

DesignCon 2005

High Performance FPGA Bypass Networks

Steve Weir, Teraspeed Consulting Group
steve@teraspeed.com

Scott McMorrow, President, Teraspeed Consulting Group
scott@teraspeed.com
(401) 284-1827

Abstract

This paper demonstrates efficient synthesis of power supply high frequency bypass networks for FPGAs, using both ordinary and low inductance capacitors. We demonstrate where low inductance capacitors may be used to reduce bypass capacitor count by up to four to one, and cite a case study using Xilinx FPGAs. We demonstrate fixturing for both board and device characterization.

Much has been written about designing bypass networks with many rules and guidelines supplied. Many of the current guidelines available are premised on assumptions that are at best true only under certain conditions. The result is that many systems have been deployed that work almost by dumb luck alone. The cost of inappropriate designs in either excessive cost, or far worse- failing systems is unacceptable. With an appreciation of the significant mechanics involved, systems can be designed to work properly without design overkill.

Author(s) Biography

Steve Weir, Consultant, Teraspeed Consulting Group

Steve is an independent consultant with over 20 years plus industry experience with a broad range of expertise. Steve holds 17 US patents, and has architected a number of TDM and packet based switching products, consults on patents and is a frequent contributor to the SI-list signal integrity reflector.

Scott McMorrow, President Teraspeed Consulting Group

Mr. McMorrow is an experienced technologist with over 20 years of broad background in complex system design, interconnect and Signal Integrity engineering, modeling and measurement methodology, engineering team building and professional training. Mr. McMorrow has a consistent history of delivering and managing technical consultation that enables clients to manufacture systems with state-of-the-art performance, enhanced design margins, lower cost and reduced risk. Mr. McMorrow is an expert in high-performance design and Signal Integrity engineering, and has been a consultant and trainer to engineering organizations worldwide.

The Task- AC Short from Near DC to FPGA Package Cut-off

Bypass networks¹ serve the lowly, but critical function of supporting IC supply voltage up to the ICs that actually provide functional value. They do so by presenting a parallel admittance across the supply rails. An ideal network would be a “short that supports DC”.

Conceptually, we can design a bypass network in four steps:

1. Determine the AC noise voltage budget.
2. Determine the AC noise current versus frequency.
3. Derive the required impedance versus frequency.
4. Design a physical layout and bypass network to meet the required profile.

The first step is straight forward, as the designer allocates some fraction of the voltage error budget to DC accuracy, line and load regulation in the VRM and interconnect, and the remainder to AC noise. The second step may be quite involved and is beyond the scope of this paper. What must be kept in mind however is that IC packaging itself acts as a multi-pole low-pass filter that greatly attenuates current flow beyond the cut-off frequency. A reasonable approach is to design the board level impedance to:

- Meet target impedance out to the package cut-off frequency.
- Exhibit stable phase (ie inductive) near the package cut-off frequency.

This paper focuses on the third and fourth steps. Voltage is maintained by imposing a sufficiently high admittance shunt across the power rails. From an AC standpoint, our goal is to approximate welding the power and ground planes together. At high frequency, vias guarantee that the shunt impedance is inductive. For purposes of bypass, and for most materials, the PCB planes also appear inductive.

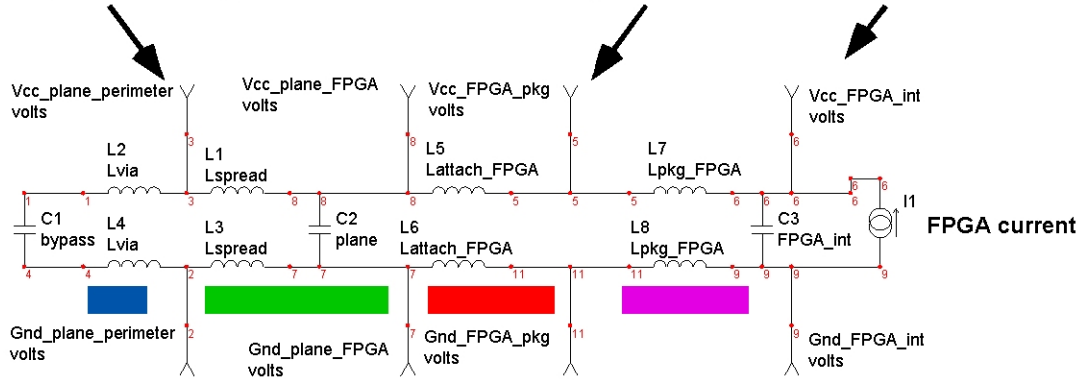
We are concerned with four sources of inductance and hence impedance at high frequency:

- FPGA attachment inductance
- Plane spreading inductance
- Bypass capacitor self inductance
- Bypass capacitor attachment inductance

Figure 1. depicts a schematic from the FPGA to the bypass capacitors:

¹ Once upon a time before power planes became commonplace, systems had power-decoupling networks. Those networks isolated noise currents between system nodes. Each filter included a series element to limit noise currents in the power wiring and local bypass capacitors to satisfy local device transient currents. The advent of printed circuit planes largely eliminated the need for series elements for digital logic, leaving only the bypass capacitors from the original decoupling networks.

Plane noise measured away from power ring Component specs voltage noise here SSN on idle I/Os visible here



FPGA attachment, plane spreading, and capacitor attachment via inductance all work to decouple the bypass capacitors from the FPGA.

Figure 1, Schematic PDS, FPGA to Bypass Capacitors




Figure 2, Partial Inductance Components, Bypass Cap to IC

Spreading Inductance-

Absent very high Er materials, over the distance between the FPGA center to the bypass capacitors, the planes provide little supplemental charge, and instead primarily behave as inductive conductors between the IC vias and the bypass capacitor vias.

