

This is the second in a series of tech online presentations about Xilinx FPGA components and their packaging.

This presentation is best viewed using the Adobe reader at 100% viewing size on a screen of at least 1024x768 resolution.

Peter, thank you for that introduction, and to Xilinx for giving me this opportunity to address you on an important topic.

BGA Package Crosstalk in Depth Ground Ball Placement

**Xilinx® Virtex-4™ FPGA and
Altera® Stratix® II FPGA**

Prepared for Xilinx

Tech On-Line

June 7, 2005

By

Dr. Howard Johnson

Theory, examples, giant scale models, and
measured lab results of crosstalk in BGA packages

© 2005 Signal Consulting, Inc. All rights reserved.
Xilinx is a registered trademark, and Virtex-4 a trademark, of Xilinx, Inc.
Altera and Stratix are registered trademarks of Altera Corporation.
All other company and product names may be trademarks of their respective companies.



This talk is about crosstalk, specifically, the crosstalk that occurs underneath a BGA package, in the spaces between the balls and your pc board. This form of crosstalk is often called SSO (simultaneously switching output noise) or ground bounce.

Crosstalk in this situation turns out to be closely related to the geometry of those balls, and geometry and positioning of your the power and ground vias. We will have a chance a little later to compare two popular FPGA layouts to discover what differences there might be between them.

In all cases the crosstalk associated with a BGA package is caused predominantly by magnetic fields, so let's start there, with a brief review of how magnetic crosstalk, also called inductive crosstalk, operates.

Outline of Talk

SSN (Crosstalk)

- BGA packages
- Pcb vias

Ball geometry

- Comparison of Xilinx and Altera power/ground pinouts



What causes crosstalk in this situation?

- Magnetic fields



Here I have set up a simple experiment involving a conductive loop, around which I can place a high-frequency current. The loop is approximately 8" in diameter, formed from 1/2" copper tubing. In the schematic view (right side of figure) you can see that the right side is also split, and wired to a signal source. I can excite the tube with a sine wave of 1 MHz. Current enters the loop at A, travels around the circle counterclockwise, through the 50 ohm load, and exits at B, returning to the source through the coaxial shield. You can think of the top half of the loop as the "signal wire" and the bottom as the "return wire".

Magnetic Field Generator



Currents circulating in the loop create a magnetic field

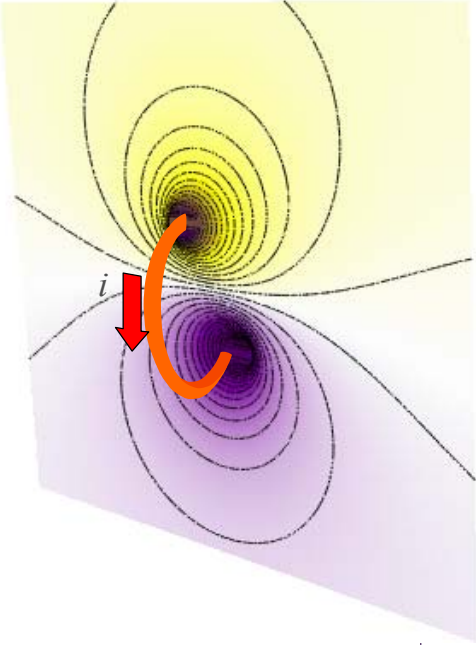
Equal amounts of current pass through the top (signal) and bottom (return)



Current flowing through the loop creates a magnetic field. The field surrounds the loop. Michael Faraday, who discovered the principle of magnetic induction, saw in his mind's eye invisible lines of magnetic force, permeating all of space, emanating from the center of the loop. The existence of these magnetic fields causes crosstalk between circuits. This form of inductive crosstalk, or magnetic-field crosstalk, is extremely common in high-speed digital systems.

Magnetic Fields

Michael Faraday imagined invisible lines of magnetic force emanating from the loop, permeating all of space




To investigate the magnetic fields surrounding this loop I have arranged a probe circuit much like the circuits on your printed circuit board.

This circuit represents a typical digital driver stuck low, whose signal passes through a short length of wire to a receiver. The driver and receiver share a common ground connection.

On the right side of the probe I feed the received signal into a coaxial cable, and on to the scope so we can observe the interference received by this circuit.


I am using a Tektronix model CSA7404 digitizing oscilloscope with an input bandwidth of 4 GHz.

Pcb Circuits Pick up Magnetic Fields



If I flip the probe over, you can see on the back side of my probe the actual circuit within. It is just a simple loop of wire. Any simple loop of wire arranged like this forms a magnetic pickup.

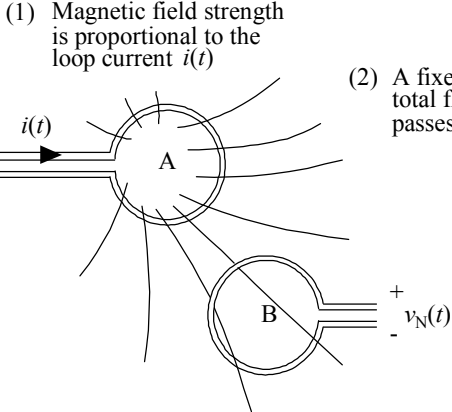
My Magnetic Field Probe



Consists of one simple loop of wire

When I hold the probe near my big conductive loop, the changing magnetic fields surrounding that loop create measurable voltages. This exact sequence of events -- a changing current creating changing magnetic fields, those fields passing through space to penetrate another circuit, and the fields at that point inducing crosstalk in the victim circuit, describes precisely how many digital crosstalk problems are caused.

