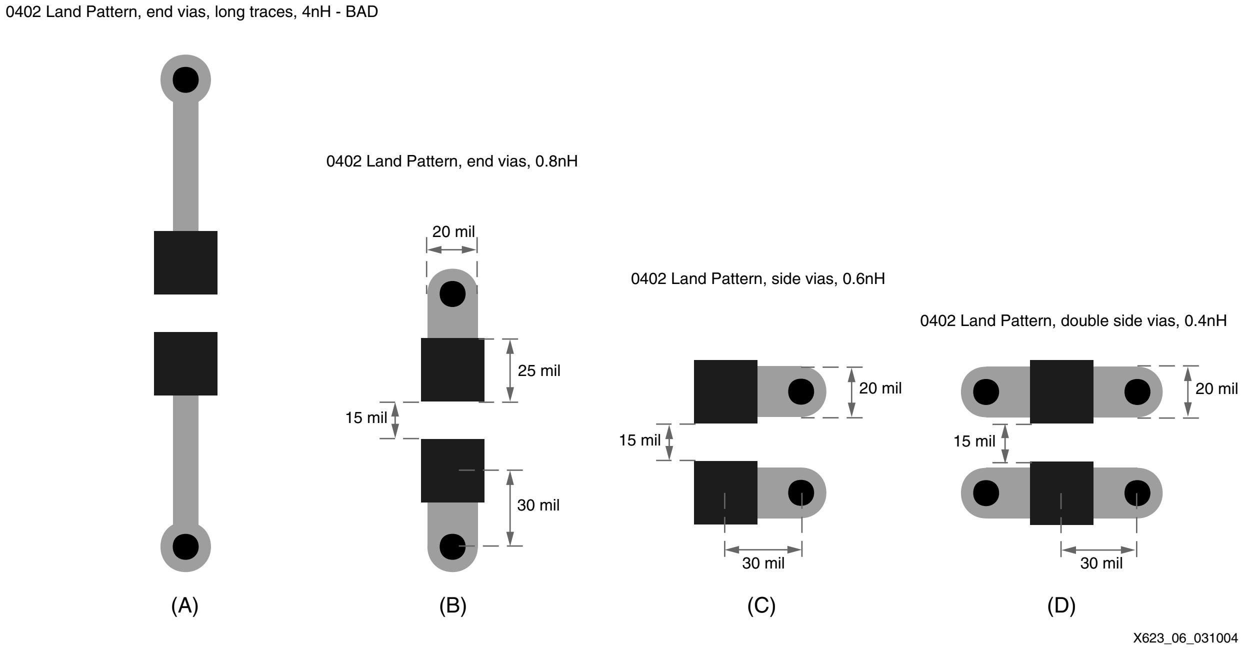
0402 Land Pattern, end vias, long traces, 4nH - BAD
0402 Land Pattern, end vias, 0.8nH
0402 Land Pattern, side vias, 0.6nH
0402 Land Pattern, double side vias, 0.4nH
20 mil
25 mil
15 mil
30 mil
20 mil
15 mil
30 mil
20 mil
15 mil
30 mil
(A)
(B)
(C)
(D)
X623_06_031004

Figure 6: **Example Capacitor Land and Mounting Geometries**

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

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

Plane Inductance

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

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

Spreading inductance acts like any other inductance — to resist changes to the amount of current in a conductor. In this case, the conductor is the power plane or planes. This quantity should be reduced as much as possible, since it retards the ability of capacitors to respond to transient currents in the device. Since the X-Y shape of the plane is typically something the designer has little control over, the only controllable factor is the spreading inductance value.

This is primarily determined by the thickness of the dielectric separating a power plane from its associated ground plane.

In high-frequency power distribution systems of the type discussed here, power and ground planes work in pairs. Their inductances do not exist independently of each other. The spacing (and the dielectric constant of the material) between power and ground planes determines the spreading inductance of the pair. The closer the spacing (the thinner the dielectric), the lower the spreading inductance. Table 1 gives approximate values of spreading inductance for different thicknesses of FR4 dielectric (**Reference #2**).

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

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

Since closer spacing results in decreased spreading inductance, it is best, wherever possible, to place V_{CC} planes directly adjacent to GND planes in the stackup. Facing V_{CC} and GND planes are sometimes referred to as "sandwiches." While the use of V_{CC} – GND sandwiches was not necessary in the past for previous technologies, the speeds involved and the sheer amount of power required for fast, dense devices demands it.