Spreading inductance may be derived using the Biot-Savart law¹. For a BGA with power and ground concentrated in central rings as with VirtexII / Pro, we can work the inductance as that of a ring. We integrate flux against the radial distance from the power / ground vias on the ring inner edge to the bypass capacitor vias on the ring outer edge. The current density in any circular filament dr is $1/(r * dr * 2\pi)$, while the incremental increase in loop distance is dr. Assuming no plane perforation, the integral of the loop is:

Equation 1

$$\mu_0 * I / 2\pi * H * \ln(R2/R1)$$

Where:

μ_0 is the permeability of free space, 31.9nH/square.

H is the plane separation in inches.

R2 is the radius from the die center to the bypass capacitor vias.

R1 is the radius from the die center to the IC power / ground vias.

It should be readily apparent from this that the spreading inductance depends on:

- 1) The **ratio** of the radius from die center to power / ground attachment to the bypass capacitor vias to the radius from the die center to the IC attachment vias.
- 2) The plane separation.

Perforation increases plane inductance, and has been described by Yang, et-al.² For small anti-pads, a crude approximation is the linear proportion of total area / perforated area.

FPGAs with Power / Ground Slug-

VirtexII and VirtexII-Pro FPGAs concentrate ground connections in a slug at the middle of the part. Around this slug a ring of core power is then wrapped with two rings of I/O power. The remaining I/O power pins, and a number of additional grounds distribute around the remaining area of the part as shown in Figure 3.

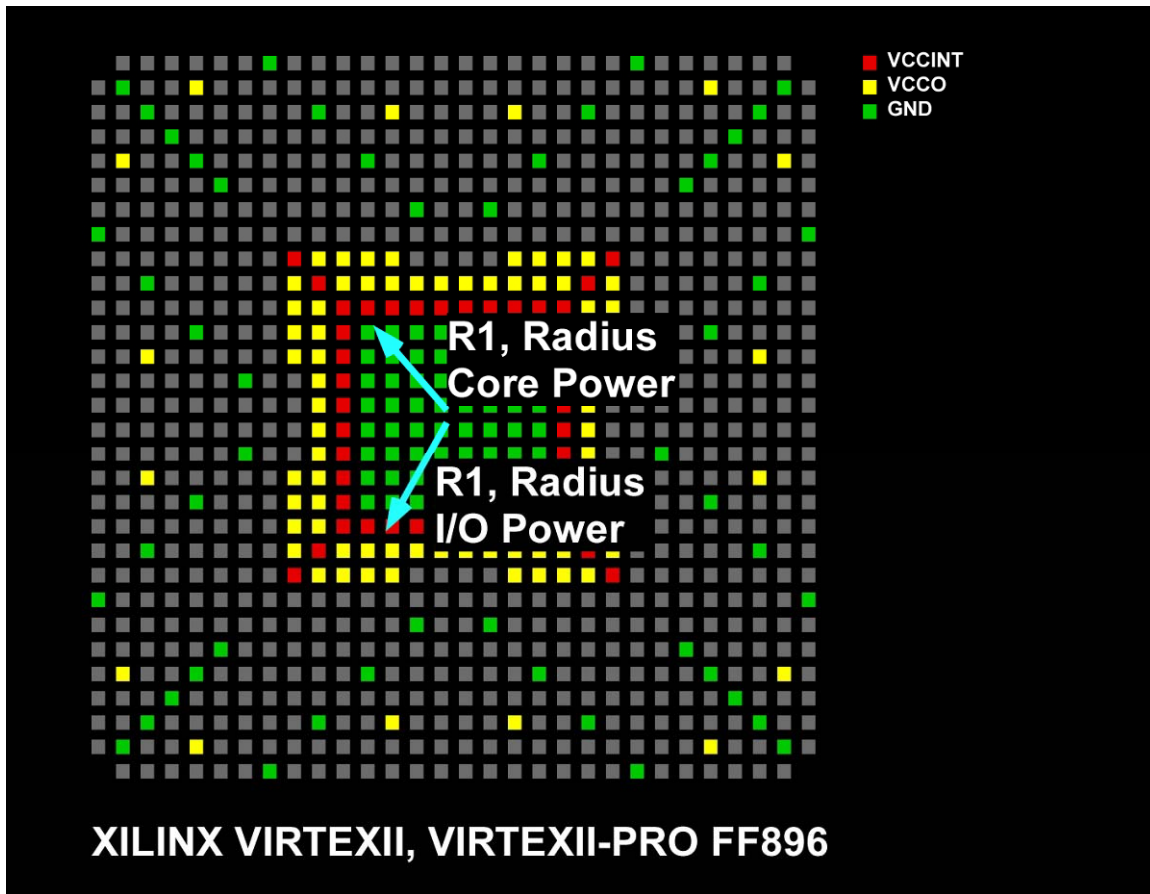


Figure 3, VirtexII / Pro(tm) FF896 Power / Ground Assignments

This sets R1 in the range of 250 – 400mils. For bypass capacitors mounted around and within 200mils of the device perimeter, R2 ranges from 700 – 1000 mils from the device center. $\ln(R2 / R1)$ therefore varies over a relatively narrow range from about 1 to 1.3.

The good news is that the capacitors are already so far away from the power attachments, that moving the capacitors further has little effect on spreading inductance. This provides latitude to locate termination networks as close to the device as possible without substantial adverse consequences. The bad news is that the spreading inductance for common plane separations is already quite high.