Origins of Magnetic Crosstalk



- (1) Magnetic field strength is proportional to the loop current $i(t)$
- (2) A fixed fraction of the total flux from loop A passes through loop B
- (3) The rate of change of flux in loop B is proportional to the rate of change of current $i(t)$
- (4) Faraday's law states that the induced voltage is proportional to the rate of change of flux in loop B

$$v_N(t) = L_M \frac{di}{dt}$$



Here I have mounted a big graphic alongside the conductive loop to show the magnetic field patterns. The current flowing in the loop follows a square wave pattern.

In the picture I am holding the pickup near the center of the loop. The changing magnetic field from the loop enters my pickup from left side, causing crosstalk (see inset).

[NOTE: there are some Morey patterns in the digital version of this picture that appear to distort the shape of the field near the center of the loop. If you could see the graphic in full resolution, the field lines at the center actually flow horizontally through the center, as you would expect.]

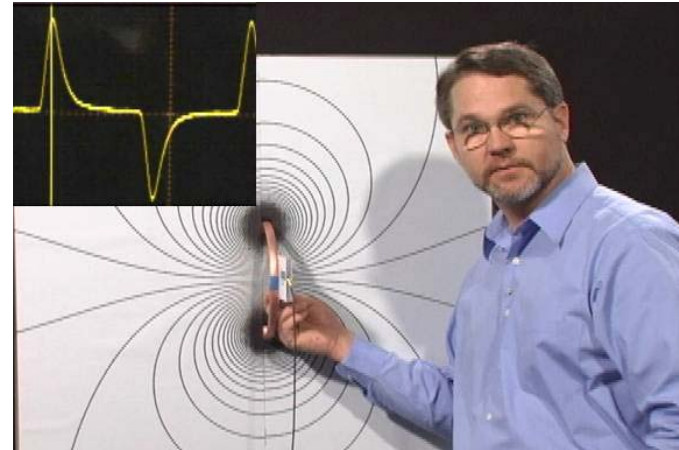
Each rising or falling edge of current causes a pulse response in the probe. You can see the probe response first goes up (on the rising edge) and then down (on the falling edge). This response follows the same form as the crosstalk in your digital circuits.

I would like to show you that the polarity of the magnetic field makes a difference. If I hold the pickup in the same place, but flip it upside down, then from the perspective of the pickup the magnetic fields will enter from the "other side". This change should reverse the polarity of magnetic interference, and thus crosstalk.

As I flip the pickup, watch the crosstalk waveform to see if it changes polarity. Right now, the crosstalk waveform pulses up, then down.

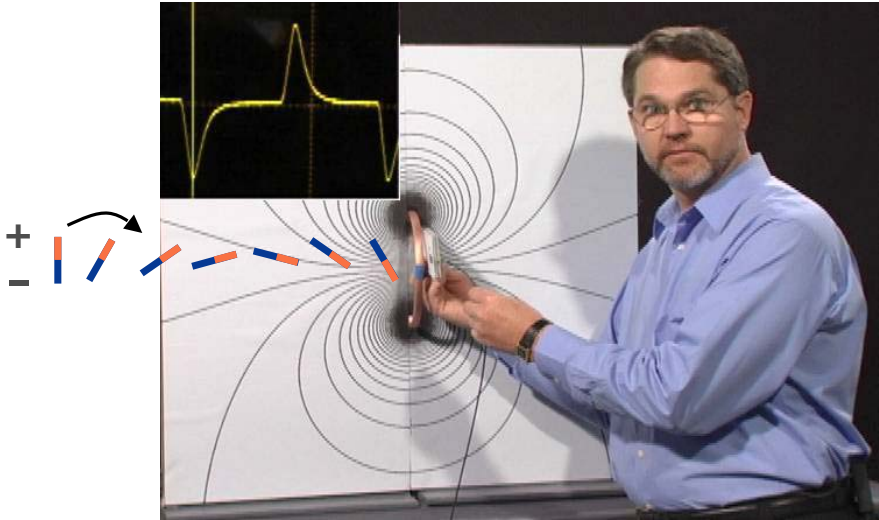
[NOTE: in the movie version of this talk you can watch the probe flip over, in the next slide I'll just have to draw in a graphic indicating how the flip is accomplished.]

Fast-Changing Magnetic Fields Cause Crosstalk



With the pickup flipped over, the crosstalk inverts polarity (i.e., now it goes *down* first and then *up*). This is one way you can prove that the pickup in this example is magnetic in nature, not capacitive. If we were experiencing capacitive crosstalk, the polarity would remain the same regardless of how I flipped the pickup circuit.