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

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

PCB Stackup and Layer Order

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

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

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

Capacitor Effective Frequency

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

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

To determine the resonant frequency of an RLC circuit, the following equation is used:

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

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

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

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

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

Example

X7R Ceramic Chip capacitor (AVX capacitor data used here)

$C = 0.01\ \mu F$

$L_{SELF} = 0.9\ nH$

$F_{RSELF} = 53\ MHz$

$L_{MOUNT} = 0.8\ nH$

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

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

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

Plugging in the values from the example:

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

$$F_{RIS} = \frac{1}{2\pi\sqrt{(1.7 \times 10^{-9} H) \cdot (1 \times 10^{-8} F)}} = 3.8 \times 10^{7} Hz$$

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

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

Capacitor Anti-Resonance

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

Capacitor Placement

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

First, increased spacing between the device and decoupling capacitor increases the distance travelled by the current in the power and ground planes, and hence, the inductance of the current path between the device and the capacitor. Since the inductance of this path (the loop followed by current as it goes from the V_{CC} side of the capacitor to the V_{CC} pin[s] of the FPGA, and from the GND pin[s] of the FPGA to the GND side of the capacitor[s]), is proportional to the loop area, decreasing its inductance is a matter of decreasing the loop area. Shortening the distance between the device and the decoupling capacitor(s) reduces the inductance resulting in a less impeded transient current flow. Because of the dimensions of PCBs, this reason tends to be less important with regard to placement than the second reason.

The second reason deals with the phase relationship between the FPGA noise source and the mounted capacitor. Their phase relationship determines the capacitor's effectiveness. For a capacitor to be effective in providing transient current at a certain frequency (for instance, the optimum frequency for that capacitor), it must be within a fraction of the wavelength associated with that frequency. The placement of the capacitor determines the length of the transmission line interconnect (in this case, the power and ground plane pair) between the capacitor and FPGA. The propagation delay of this interconnect is the relevant factor.

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

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

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

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

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

Example

0.001 µF X7R Ceramic Chip capacitor, 0402 package

L_{IS} = 1.6 nH

$$F_{RIS} = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{1.6\times10^{-9}\times 0.001\times10^{-6}}} = 125.8\,MHz$$

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

$$T_{RIS} = \frac{1}{F_{RIS}} = \frac{1}{125.8\times10^{6}} = 7.95\ ns \qquad \text{Equation 2}$$

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

$$\lambda = Wavelength = \frac{T_{RIS}}{V_{PROP}} \qquad \text{Equation 3}$$

where $V_{PROP} = 166\times10^{-12}\frac{s}{inch}$

$$\lambda = \frac{T_{RIS}}{V_{PROP}} = \frac{7.95\times10^{-9}}{166\times10^{-12}} = 47.9\ inches$$