Using 3.2 as a typical value for R2/R1, the spreading inductance reduces to:

Equation 2

- $\mu_0 * h / 2\pi * 1.15 = 5.9\text{pH/mil}$, unperforated, or

Taking a crude approximation of perforated inductance as unperforated area / perforated area:

Equation 3

- $\mu_0 * h / 2\pi * 1.15 * 1.25 = 7.4\text{pH/mil}$, perforated 1mm pitch 20mil anti-pads

An example 14 mil dual strip line sandwich of: Gnd – Sig – Sig – Power sets a floor of 104pH with zero via and zero bypass capacitor impedance. Figure 3 charts the total inductance reflected back to the plane at the power / ground ring:

**Example Loop Inductance:
14mil Plane Separation, 250mil R1, 800mil R2**




Figure 4, Spreading Inductance Effects, 14 mil Plane Spacing

The legend at the bottom of the chart reports incremental admittance resulting from incremental addition of capacitors in groups of ten normalized to admittance with only the first ten capacitors. Increasing the count from ten capacitors to twenty only reduces the inductance at the IC power / ground ring by 25%,

ie a 33% increase in admittance. Absent the contribution of plane spreading inductance, the admittance would have increased by 100%. Consequently, the additional capacitor / vias are only 33% as effective as they otherwise would be under a condition of zero spreading inductance. The situation degrades rapidly with additional capacitors. Beyond 30 capacitors, even massive numbers of additional capacitors provide very little benefit to FPGA power.

High spreading inductance effectively isolates capacitors located outward of the FPGA. “Carpet bombing” the PCB will reduce the PCB impedance but does little to improve transfer impedance where it counts most: at the IC power / ground connections. We could TIG weld the planes together at the periphery with little improvement. Design guidelines that call for high capacitor counts assume negligible effect from combined spreading and FPGA attachment inductances.

Failing to account for spreading inductance can also lead to significant measurement error. Referring to Figure 4, we see that high FPGA attachment inductance, resulting from poor via choices and/or planes located far from the device, and/or high spreading inductance, effectively isolates the bypass capacitor network from the FPGA. Insertion loss measured from the FPGA center to the capacitor array can be very high, while measured noise at or near the capacitor array will be low. The power system erroneously viewed this way may look more than adequate, while the FPGA starves.

To improve bypass performance as viewed at the IC, we must reduce the FPGA attachment and/or spreading inductances. We must either go to thinner dielectrics or locate bypass capacitors closer to the FPGA power / ground connections.

Backside capacitors-

For a number of board geometries, significant gains may be realized by mounting capacitors on the backside of the PWB.

In the case of VirtexII/Pro packages core power may be bypassed using four vias per capacitor by plugging via holes and setting capacitors in the power ring. In the example pattern shown in Figure 4, applicable to both the FF896 and FF1152 packages, twenty-eight ordinary 0402 capacitors bypass core power. Four each power and ground vias are added to cover the gaps caused by break-out fanning.

Despite the deleterious effects of shared vias and mutual coupling between the capacitors, less than 100pH is realizable on 0.062” boards, for planes at the center of the stack-up. This is significantly better than what can be realized with arbitrary capacitors at the periphery with a 0.014” plane separation. When combined with a modest count of peripheral capacitors, similar results to 0.003” dielectric with only peripheral capacitors can be had, and may spare the cost of a thin dielectric layer pair otherwise needed.

However with thick boards where the power plane is located close to the IC, the utility of backside capacitors can be very questionable. For a balanced stack-up, the effective plane cavity that defines via inductance approaches the board thickness, and inductance rises nearly linearly with thickness.² Were one able to manufacture an unbalanced board, with the planes only near the IC side of the board, via inductances could skyrocket to the order of 10nH for backside attachments.

² See Equation 6. For thick boards that are balanced, H1 in the equation is the height to the ground plane closest to the decoupling capacitor, and H2 is the cavity height from this ground to the power plane closest to the FPGA.




Figure 5, Example Under Chip Bypass, Vccint

Spreading Inductance With Thin Dielectric-

Spreading inductance improves linearly with height reduction. 3mil material is widely available at relatively low-cost:

**Example Loop Inductance:
3mil Plane Separation, 250mil R1, 800mil R2**




Figure 6, Spreading Inductance Effects, 3 Mil Plane Spacing

With 3 mil planes up to forty conventional capacitors remain effective.

Similarly, very thin dielectric, such as 3M C-Ply, at 0.63 mils permits gainful use of more than 100 conventional capacitors, or alternatively dramatic reduction in required capacitor count:

**Example Loop Inductance:
0.63mil Plane Separation, 250mil R1, 800mil R2**

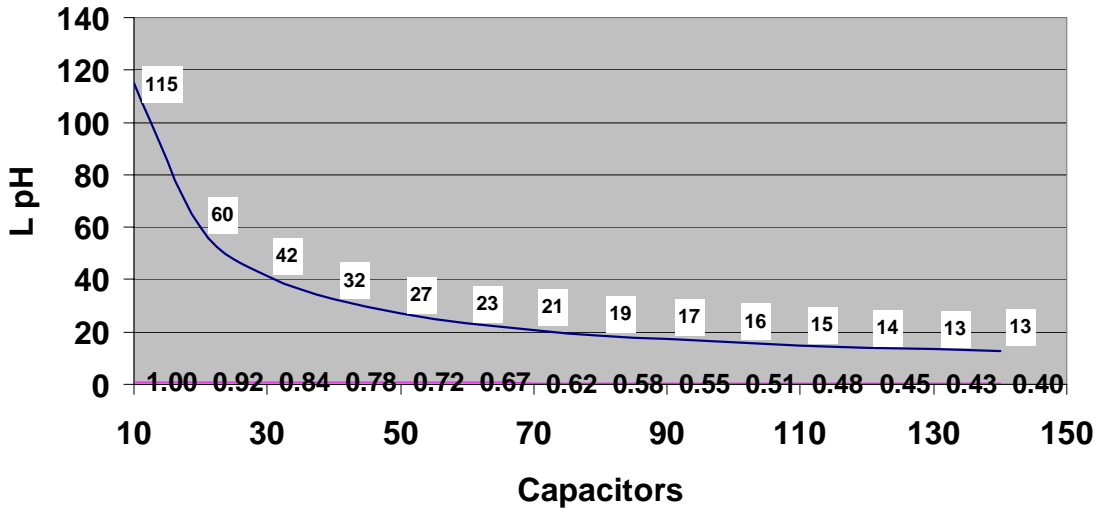


Figure 7, Spreading Inductance Effects, 0.63 Mil Dielectric

With thin material, incremental capacitors remain 50% effective out to 100 parts / 16pH / 5mΩ @ 50MHz. The impedance floor is an order of magnitude less than with 14 mil planes.