Proving Crosstalk Is Inductive

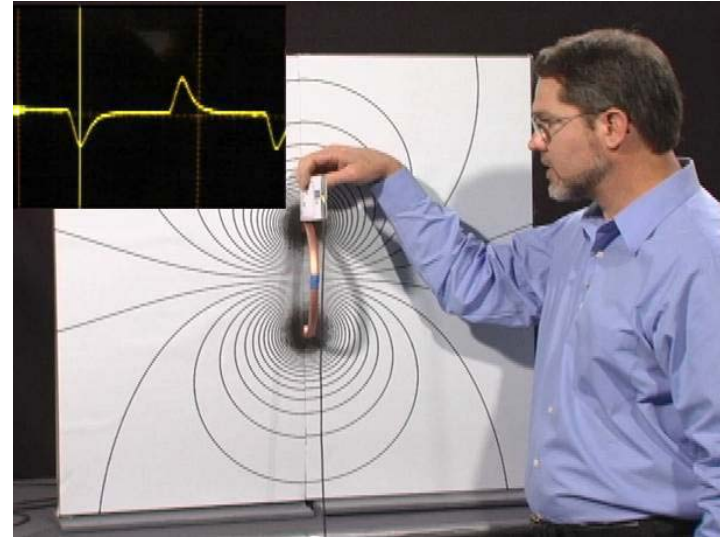


Here is another way I can show the pickup is magnetic. If I place the pickup near this top half of the loop, I get a certain level of crosstalk.

What do you suppose will happen if I hold it near the bottom half?

[NOTE: another nice aspect of the movie is that you can watch the scope waveforms change, in real time, as the probe moves about.]

Crosstalk Near Signal Wire

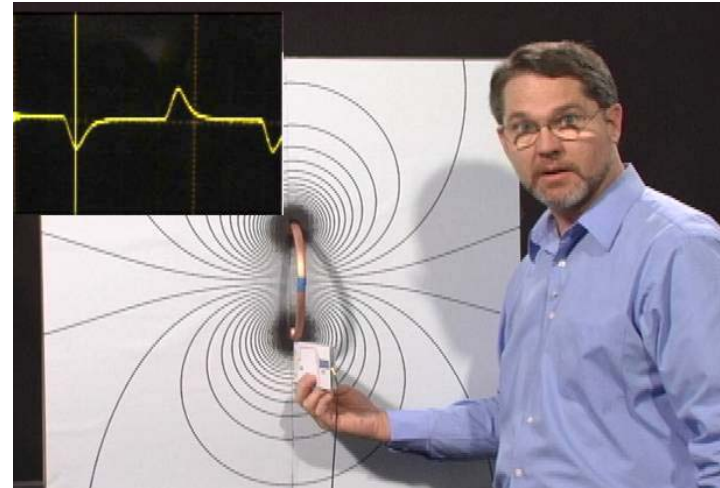


Near the bottom half, I get the same level of crosstalk. That is a very important conclusion. Crosstalk occurs not only near signal wires, but near return wires. In a practical pc board the returns can be ground or even power connections.

Capacitive crosstalk occurs only near signal wires, because they carry fast-changing voltages. The return wires are electrically "quiet", and so create no capacitive crosstalk.

Inductive crosstalk is different. Because both signal and return connections carry current, signal and return pathways are equally effective at delivering inductive crosstalk. In a high-speed, low impedance circuit, you must therefore think both about the placement of signal connections, and also return-path connections.

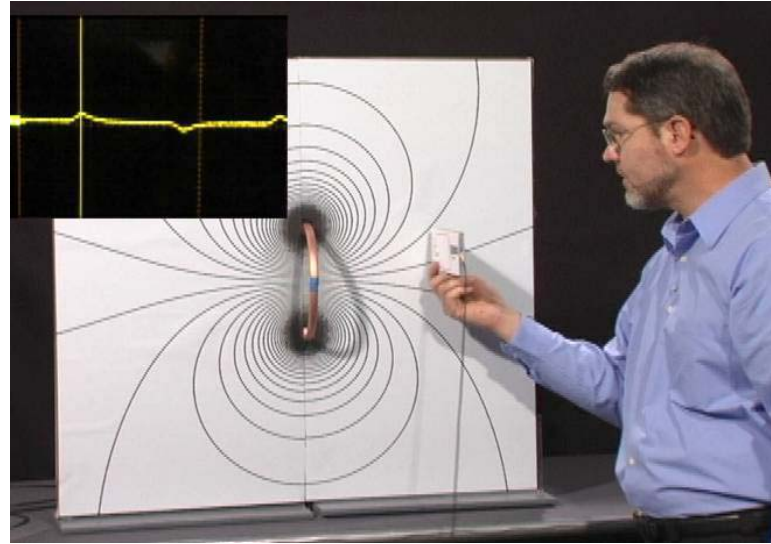
Equal Amount of Crosstalk Near the Return Wire



Let's try moving further away from the loop.

As I move the pickup, the amplitude of the interference corresponds to the intensity of the magnetic field at any point. Near the loop where the lines are most intense we pick up the most signal... at this point further away here we get less. That illustrates a very important principle: moving away from the aggressive circuit dramatically reduces crosstalk.

Moving Further Away




Here is the schematic view of our noise-generating circuit.

Right now current passes across the top half of the loop, through the load, then returns on the bottom half. Next I will change the path of current to see what influence that has on crosstalk.

Current passes across the top half of the loop, through the load, then returns on the bottom half.