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

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

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

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

Example Capacitor Layout

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

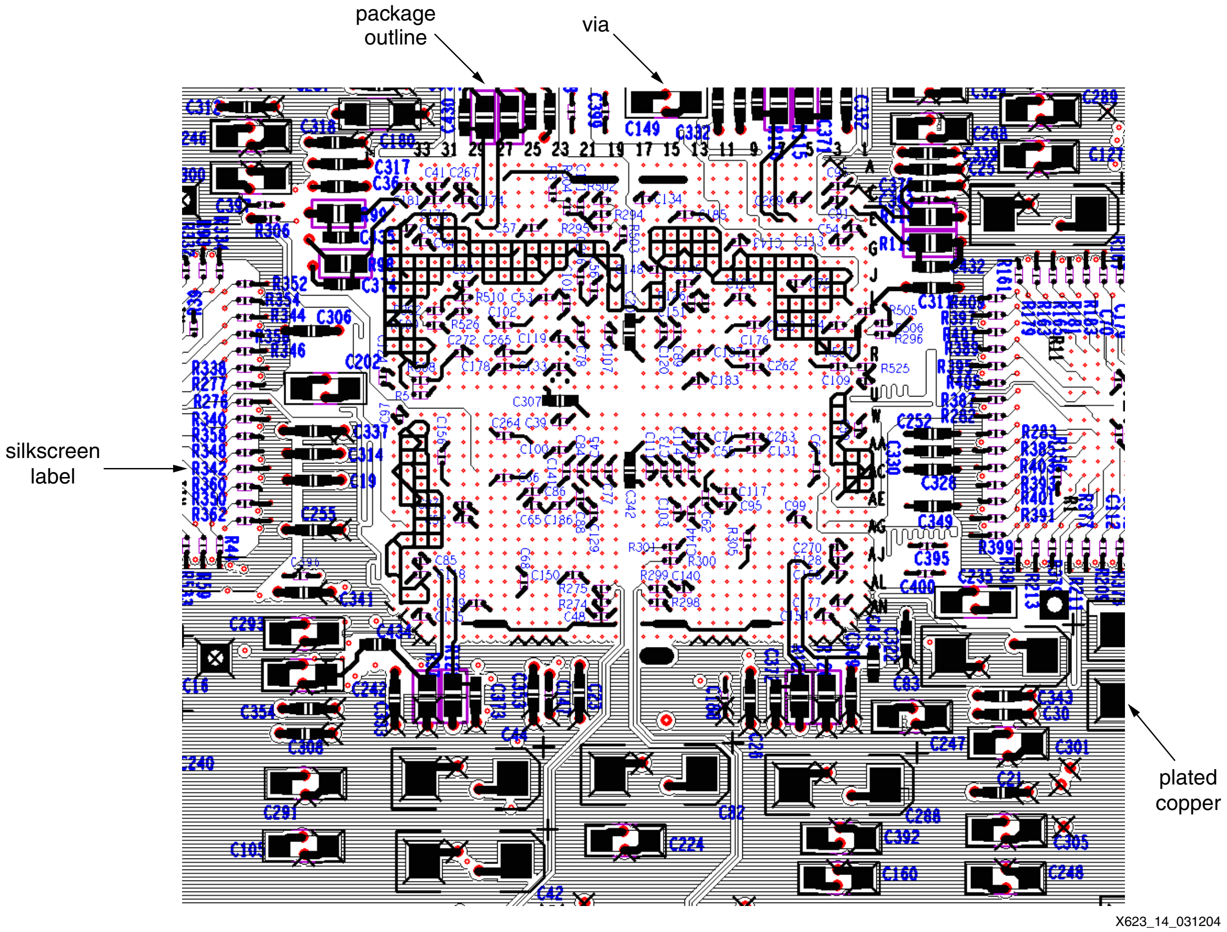

Figure 7: **Example PCB Layout showing Capacitor Placement on the Bottom Surface**

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

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

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

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

PDS Design and Verification

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

Step 1: Determining Critical Parameters of the FPGA

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

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

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

Pro-rating V_{CCO} Pins

The number of V_{CCO} pins used by a device can be determined based on the Simultaneously Switching Output (SSO) restrictions given in the device documentation (data sheet and user guide). A budget is calculated on a per-bank basis using these restrictions. The utilization of I/O resources in a bank determines the percentage of the budget used. This percentage effectively represents the percentage of V_{CCO} pins used by the device.

Examples: Using an XC2V3000 FF1152

Single bank and full device examples are provided.

Single Bank Example

In a hypothetical design, Bank 0 has 80 outputs in it. Each is configured as a 3.3V LVCMOS 12 mA Fast driver.

The limit for 3.3V LVCMOS 12 mA Fast drivers in the SSO table of the datasheet is 10 per V_{CC}/GND pair. There are 13 V_{CCO} pins per bank on this device. Therefore, the limit for this type of I/O driver is 130 per bank.

This bank uses 80 outputs. Therefore the percentage of the total bank 0 budget that is used is:

Bank 0 Percentage Used = Used/Limit = 80/130 = 62%

Full Device Example

In this example, the utilization of all I/Os for one device is listed in Table 2, as well as the per-bank SSO limits for each standard (Table 3), computed from the number per V_{CC}/GND pair SSO limits in the *Virtex-II Platform FPGA User Guide* (**Reference #3**).

Table 2: **I/O Utilization for Each Bank in Full Device Example**

Bank Number	Voltage	I/O Utilization	I/O Standard
Bank 0	3.3V	80	LVCMOS_12F
Bank 7	3.3V	80	LVCMOS_12F
Bank 1	1.5V	16	LVDCI
Bank 6	1.5V	16	LVDCI
Bank 2	1.8V	32	HSTL_1
		45	LVCMOS_12F
Bank 3	1.8V	32	HSTL_1
		45	LVCMOS_12F
Bank 4	1.8V	32	HSTL_1
		45	LVCMOS_12F
Bank 5	1.8V	32	HSTL_1
		45	LVCMOS_12F

Table 3: **SSO Limits for I/O Standards of Full Device Example**

I/O Standard	SSO Limit Per Bank
3.3V LVCMOS_12F	130
1.5V LVDCI	130
1.8V HSTL_1	260
1.8V LVCMOS_12F	117

The budget for banks 0, 7, 1, and 6 are computed as in the single-bank example. Banks 2, 3, 4, and 5, however, all have two I/O standards in them. For these banks, the budget is computed for each standard separately, and then the two are combined.

For Banks 2, 3, 4, and 5:

1.8V HSTL_1:

% used = used/limit = 32/260 = 13%

1.8V LVCMOS_12F:

% used = used/limit = 45/117 = 39%

Total budget for each bank:

13% + 39% = 52%

Table 4 summarizes the budgets for each bank of the device.

Table 4: **Budgets for Each Bank in Full Device Example**

Bank Number	Budget
Bank 0	62%
Bank 7	62%
Bank 1	12%
Bank 6	12%
Bank 2	52%
Bank 3	52%
Bank 4	52%
Bank 5	52%

The number of V_{CCO} pins used in a bank (Table 5) is simply the number of V_{CCO} pins in a bank times the percentage of the SSO budget used.

Table 5: **Number of V_{CCO} Pins Used**

Bank Number	Calculated	Number of Pins Used
Bank 0	13 pins × 62%	8 pins
Bank 7	13 pins × 62%	8 pins
Bank 1	13 pins × 12%	2 pins
Bank 6	13 pins × 12%	2 pins
Bank 2	13 pins × 52%	7 pins
Bank 3	13 pins × 52%	7 pins
Bank 4	13 pins × 52%	7 pins
Bank 5	13 pins × 52%	7 pins

Step 2: Designing the Generic Bypassing Network

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

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

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

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

A ratio of capacitors giving a relatively flat impedance is one where the quantity of capacitors is roughly doubled for every decade of decrease in size. In other words, if the bottom three values in the network were 1.0 µF, 0.1 µF and 0.01 µF, the network might have two 1.0 µF capacitor, four 0.1 µF capacitors, and eight 0.01 µF capacitors.