Spreading Inductance Demonstration Fixture-

To demonstrate the effects of spreading inductance on transfer impedance, we constructed a set of test boards that contain only bypass capacitors, power and ground vias for a Xilinx FF896 package, and probe points. These boards are simple 0.062", 4 layer construction with inner planes at 0.012", and 0.050" inches respectively, 0.038" plane spacing. Vias are 8 mil diameter. The BGA periphery is "carpet bombed" with 104 0402 capacitors to match design guidelines.

For our measurements, we utilized an Agilent 8753A VNA. Because we are only interested in fairly pedestrian frequencies up to 100MHz or so, we employed ordinary 0.025" square post connectors on 0.100" centers for convenient access at 23 points around each board. Simple adapter cards feed vertical launch SMAs to 2 pin female headers. We calibrated the entire cable assembly back to back and note less than 0.1dB insertion loss from 100KHz to 300MHz prior to calibration. At 50MHz back-to-back probes show 0.0147dB loss, corresponding to a parasitic inductance of 67pH / probe.

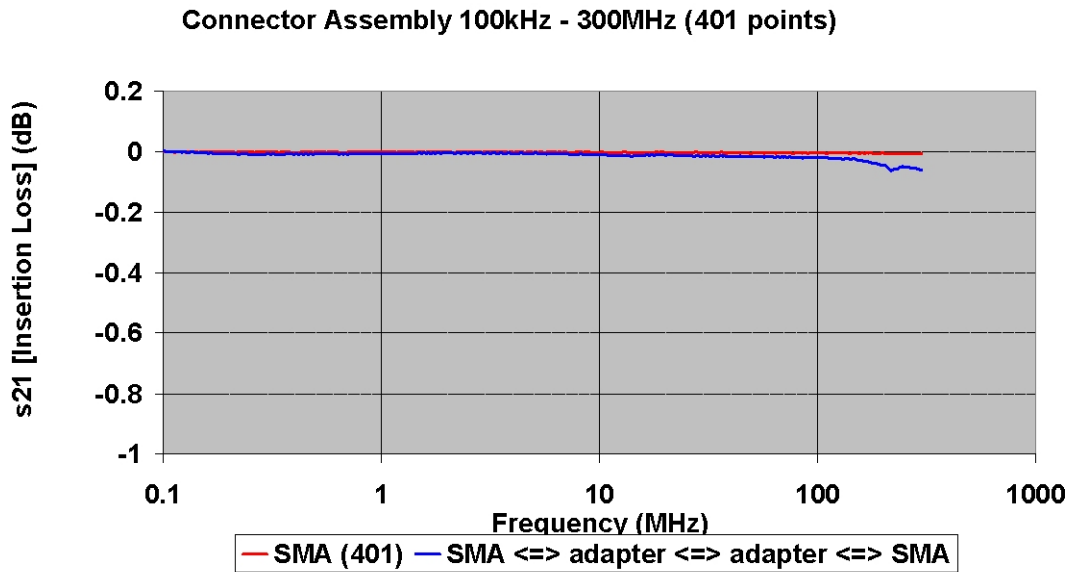


Figure 8, Cable and Probe Assembly Insertion Loss, Prior to Calibration

For core power, the FPGA draws current distributed around the center ground slug. We can model this as a single current source centered at the middle of the slug as depicted graphically in Figure 9:




Figure 9, Model FPGA vs. VNA

Figure 10. demonstrates dramatic differences between perceived transfer impedance at various measurement points, and what the FPGA actually sees for two boards, one populated with 132

conventional capacitors, ie one capacitor for each power pin as per recommendations, and another with 20 X2Y capacitors, and 28 conventional capacitors on the board underside.

**PCB #6 104 0402 outer + 28 inner vs #7 20 X2Y outer + 28 ordinary inner
100kHz - 300MHz Xilinx test Boards {port 1 to 2}**