Path of Current




I have clipped our probe in place so we can continuously observe the crosstalk. The loop is on the left, the probe is clipped on a wooden stand a few inches to the right of the loop. I have increased the gain a bit to make the crosstalk appear full-sized for this demonstration.

I have in my hand a new wire. I shall next add this new wire to the apparatus, changing the path of the returning signal current. My new wire is the size of a half-loop.

Observe the size of the crosstalk now, and then see what happens when I add the wire.

Try an Alternate Return Pathway



Here I have connected the new wire in parallel with the return half of the loop, going from the ground connection at the source to the ground side of the 50-ohm load. As you can see, crosstalk is markedly reduced.

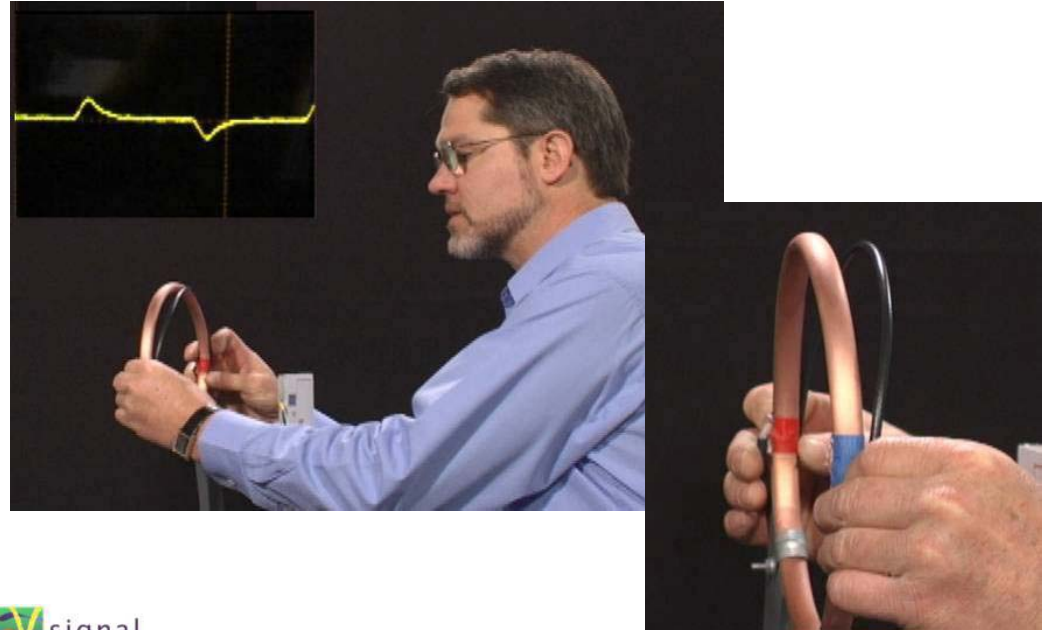
This new connection provides an alternate pathway for returning signal currents. The signal pathway is unchanged, but, after passing through the load, the returning current can choose whether it wishes to return to the source on the old path, along the bottom of the loop, or along the new path, running on the new wire adjacent to, but not touching, the signal pathway.

Taking the signal path in conjunction with the bottom pathway comprises a large loop, a large inductance. On the other hand, if the current goes out and then reverses direction, returning on the new conductor, it traverses an overall path with a smaller exposed loop area, a smaller inductance.

Given this choice, current at high speeds always chooses the path of least inductance. Any time you provide an alternate return pathway, closer to the path of outbound current, most of the returning signal current is swept up into the new path. Your circuit now produces a pattern of outbound current creating magnetic fields, plus a pattern of equal and opposite return current creating equal and opposite magnetic fields, both field patterns superimposed almost directly on top of each other, in almost perfect cancellation. The resulting crosstalk is dramatically reduced.

For the next experiment, instead of using a simple wire as a return path, why not try something more effective, like a solid plane?


Applying Wire



The solid plane, used as a conduit for returning signal current, is the most effective object you can use in a digital pcb to mitigate crosstalk. This plane is touching from ground to ground, in parallel with the existing return connection.

You can see in the detail view that the plane shorts out the bottom half of the loop, modifying the return path, but it is not shorting out the signal wire. It provides a wide, low-inductance path for returning signal current, reducing crosstalk to levels below the sensitivity of my scope.

Applying a Copper Plane



Let me be clear that these remarks apply to the operation of low-impedance circuits, like 50-ohm drivers and receivers. Any time you bring such circuits out of the plane of the pcb, exposing them to free space, going through a package pin, a connector, or a component body, in those cases the magnetic crosstalk almost always overwhelms the capacitive crosstalk.

Why do we care about the difference between magnetic and capacitive crosstalk? Because it helps you understand how crosstalk functions, and therefore how to mitigate crosstalk.

I am aware that many of us, when facing a crosstalk problem, tend to think first about capacitive crosstalk. I'd like to address the question of why we have this bias in our thinking.

Our electrical engineering curricula was developed during the tube era. Tubes are very high-impedance devices. In the Eniac computer, for example, the circuits used huge voltages (hundreds of volts) and very small currents (microamps). Such circuits produces huge electric fields, tiny magnetic fields.

Given such large voltages, even the tiniest amount of parasitic capacitance from one plate to the next grid causes enough crosstalk to destroy the operation of your computer. As a result, we have passed down to use from generations of earlier engineers the warning, "Watch out for capacitive crosstalk".