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

A set of percentages is helpful for calculating these ratios based on the total number of capacitors (Table 6).

Table 6: **Capacitor Value Percentages for a Balanced Decoupling Network**

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

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

1.5V Supply Example

In this example, the 1.5V supply for the Virtex-II device supplies V_{CCO} for banks 1 and 6, and V_{CCINT}. There are 44 V_{CCINT} pins on this device. Banks 1 and 6 were previously calculated to use two pins each. Adding the 44 V_{CCINT} pins and the four V_{CCO} pins for Banks 1 and 6 equals 48 pins. Therefore, there should be 48 capacitors total on the 1.5V supply. Table 7 shows how the quantity of each value of capacitor is determined.

Table 7: **Calculation of Capacitor Quantities for 1.5V Supply Example**

Capacitor Value	Calculated	Quantity of Capacitors
680 µF	48 pins x 4% = 1.92	2
2.2 µF	48 pins x 14% = 6.72	7
0.47 µF	48 pins x 27% = 12.6	13
0.047 µF	48 pins x 55% = 26.4	26

This calculation gives a first-pass estimate of the capacitors necessary for the 1.5V supply. Changes can be made to the exact number of capacitors to accommodate different values and to make the supply more symmetric (e.g., using eight 2.2 µF capacitors instead of seven for a more standard the PCB layout). Capacitor values can also be modified according to the specific constraints of the design (e.g., a pre-existing BOM of capacitors). This process of capacitor selection must be repeated for each supply.

Step 3: Simulation

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

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

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

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

Figure 8 shows a simple impedance plot from a simulation of the parallel combination of these capacitors, taking into account their parasitics and the approximate parasitics of the PCB. An equivalent SPICE netlist is included in Appendix C: SPICE Simulation Examples. Table 8 lists the capacitor quantities, values, and parasitic values used in the simulation. The RLC characteristics of the V_{CC} and GND planes of the PCB are not taken into account.

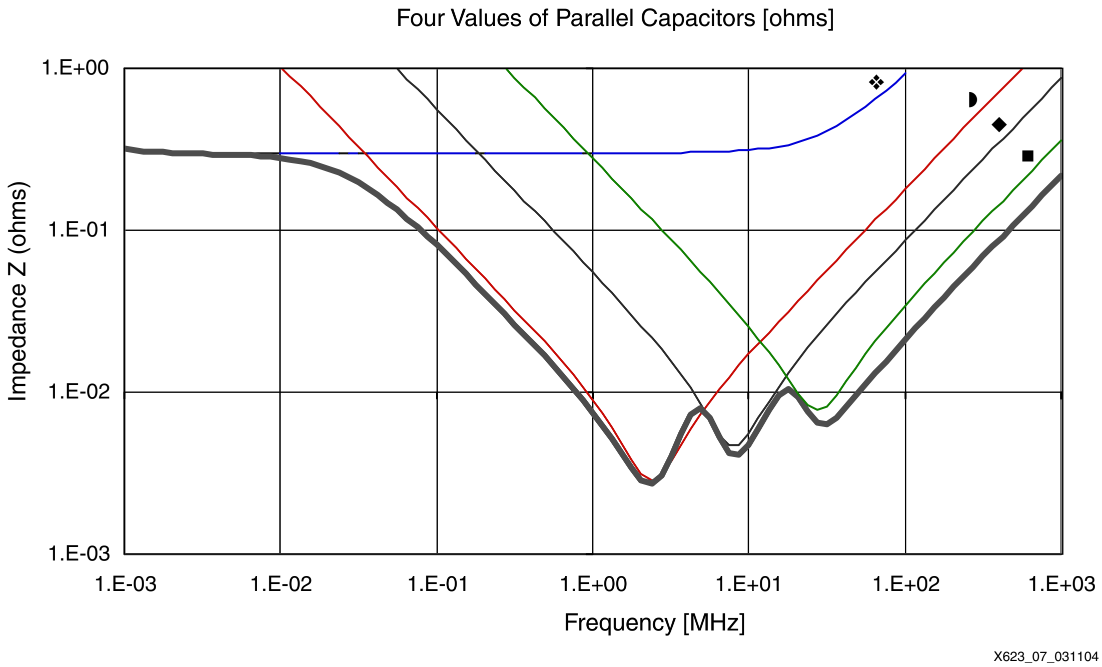

Figure 8: **PDS Impedance Versus Frequency Plot**

Table 8: **Values Used in Impedance Plot of Figure 8**

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

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

Step 4: Building the Design

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

Step 5: Measuring Performance

In the performance measurement step, measurements are made to determine whether the PDS is adequate for the devices it is serving. Determining whether or not a bypassing network is adequate for a given design is relatively simple. The measurement is performed with a high-bandwidth oscilloscope (1 GHz oscilloscope and 1 GHz probe at minimum), on a design running realistic test patterns.