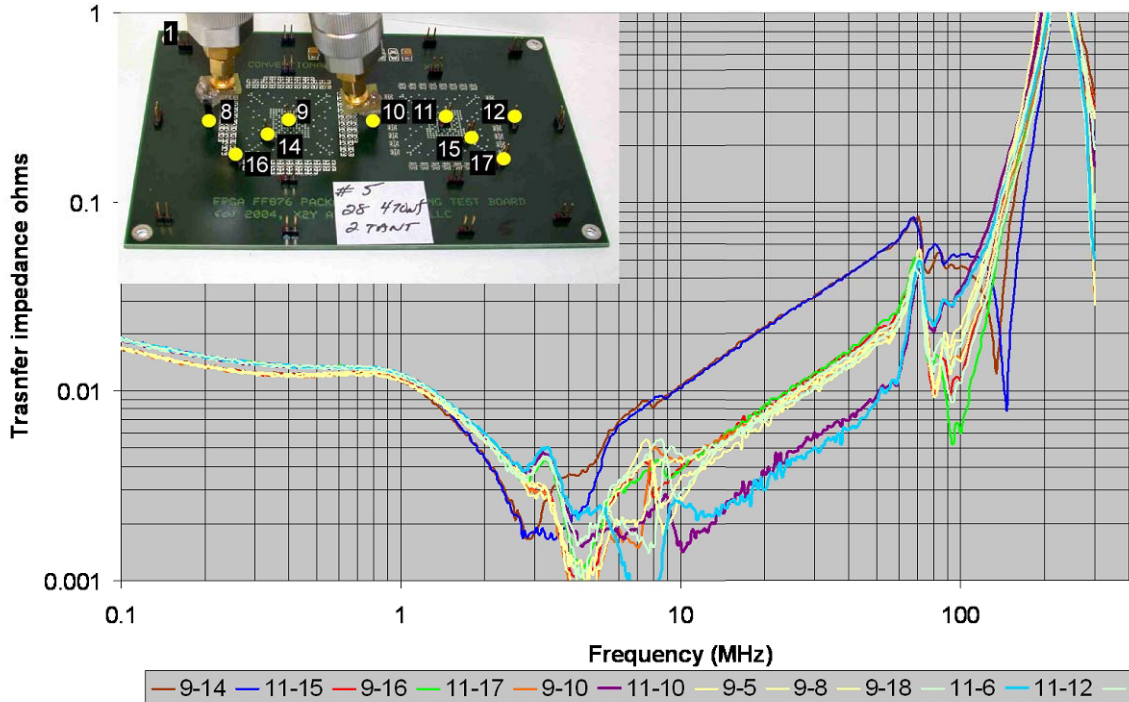


Figure 10, Insertion Loss Translated to |Z|shunt, Test Boards #6, #7

These boards demonstrate several points:

- 1) As seen comparing 9 to 14 against 9 to 16, or 11 to 15 against 11 to 17, the transfer impedance as viewed by the FPGA across the power / ground attachments is dramatically higher than the impedance viewed from the FPGA center to the perimeter capacitor radial.
- 2) As seen in measurements from points 11 to 12 or 11 to 10 on the X2Y™ test card versus 9 to 8 or 9 to 10 on the conventional test card, local transfer impedance with X2Y™ capacitors is far lower than with five times as many conventional capacitors. Note that this is a very localized effect.

Via Inductance-

Via inductance above the upper plane for either the bypass capacitors or the FPGA attachment may be calculated by summing inductance above the uppermost plane, with the inductance between the planes, per closed-form formulas by Dr. Howard Johnson, www.sigcon.com for each³:

Above the uppermost plane-

Equation 4

$$L_1 = \frac{\mu_0}{2\pi} * (H_1^2 * (2/D - 1/S))$$

Between the planes-

Equation 5

$$L_2 = u_0/2\pi * 2 * H_2 * \ln(2*S/D)$$

Setting K_1 to D/S and rearranging, in English units for capacitors on the component side-

Equation 6

$$5.08nH * ((H_1^2 * (2 - K_1) / (S * K_1)) + (2 * H_2 * \ln(2 / K_1)))$$

H_1 Via length above the uppermost plane

H_2 Plane to plane separation

D Via diameter

S Via separation, on centers

D by definition is always less than S , and is rarely more than $0.5 S$. Consequently, $2 / K_1$ is almost always > 4 . As a result, for via section that lies between the planes, changes in K_1 result in less than $1/4^{\text{th}}$ proportionate via inductance changes. However, for the portion between the plane cavity and the board surface, sensitivity to K_1 is greater than linear.

For capacitors mounted on the backside, two via segments lie outside the plane cavity, and the effective plane cavity H_2 extends from the ground plane closest to the backside to the bypassed power or ground plane closest to the FPGA. As a result, for symmetric stack-ups, via inductance increases nearly linearly with board thickness.

For H_3 equal to the bottom side via segment outside the plane cavity:

Equation 7

$$5.08nH * (((H_1^2 + H_3^2) * (2 - K_1) / (S * K_1)) + (2 * H_2 * \ln(2 / K_1)))$$

The final component of attached capacitor inductance is any surface trace, such as needed to provide for a solder mask dam. For trace width similar to height above the nearest plane 10pH/mil provides a workable estimate.

Capacitor Self Inductance-

Capacitor inductance is a strict function of device geometry. Extremely cheap on a unit cost basis, mounted, conventional SMT capacitors provide better performance as the size of the package shrinks until via spacing limits are reached. Below the 0402 package, either microvias, or fanned vias are needed, trading device gains for board level losses. Despite the same planar outline, inductance between capacitor manufacturers does vary due to height variations. We have variations of up to 100pH between manufacturers of 0603 devices. Common values for 0603 capacitors are 500pH, while 400pH is common for 0402 devices.

The X2Y™ Capacitor-

Since the early 1990s capacitors with improved geometries directed at lower inductances have been introduced to the market. Many have been disappointing, while others perform well but at very high cost. More recently the X2Y™ style capacitors have become available from multiple manufacturers in attractive values. X2Y™ capacitors have four terminals, two each on each axis. Internally, these devices form two separate but tightly balanced capacitors. In sizes from 0603 through 2220, X2Y™

capacitors are the same shape and size as their conventional counterparts, with the important distinction that they include terminals on all four sides of the part. Consequently, the manufacturing process is not exotic.

When configured as bypass capacitors, current flow cuts the path across each of the four corners of an X2Y™ capacitor. This affords: multiple current paths, each path much shorter than its conventional SMT counterpart in either normal or reverse geometry. The result is a drastic reduction in inductance for a given package size. We have measured X2Y™ 0603 capacitors in microstrip fixtures with under 60pH inductance, or roughly 1/8th that of its conventional 0603 counter part and about 1/7th of a high quality 0402.

Mounted Capacitor Inductance-

A capacitor that drops inductance from 400pH to 60pH would hardly provide useful benefit if connected through 2nH via inductance. Mounted X2Y™ capacitor inductance scales due to the use of multiple vias per capacitor.




Figure 11, Induction Loops, X-Y Plane, X2Y vs. Conventional

The appropriate arrangement with X2Y™ is to place a via adjacent to each corner of the A and B terminals as well a via adjacent to each of the G1 / G2 terminals.