We no longer make high-impedance circuits. We make low-impedance circuits, 50-ohm stuff. Little tiny voltages—huge currents. Any time you have a circuit with an impedance less than the "impedance of free space", a value set by basic physics at 377 ohms,

Digital PCB Crosstalk Is Not Capacitive



slide 3.17

most of the electromagnetic energy surrounding that circuit will be in the magnetic field mode, not the electric-field mode. That is what makes magnetic coupling so important in digital pc-boards.

Sure, capacitance still matters for high-impedance work. For example, on-chip your circuit impedances are quite high, some few thousands of ohms. There, capacitance crosstalk matters. But for any digital pc-board crosstalk problem you face, any time you expose your low-impedance circuits to the air through a package pin, a connector, or a component body, think first about inductive crosstalk. Identify the path of outbound current and the path of returning signal current. Causing the return current to flow home closer to the outbound path improves crosstalk.

The concept of return current explains why connectors have so many ground pins.

The figure illustrates a connector thinly populated with ground pins. Several signals share a common ground return, merging their currents at that location, causing crosstalk.

If you add ground pins to this connectors, then every signal current passing through the connector may choose how best to return.

For example, if you ground the top pin in the figure (pin 1), then current associated with pin 2 (the next pin down) can choose to return in two ways:


(1) On the old ground pin, making a big loop, a big inductance.

(2) On the new ground pin, adjacent to the outbound signal pin, making a smaller loop, a smaller inductance.

Given that choice, returning signal will flow predominantly on the lowest-inductance path, right next to the outbound signal pin. In that case the magnetic fields generated by the outbound signal current are cancelled, to a large degree, by equal and opposite currents flowing on the adjacent ground pin.

Interleaving a number of ground pins along the body of your connector naturally encourages the current associated with each individual signal to flow on its nearest ground, disengaging the signaling loops from each other, reducing crosstalk.

Why Connectors Have Ground Pins



This ground pin carries all three returning signal currents

Our theoretical review is now complete. Let's go make some measurements.

I have constructed a working physical model we can use to actually measure inductive crosstalk. The model can directly compare the performance of alternative ground-ball layouts.

This model makes use of the the 3-D rule of scaling. This rule governs the behavior of all connectors, vias, and other components generally classified as "lossless" (i.e., they do not involve resistive heating or significant amounts of signal radiation). First laid out in mathematical form by Maxwell in his seminal work, "A Treatise on Electricity and Magnetism", 1873, the 3-D rule of scaling says that any lossless circuit, when scaled down in physical size by a factor of 10 in all three dimensions, produces a new circuit that works precisely the same as the old circuit, only 10 times faster. This rule is the key to understanding why chips go so fast – the smaller you make them the faster they work.

In our case I shall scale in the other direction, making a model of a BGA package 57 times bigger than real life, which when operated displays all the properties of a real, working BGA model only 57 times slower than in real life. In this model, all the inductances and capacitances associated with high-frequency behavior scale precisely.

I've chosen to use a scale model for several reasons. First, so that you can see it, and so that I can get my hands into it to probe the current flowing in individual balls. Next, because it slows down all the effects to the point where we can easily instrument the model to make precise measurements.

BGA Model

Demonstrate crosstalk
Compare alternative ground-ball layouts

This model makes use of the
➤ 3-D rule of scaling



Let me point out that this physical model is every bit as good as the model used by a 3-D field solver. The field solver simply scales the problem into the software domain, using bits to represent physical quantities instead of the quantities themselves. Had I used a field solver to compute a problem of this complexity, with some significant dimensions only a few mils across and others as large as 23 inches, it would have taken a simulation grid of at least 23,000 points in each dimension, x, y, and z to yield any useful answers. A grid that size takes quite a while to compute; mine renders the answers instantaneously.

This model comprises a total of 100 balls, representing the top left corner of a Virtex 4 FPGA in the FF1148 package package.

BGA Crosstalk Model



Here you see part of the interior construction of the model. All the parts were hand-machined by my model maker, Bernie Hosey.

Interior Construction



Model, partially assembled



Modelmaker
Bernie Hosey



The BGA package section appears at the top of this detailed view. The BGA balls sit on round pads. The pads connect to vias, and the vias lead down to the bottom of the model, representing a pc-board ground plane buried deep within a typical PCB. I left the vias exposed so you can see how they are arranged.


Detail Showing Balls and Vias



The finished assembly, shown here sitting up on its side, includes numerous SMA connections for monitoring signals within the model.

The next slide shows a schematic view of how the model works.

Back Side of PCB




The figure illustrates a sample of four vias taken from the model.

When operating this model, I imagine switch C to be closed. I then connect a signal source at the bottom right corner of the diagram, injecting current into via F. The source forces current through switch C, returning on some combination of the ground vias, only one of which is show here. The path of current is shown as a red dotted line.

Current traversing that path fills the entire space between the BGA substrate and the pcb with fast-changing magnetic fields, causing magnetic crosstalk for the victim circuits D and E. I can observe the resulting crosstalk waveforms at these locations using my Tektronix digitizing scope.