Noise Magnitude Measurement

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

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

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

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

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

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

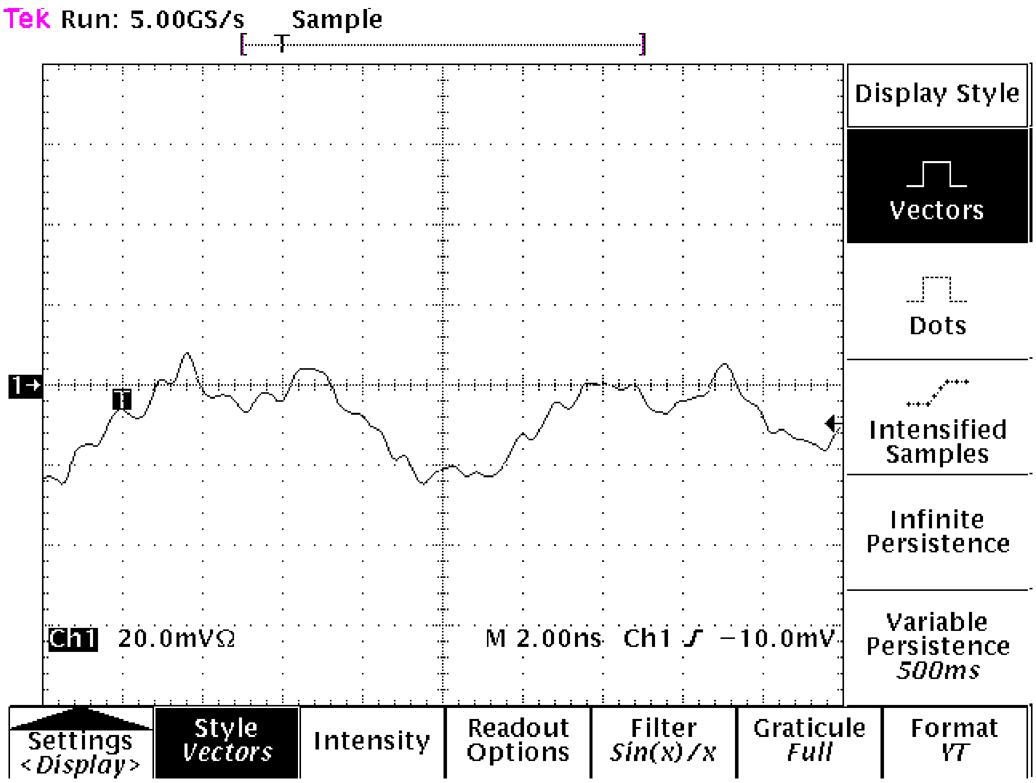
Tek Run: 5.00GS/s Sample

Ch1 20.0mVΩ M 2.00ns Ch1 ƒ −10.0mV

Display Style: Vectors

Style: Vectors

1→
T
Display Style
Vectors
Dots
Intensified Samples
Infinite Persistence
Variable Persistence 500ms
Settings <Display>
Style Vectors
Intensity
Readout Options
Filter Sin(x)/x
Graticule Full
Format YT