As depicted in Figure 11, from an induction standpoint, a six via configuration for X2Y™ results in four, small radii induction loops in the X-Y plane. Because there is only a single via for each: the G1 and G2 pads, the relative via inductance is found from the relative AC current through each:

Equation 8

$$L_{\text{ATTACH_X2Y}} = L_{\text{VIA}}/2 + L_{\text{VIA}}/4 = L_{\text{VIA}} / 1.35 = L_{\text{VIA_PAIR}} / 2.7$$

Equation 9

$$L_{\text{X2Y_MOUNTED}} = 60\text{pH} + L_{\text{VIA_PAIR}} / 2.7$$

The number of conventional *mounted* capacitors that may be replaced by a single *mounted* X2Y™ may be found by substitution:

Equation 10

$$N = 2.7 * (L_{CONV_ESL} + L_{VIA_PAIR}) / (162pH + L_{VIA_PAIR})$$

For $L_{VIA_PAIR} \gg L_{CONV_ESL}$, N has a lower limit at 2.7.

For $L_{VIA_PAIR} \ll L_{CONV_ESL}$, N approaches $L_{CONV_ESL} / 60pH$.

Table 1, Mounted Inductance, Comparative Conventional and X2Y³

	Capacitors on Component Side							Capacitors on Back side ⁴		
H1	0.005	0.020	0.005	0.020	0.005	0.012	0.012	0.005	0.005	0.005
H2	0.014	0.003	0.003	0.001	0.001	0.038	0.038	0.014	0.003	0.001
S	0.03	0.03	0.03	0.03	0.03	0.032	0.044	0.03	0.03	0.03
D	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.01	0.01	0.01
K1 D/S	0.33	0.33	0.33	0.33	0.33	0.63	0.45	0.33	0.33	0.33
L / via pH	318	393	76	217	40	590	629	1580	1530	1540
L 0603	1052	1290	662	935	579	1500 ⁵	1760	3670	3560	3590
L 0402	952	1190	552	835	479	1400	1660	3570	3460	3490
L X2Y	267	355	117	223	90	435	531 ⁶	1250	1210	1220
Caps req'd 0603	3.9	3.6	5.6	4.2	6.5	3.4	3.3	2.9	2.9	2.9
Caps req'd 0402	3.6	3.3	4.7	3.7	5.3	3.2	3.1	2.9	2.9	2.9
Caps req'd X2Y	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0

As shown in Table 1, under practical conditions, ratios range from 2.9:1 for backside mounted capacitors to 6.5:1 for layer 2/3 plane pair arrangements using 0603 or 5.3:1 using 0402s. For any ratio above 3:1, X2Y™ bypass networks reduce total via count.

Conventional Capacitor Reduction Techniques-

Table 1 illustrates that significant component reduction is possible with 0402 capacitors versus 0603s when thin dielectric is used with planes located close to the surface as is needed for low impedance voltage rails. For either thick dielectric or cavities that begin far from the surface, 0402s offer little advantage over 0603s.

Another way to reduce conventional capacitor count is use of two vias per pad, one on either side. As shown in Table 2, this technique affords the most benefit at lowest via count penalty for capacitors mounted on the backside, ie moderate impedance rails. For thin dielectric planes close to the surface, the benefits are small, and via count nearly doubles.

³ Solder mask dam spacing has not been accounted for in this table.

⁴ Symmetric construction and 0.093” boards, w/ ground plane closest to the IC assumed.

⁵ 0603 Measured value 0603 1520pH

⁶ Measured value 530pH

Table 2, Mounted Inductance, Comparative Conventional 4 Via / Capacitor ⁷

	Capacitors on Component Side					Capacitors on Back side ⁸		
H1	0.005	0.020	0.005	0.020	0.005	0.005	0.005	0.005
H2	0.014	0.003	0.003	0.001	0.001	0.014	0.003	0.001
S	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
D	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
K1 D/S	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33
L / via pH	318	393	76	217	40	1580	1530	1540
L mounted	676	793	476	617	439	1980	1930	1940
Total Caps	0.71	0.67	0.86	0.74	0.92	0.56	0.56	0.56
Total Vias	1.42	1.34	1.73	1.48	1.84	1.11	1.12	1.11

Mutual Coupling, Arrays of Conventional Capacitors-

When multiple capacitors are placed in close proximity, fairly strong mutual coupling can occur that erodes the incremental effectiveness of additional capacitors. For conventional capacitors, the best results are obtained with capacitors placed in a single column, end to end. Measured results indicate approximately 7% deterioration in admittance for this arrangement using 0402 capacitors on 2mm grid.

If more than one column is needed alternating via polarity and alternating via / capacitor columns as shown in Figure 12 degrades admittance by less than 16% compared to scaled results using one capacitor. Avoid orienting multiple columns with the same polarity as this can easily reduce admittance by 24% or more compared to scaled results using one capacitor.




Figure 12, Mounting Options, Conventional Capacitor Arrays

⁷ Assumes negligible degradation due to same polarity mutual inductance.

⁸ Symmetric construction and 0.093” boards, w/ ground plane closest to the IC assumed.

Capacitor SRF-

Ignoring design of specific frequency notches, the bypass goal is to apply a minimum admittance across the power rails over a wide frequency band. Ultimately, the mounted inductance and the mounted inductance alone sets the high frequency impedance. Given a high frequency impedance target $|Z|$, and a mounted inductance L_{mount} , the number of capacitors required is simply:

Equation 11

$$N \geq j\omega_{\text{HF}} * L_{\text{MOUNTED}}/|Z|$$

The lowest frequency at which this network is less than $|Z|$ is:

Equation 12

$$j\omega_{\text{LF}} \geq 1/(N * C_{\text{CAP}})$$

Where:

$|Z|$ is the target impedance

$j\omega_{\text{HF}}$ is the high frequency limit.