Jamming current in this way from the outside through the model is not an artificial mechanism; it is exactly what happens in a real application. Inside a real BGA package at location C there is an IO switch. When the switch closes, the source of power lies outside the chip, stored somewhere in the power supply or load circuit. That source of power drives current through switch C as long as the load and ground voltages differ. The chip is not a source of power. It is just a passive switch. The real source of power comes from outside the chip.

My experiment leaves the IO circuits within the BGA package turned on (i.e., connected to ground) all the time. My external source controls the current. Since the switches all remain on, all the time, I can model the BGA package itself simply as



a big hunk of solid metal. The top side of every ball connects directly to that metal plate.

This configuration simulates an ideal BGA package with perfect power and ground bypassing and zero crosstalk internal to the BGA routing. The model just looks at the inductive crosstalk pickup up by the balls and vias underneath the package.

Six of the via locations in my model represent fixed ground vias. These I've marked with green dots. Alternately, they could represent power balls. Assuming good bypassing internal to the package, either power or ground balls serve equally well as return connections.

Six Via Locations Represent Fixed Grounds




This is a bottom view of the model. This detailed view shows how one via attaches to the simulated pcb. That big screw head on the right of the figure holds the via in place. The screw sits on an insulating washer that prevents the via from shorting directly to the metal plate.

This particular via comes equipped with two switches. Right now, the switch on the right is turned "on". Turning this switch "on" connects this particular via to the pcb ground. With this switch "off" you would have no connection. In that way the switches allow you to dynamically change the configuration of return connections in the model.

These tiny switches have negligible series inductance compared to the natural inductance of the ground ball and its via. I wish I could do this to a real BGA package, but it would require a switch only 5 thousandths of an inch across.

Other Grounds are Configurable

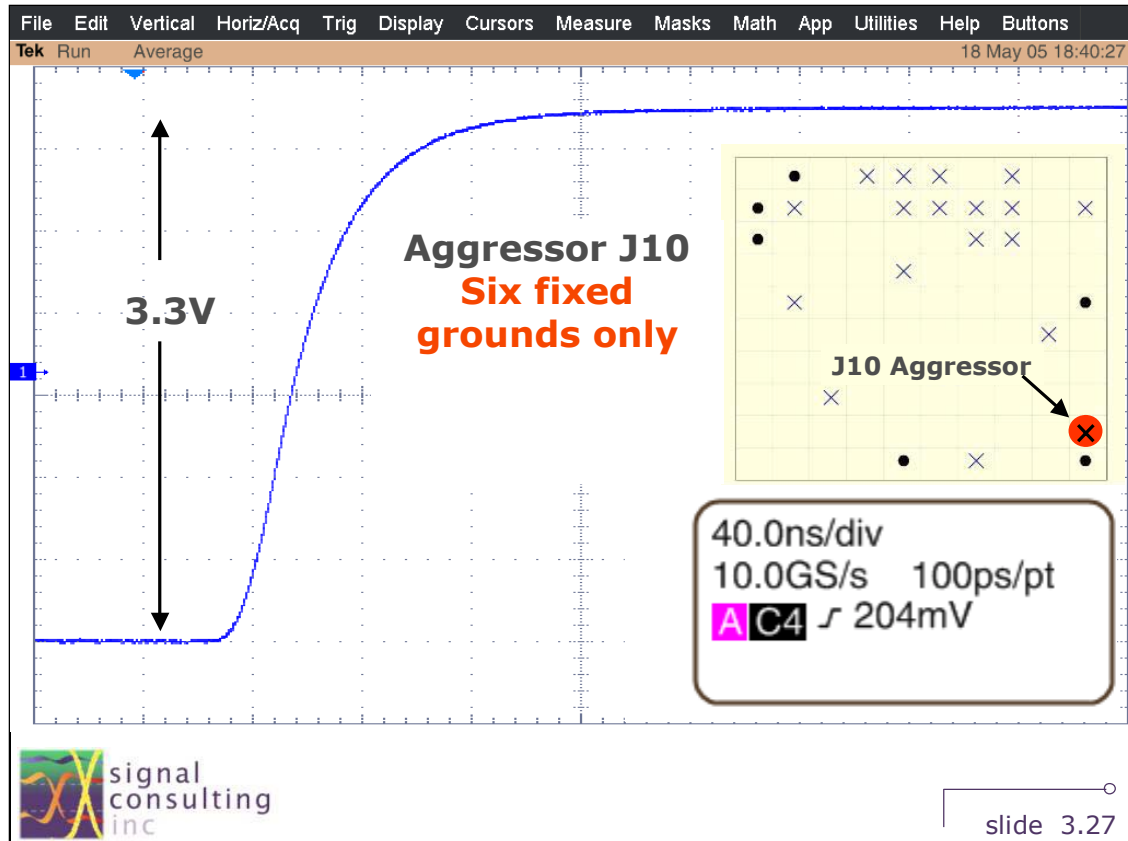


The first thing let's do is investigate where return currents flow in this experiment.

I'll drive ball J10 down near the bottom right corner of the apparatus. The ball pattern of the model (top view) is illustrated on the right side of this figure.

At the moment the only ground balls in use are the six positions all marked in the figure with black dots. (A2, B1, C1, E10, K5, K10). The locations marked "x" represent SMA connections where we can sample the crosstalk. As you can see, I populated some areas thoroughly with SMA connections, and some more sparsely, simply to save effort in building the model.