x623_08_080502

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

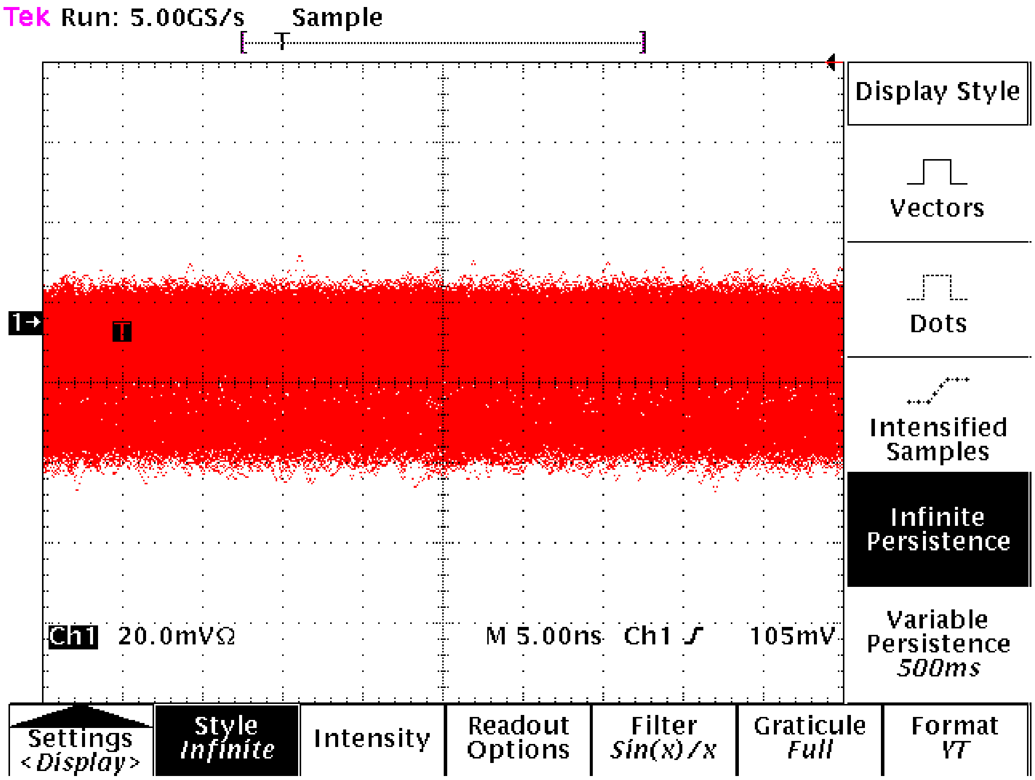
Tek Run: 5.00GS/s Sample

Ch1 20.0mVΩ M 5.00ns Ch1 ƒ 105mV

Display Style: Infinite Persistence

Style: Infinite

1→
T
Display Style
Vectors
Dots
Intensified Samples
Infinite Persistence
Variable Persistence 500ms
Settings <Display>
Style Infinite
Intensity
Readout Options
Filter Sin(x)/x
Graticule Full
Format YT

x623_09_090502

Figure 10: **Infinite Persistence Measurement of Same Supply**

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

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

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

Noise Spectrum Measurements

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

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

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

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

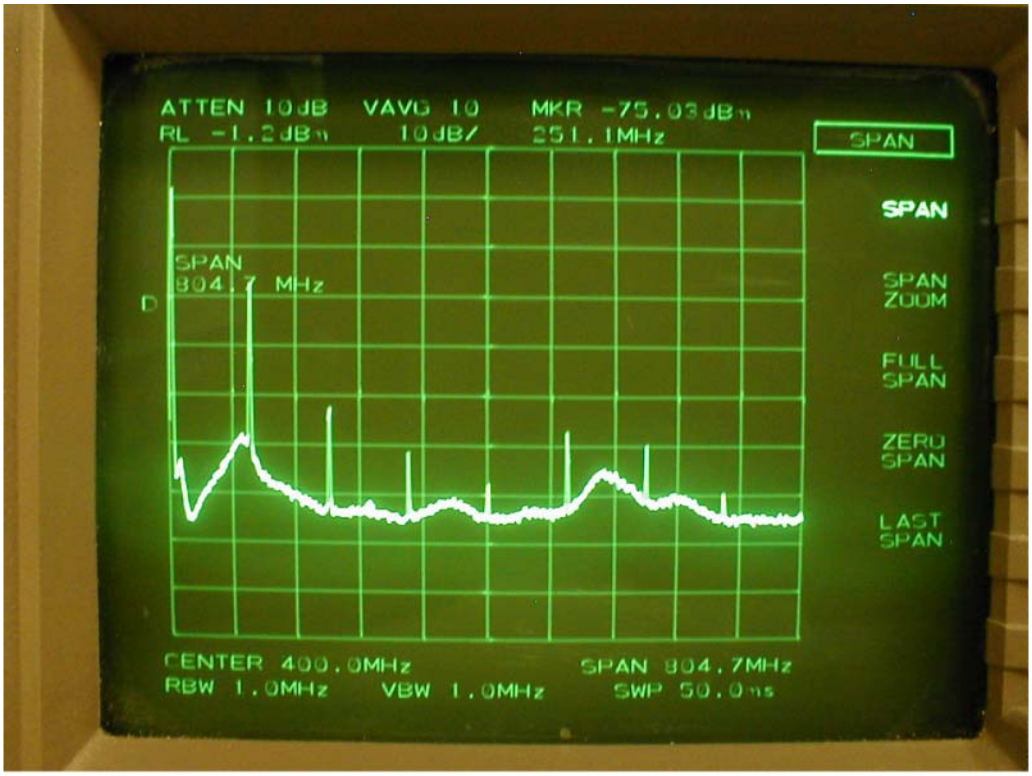

Figure 11: **Screenshot of Spectrum Analyzer Measurement of V_{CCO}**

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

Step 6: Optimum Bypassing Network Design (Optional)

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

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

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

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

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

Other Concerns and Causes

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

Possibility 1: Excessive Noise from Other Devices on Board

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

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

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

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

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

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

The spreading inductance of the planes is controlled by the plane spacing and by the dielectric constant of the material between them. See section on "Plane Inductance".

When boards are exceptionally thick (greater than 90 mils or 2.3 mm), vias have higher parasitic inductance. In these cases, the following changes to the design should be considered. The first is to move the V_{CC}/GND plane sandwiches close to the top surface the FPGA is on. The second is to place the highest frequency capacitors on the top surface. Both changes together reduce the parasitic inductance of the relevant current path.