$j\omega_{\text{LF}}$ is the low frequency limit

L_{MOUNTED} is the mounted capacitor inductance,

C_{CAP} is the device capacitance and

N is the number of capacitors

Altering the capacitance and hence the SRF has no effect on $j\omega_{\text{HF}}$ whatsoever. The only “benefit” a lower capacitance and hence higher SRF brings is the potential to increase the number of capacitors, N , required to satisfy the low frequency target.

Multiple values of capacitors such as “capacitors by the decade” are often used with the intent of realizing a broadband filter response. Too often implementations of this concept actually degrade frequency response while complicating the design. Implemented improperly, parallel anti-resonance between frequency bands can create large impedance peaks that violate impedance requirements.

The first design choice should be to avoid multiple MLCC capacitor values unless necessary. Capacitor quantity is determined by inductance needs. The largest economically attractive capacitance should be selected for the high frequency capacitor. This is often one or two values less than the maximum capacitance available in a given package and voltage rating.

Useful application of multiple capacitor values is when necessary to extend bypass impedance low frequency cut-off.

Example Networks-

One well-known FPGA application note demonstrates a network that meets a target of about 10m Ω over a broad range by using multiple value MLCC capacitors in three package sizes. The inductance is via dominated, dictating the number of via holes.

The same or better response is available at lower cost and complexity by simply using the same number of capacitors (via holes), or using far fewer X2Y™ capacitors while retaining a constant number of via holes:

Table 3, Example Decoupling Networks

Size	Pkg	ESL mounted	ESR	QTY		
				App Note	0402	X2Y
680uF	E box	2.8nH	570mΩ	2 ⁹	-	-
2.2uF	0805	2.0nH	20mΩ	7	-	4
220nF	0603	1.8nH	60mΩ	13	-	-
22nF	0402	1.5nH	200mΩ	26	-	-
470nF	0402	1.5nH	10mΩ	-	46	-
940nF	X2Y 0603	0.53nH	5mΩ	-	-	14
Vias				92	92	92

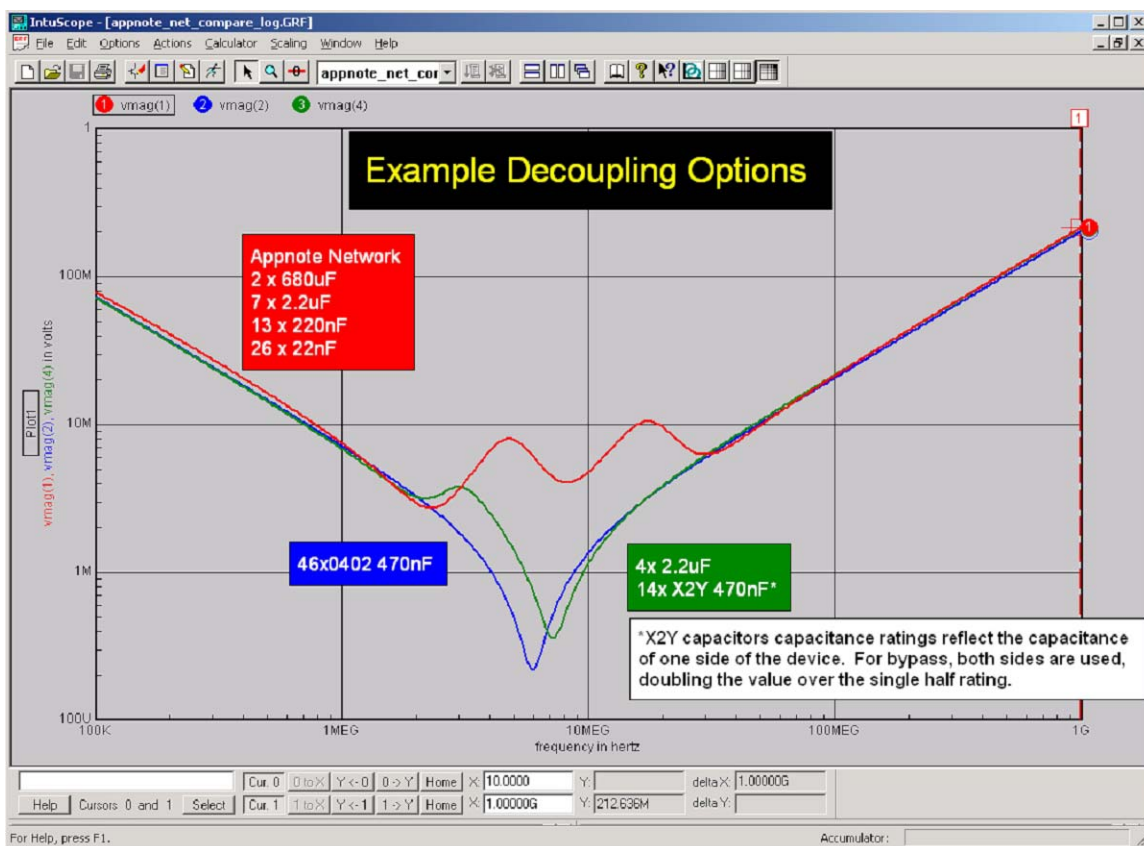


Figure 13, Performance w/o Planes, Example Decoupling Networks

While all three options meet the intent of the application note, the multiple value conventional capacitors specified provide no advantage over either 0402 only network, or a much lower capacitor count network using X2Y™ capacitors¹⁰.

Because the X2Y™ network eliminates so many capacitors, a few conventional 2.2uF capacitors are necessary to extend low frequency response as shown.

⁹ Due to high ESR, these capacitors do not materially affect impedance above 40KHz.

¹⁰ As of this writing 470nF circuit 1 rated 0603 X2Y capacitors are scheduled for general availability Q1 2005.