This picture shows what the source signal looks like when connected to drive a 50-ohm load. The source risetime is about 57 ns, which at our scale translates to a standardized 1-ns rise/fall time (and 3.3-volt signal amplitude) used for all my time-domain experiments with this model.



To probe the currents flowing within the model I have built this tiny current probe. It sits on the end of a long stick.

The probe is just a coil of wire, passing the received signal out through a coaxial cable to the scope. If we place this probe next to a return ball, it reads a signal proportional to the current flowing in that individual ball. (Wish I could do this on a real BGA).


A Tiny Current Probe



I place the probe near a ball to measure the current flowing on that individual connection. The top trace is a record of the source signal. The bottom trace shows the profile of current flowing at position E10. Those of you used to magnetic field measurements know that the B-field probe responds to changes in magnetic intensity, and so when the current steps from "off" to "on" the probe responds with a pulse waveform.

Watch out for the pin numbers in this view, since we are looking at the back side of the pcb the pins go right to left, 1-10, in this view.

Measuring Current





This chart shows the relative distribution of return current on the six return pins in play during this experiment.

The source pin J10 shows positive current, flowing into the model. Return paths E10 and K10 show negative current, returning from the model.

Here you can see there is more return current flowing at position K10, for instance, than there is at position E10.

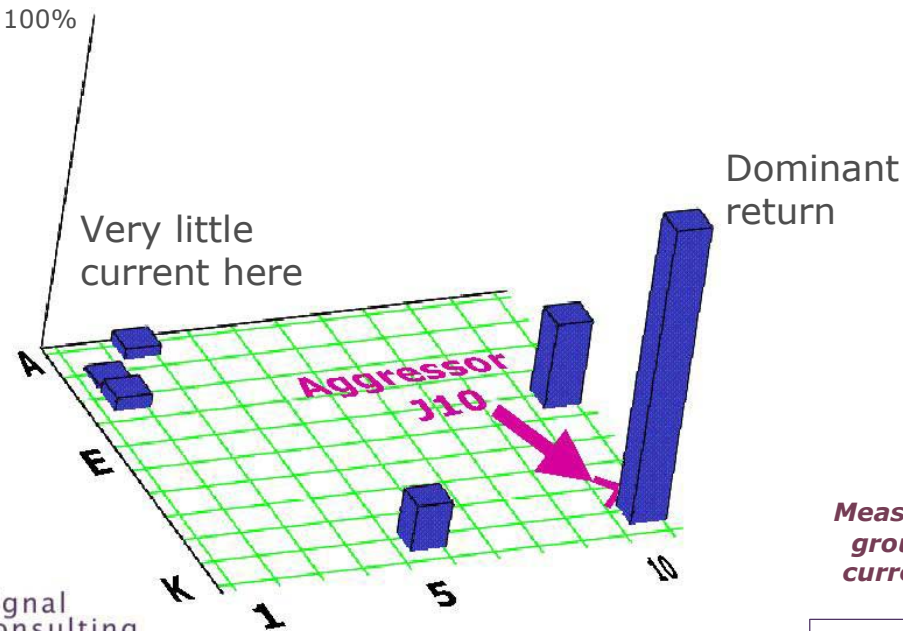
Aggressor J10
Six fixed grounds only



Most of the returning signal current concentrates on the nearest return pins. Position K10, being the closest return, acts as the dominant return path. This confines current to a region near the aggressor, limiting crosstalk elsewhere.

Next, let me show you what happens if you drive the model from a different position.


Full Distribution of Return Currents



Driving position B7, the returning signal current splits up, appearing in substantial amounts all over the chip. Even the pins way down on row K carry quite a bit of the return current. In this configuration, return current does not fall steeply as you move away from the aggressor. That means that the crosstalk associated with this geometry won't fall very steeply with distance, either, and that is the main problem with widely-separated returns.

In the next view, instead of ground currents, let's look at crosstalk for several locations near B7.