Possibility 3: I/O Signals in PCB are Stronger Than Necessary

If noise in the V_{CCO} PDS is still too high after making refinements to the PDS, the I/O interface power can be scaled back. This goes for both outputs from the FPGA and inputs to the FPGA. In some cases, excessive overshoot on inputs to the FPGA can reverse-bias the clamp diodes in the IOBs. This can put large amounts of noise into V_{CCO}. If this condition is occurring, the drive strength of these interfaces should be decreased, or termination should be used (both on input and output paths).

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

Excessive noise in the PDS can be caused by I/O signal return currents. For every signal transmitted by a device into the PCB (and eventually into another device), there is an equal and opposite current flowing from the PCB back into the device's power/ground system. If there is no low-impedance return current path available, a less optimal, higher impedance path is used. When this occurs, voltage changes are induced in the PDS.

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

Conclusion

This application note is an overview of the important principles of power distribution systems, followed by a step-by-step process for designing a PDS. This is an iterative method of PDS design, where the designer first creates a generic network, simulates and refines it, then measures it, and then refines it again based on the measured results is described. When this method fails to give an acceptable result, other possible contributors to the problem are explored. Through the use of this method, all PDS problems can be resolved.

References

1. Xilinx, Inc., *RocketIO Transceiver User Guide*, UG024
2. Larry D. Smith "Decoupling Capacitor Calculations For CMOS Circuits," Proceedings EPEP Conference, November 1984
3. Xilinx, Inc., *Virtex-II Platform FPGA User Guide*, UG002
4. Frederick W. Grover Ph.D. "Inductance Calculations: Working Formulas and Tables", D. Van Nostrand Company, Inc. 250 Fourth Avenue New York 1946

Appendix A: Glossary

Land: A section of exposed metal on the surface of a PCB where surface-mount devices are soldered

Network Analyzer: An instrument used to measure the frequency-domain characteristics of electrical networks. The electrical characteristics of power distribution systems are often measured using a network analyzer.

Oscilloscope: An instrument used to show the time-domain voltage of a signal. Power supply noise is the signal measured when establishing the magnitude of noise voltage on a power supply.

Sandwich: A pair of planes in a PCB stackup separated only by dielectric material, no signal layer is in-between. In most cases, one of these planes is at ground potential and the other plane carries power. Also known as buried capacitance.

Spectrum Analyzer: An instrument used to measure the frequency content of a signal. Power supply noise is the signal measured when establishing the characteristics of a power distribution system.

Stackup: The series of layers in a PCB is often referred to as a stackup. Multi-layered boards are comprised of alternating layers of signal routing or plane metal and dielectric material. The dielectric material also serves as a structural substrate.

Via: A vertical connection in a PCB, usually formed by drilling a hole through the PCB and plating the walls of this hole with conductive material. Vias make electrical connections between different layers of a PCB. Vias can represent impedance discontinuities when they are in a signal path, and represent additional parasitic inductance when they are in a power distribution path (both are undesirable). The parasitic inductance formula is shown in Appendix B: Calculation of Via Inductance.

Voltage Ripple: Power supply noise is often referred to as voltage ripple. The maximum voltage ripple corresponds to the maximum amount of power supply variation allowed by a part's absolute maximum ratings.

Appendix B: Calculation of Via Inductance

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

$$L = 5.08 \times h \times \left[\ln\left(\frac{4 \times h}{d}\right) - 0.\dot{7}5\right] \qquad \text{Equation 5}$$

Example

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

h = 0.062 in
d = 0.003 in

$$L = 5.08 \times h \times \left[\ln\left(\frac{4 \times h}{d}\right) - 0.\dot{7}5\right]$$

$$L = 5.08 \times 0.062 \times \left[\ln\left(\frac{4 \times 0.062}{0.003}\right) - 0.\dot{7}5\right]$$

L = 5.08 x 0.062 x 3.67 nH

L = 1.15 nH

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

Appendix C: SPICE Simulation Examples

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

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

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

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

Two other details complete the SPICE decks:

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

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

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

HSPICE Netlist

The HSPICE Netlist is available on the Xilinx FTP site:
http://www.xilinx.com/bvdocs/appnotes/xapp623.zip