New Package Architectures-

The Virtex4 packaging architecture departs from the center ground slug and surrounding power rings of VirtexII / Pro™. For Virtex4™, Xilinx redistributes ground and power balls throughout the package, providing multiple benefits while somewhat complicating analysis.

An immediate benefit is that the mean radius from the die center to any given power or ground connection on the PCB is greatly increased, which in-turn reduces the R2/R1 ratio from equation 1. With Virtex4™, for any given plane separation, spreading inductance is substantially lower than for a comparable VirtexII/Pro™ counterpart.

Figure 14 illustrates the power and ground assignments for the FX100 in an FF1152 package:




Figure 14, Xilinx Virtex4 FF1152 Power / Ground Layout

Note that logic ground (green) and I/O power (yellow) distribute very evenly throughout the package, and that core power extends well away from the device center.

High-speed signal integrity benefits from the close proximity of signal ground to all I/Os. And finally, Rocket I/O™ pads are located with their dedicated power and grounds at the package periphery. All of these features combine to yield marked reductions in system level packaged impedance.

Bypass Network Synthesis Procedure-

We now present a step-by-step procedure for developing a bypass network on an IC by IC basis:

1. Assign a target impedance to each supply as seen at the package power / ground interface.

2. Assign power / ground plane pair order from top to bottom in order of increasing target impedance.
3. Allocate no more than 25% of the impedance budget at the high frequency cut-off (50 MHz) to each: device via attachment and spreading inductance.
4. Determine maximum plane separation allowable to remain under 25% budget. For VirtexII™ and VirtexII / Pro™ spreadsheet analysis is adequate. For more complex packages such as Virtex4™, a 3D simulation can be performed, or analysis data / services maybe purchased from Teraspeed, LLC. Teraspeed utilizes CST Microwave Studio™ to extract extremely accurate package / PCB data.
5. Checkpoint- Confirm that the results from step 4 are physically realizable.
6. From package data, data from step 4, and intended via drill size, determine the maximum allowable upper plane depth in the PCB stack-up.
7. Checkpoint- Confirm that the stack-up requirements are realizable.
8. Assign the detailed stack-up.
9. Compute plane spreading and FPGA attachment inductances from stack-up in step 8.
10. Subtract consumed inductance from budgeted inductance. This is the maximum inductance allowable from the mounted bypass capacitor array.
11. Select high frequency decoupling capacitor type: conventional, or low inductance.
12. Based on selected decoupling capacitor type, determine quantity of capacitors required to meet target inductance.
13. Checkpoint- Confirm that the required capacitor count is physically realizable. Iterate steps 3-12 until a physically realizable solution is found, or redefine device / packaging requirements.
14. Determine low-frequency bypass cut-off. The required cut-off frequency depends on the particulars of the VRM and bulk capacitors used. For a Butterworth response,

Equation 13

$$C_{MLCC} \geq 2 * L_{CBULK_MOUNTED} / (ESR_{CBULK} + ESR_{MLCC})^2$$

15. Checkpoint- Determine whether a single value MLCC provides sufficient capacitance to meet the low-frequency cut-off. If so, assign capacitance equal or greater to requirement, and exit.
16. Determine capacitance shortage.
17. Select lowest inductance capacitor package that can make up the shortfall in a reasonable number of devices. Balance trade-offs between quantity, unit cost, and mounted inductance. Models can be simulated in SPICE, built-up with power integrity tools, or assistance is available from Teraspeed, LLC.
18. Checkpoint- Confirm that antiresonant peaking between the capacitor array selected in step 17, and the high frequency capacitor array remains below the target impedance. If so, exit.
19. Correct anti-resonant peaking by any of the following methods:
 - Increase the count of high frequency capacitors.
 - Increase the count of capacitors from step 17.
 - Reallocate high frequency capacitors to an intermediate value capacitor with mounted SRF approximately equal to antiresonant peak.

Equation 14

$$C = 1 / (j \omega_{ANTIRESONANT_PEAK}^2 * L_{MOUNTED})$$

20. Simulate solution from 19, and iterate as required.

Conclusions-

High performance FPGA bypass networks consist of all elements:

- FPGA attachment vias
- Power / ground planes
- Bypass capacitor attachment vias
- Bypass capacitors

A properly designed, high-performance power distribution system: accounts for, and budgets inductance to each of these elements.

Excessive via attachment, and/or plane spreading inductance between an FPGA and the bypass network defeats even “TIG welded planes”.

X2Y™ capacitors make possible dramatic component count reduction with like or better performance, and often via count reduction as well. The lower the impedance required, the better that low inductance capacitors perform compared to conventional capacitors.

Backside capacitor mounting under an FPGA can mitigate the need for thin dielectric planes under certain conditions.

Multiple ceramic capacitor values should be used only when necessary to extend the low frequency reach of the bypass network.

Virtex4™ packaging represents a major advancement in PDS to FPGA interface.

Need expert help?

Teraspeed Consulting Group, LLC.

www.teraspeed.com

(401) 284-1827

¹ Resnick, Halliday “Physics”, 1977, John Wiley and Sons

² Zhiping Yang, Jim Zhao, Sergio Camerlo, and Jiayuan Fang, Impact and Modeling of Anti-Pads on PowerDelivery System, Proceedings of the 12th Topical Meeting on Electrical Performance of Electronic Packaging, pp. 117-120, Oct. 2003.

³ Dr. Howard Johnson, “High-Speed Signal Propagation: Advanced Black Magic”, pp 259, Prentice Hall, 2003.

Xilinx™, VirtexII™, VirtexII Pro™, Virtex4™, Rocket I/O™ trademarks, Xilinx, Inc.

CST Microwave Studio™, trademark Mentor Graphics, Inc.

X2Y™ trademark, X2Y Attenuators LLC