HSPICE Output

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

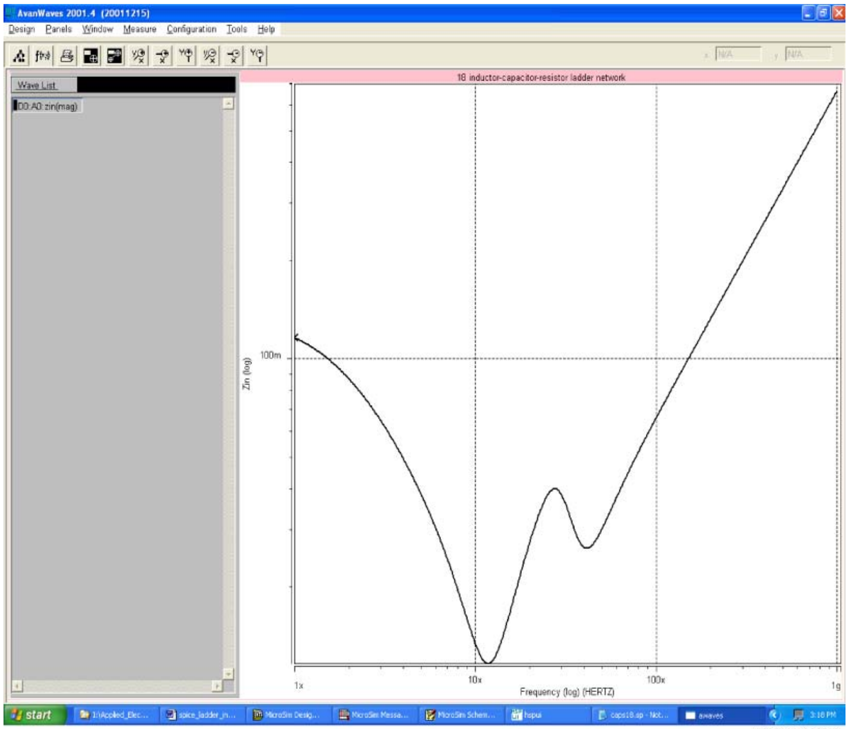
AvanWaves 2001.4 (20011215)
Design Panels Window Measure Configuration Tools Help
N/A
N/A
Wave List
D0:A0:zin(mag)
18 inductor-capacitor-resistor ladder network
100m
Zin (log)
1x
10x
100x
1g
Frequency (log) (HERTZ)
start
1:\Applied_Elec...
spice_ladder_n...
MicroSim Desig...
MicroSim Messa...
MicroSim Schem...
hspui
caps18.sp - Not...
awaves
3:18 PM

x623_11_080502

Figure 12: **HSPICE Output**

Schematic Circuit

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

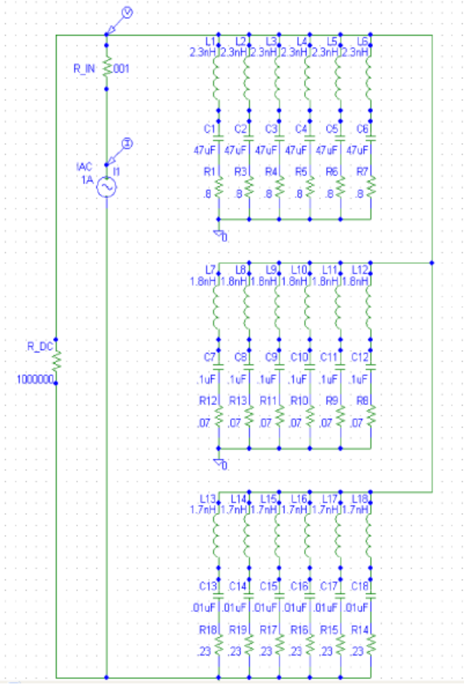
V
R_IN
.001
I
IAC
1A
I1
R_DC
1000000
L1 L2 L3 L4 L5 L6
2.3nH 2.3nH 2.3nH 2.3nH 2.3nH 2.3nH
C1 C2 C3 C4 C5 C6
47uF 47uF 47uF 47uF 47uF 47uF
R1 R3 R4 R5 R6 R7
.8 .8 .8 .8 .8 .8
0
L7 L8 L9 L10 L11 L12
1.8nH 1.8nH 1.8nH 1.8nH 1.8nH 1.8nH
C7 C8 C9 C10 C11 C12
.1uF .1uF .1uF .1uF .1uF .1uF
R12 R13 R11 R10 R9 R8
.07 .07 .07 .07 .07 .07
0
L13 L14 L15 L16 L17 L18
1.7nH 1.7nH 1.7nH 1.7nH 1.7nH 1.7nH
C13 C14 C15 C16 C17 C18
.01uF .01uF .01uF .01uF .01uF .01uF
R18 R19 R17 R16 R15 R14
.23 .23 .23 .23 .23 .23

x623_12_080602

Figure 13: **Schematic Circuit**

Appendix D: EDA Tools for PDS Design and Simulation

Table 9 lists the some vendors of EDA tools for PDS design and simulation.

Table 9: **EDA Tools for PDS Design and Simulation**

Tool	Vendor	Website URL
SIwave	Ansoft	http://www.ansoft.com
Specctraquest Power Integrity	Cadence	http://www.cadence.com
Speed 2000	Sigrity	http://www.sigrity.com
Star HSPICE	Synopsys	http://www.synopsys.com
UCADESR3.exe	UltraCAD	http://www.ultracad.com

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
08/08/02	1.0	Initial Xilinx release.
04/21/03	1.1	Updated for smaller file size.
06/11/03	1.2	Minor text changes for clarity.
04/05/04	2.0	Replaced Figure 5, Figure 6 and Figure 8. Added new sections called "PCB Stackup and Layer Order", "Capacitor Anti-Resonance", and "Example Capacitor Layout" (the new Figure 7 is part of the "Example Capacitor Layout" section). Removed all references to PSPICE. Made additional text changes for clarity.
02/28/05	2.1	Revised the denominator in the F_{RIS} equation shown on page 9, under "Capacitor Effective Frequency."