Distribution of Return Currents from Alternate Location



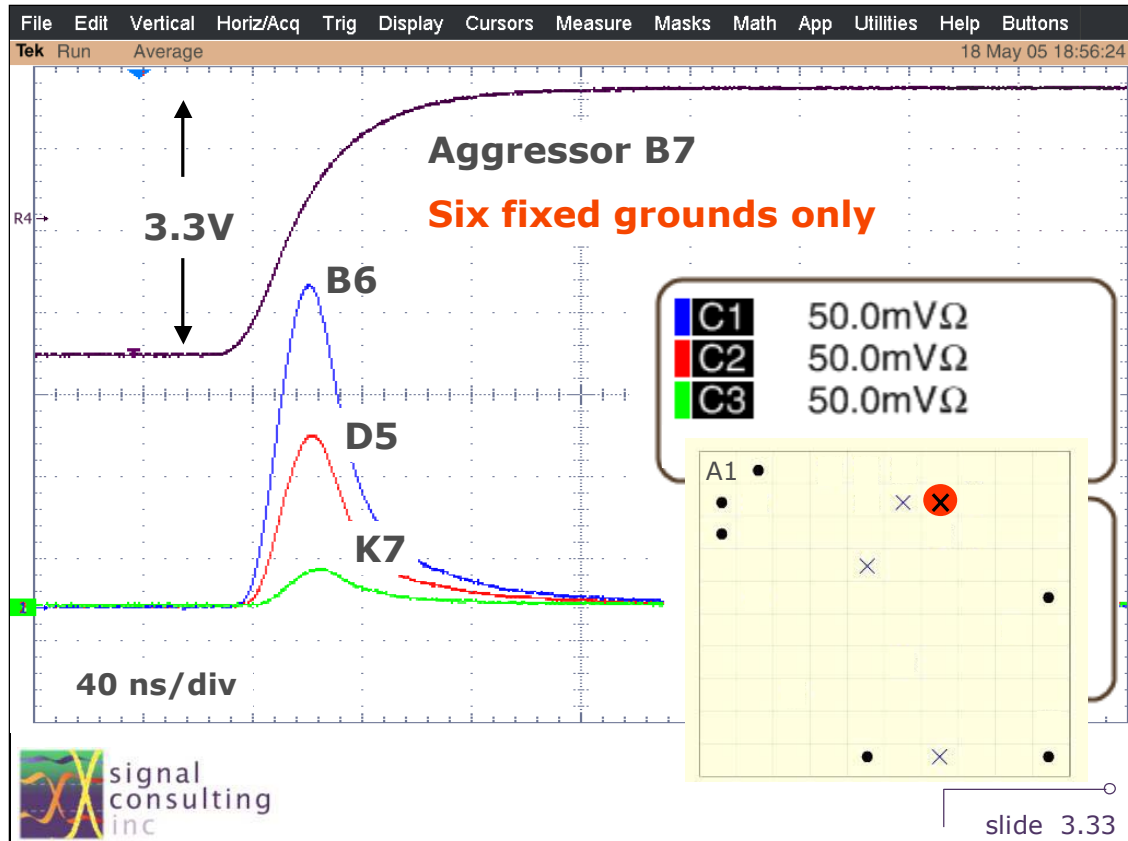
Here I am still driving position B7, marked in red on the pin diagram.

The scope monitors positions: B6, D5 and K7, all marked with "x".

The top trace shows crosstalk measured at ball B6, adjacent to B7.

The next trace shows a more remote ball, D5, and finally a very remote ball on the bottom row, K7.

Every time the aggressor experienced a surge of changing current, we get a pulse of crosstalk at each of these victim locations.

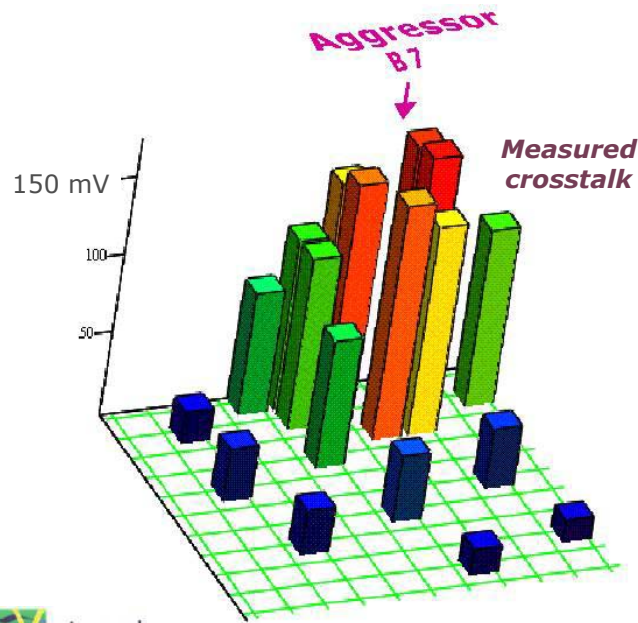


Here is the complete record of crosstalk from B7 to every other position in our test setup.

For this measurement my input stimulus is a 6 MHz sine wave, at a current amplitude of 20 mA (40 mA peak-to-peak). At our 57:1 model scale this corresponds to a real-life frequency of 350 MHz, roughly comparable to the meat of the spectrum associated with 1-ns rising and falling edges coming from 2.5-V logic driving end-terminated 50-ohm transmission lines. I like to use sine waves for these sorts of test because the sine wave avoids having to define the precise profile of your rising edge, as that affects in a subtle way the precise amount of crosstalk measured.

Near the aggressor the pair-to-pair crosstalk is about 150 mV p-p. As you move away from the aggressor the crosstalk falls off, but not very precipitously. It is important that crosstalk fall to very low levels in the outlying sections of your chip. With 500 IO circuits operative at any given time, even 0.1% crosstalk (one part in one thousand) from each of 500 sources aggregates to 50%, a clearly unacceptable level. Crosstalk in a big BGA needs to fall off to much less than one part in one thousand. In this basic setup crosstalk doesn't seem to fall off fast enough.

Basic Six-Ground Crosstalk



Basic Test Configuration

A1	•	×	×	×	×		
	•	×		×	×	×	×
	•				×	×	
		×		×			•
				×		×	
		×					×
							•
			•	×			•



Comparison of Crosstalk

Xilinx® Virtex-4™ FPGA

- **with Sparse Chevron power/ground pattern**
- **100 balls at top left corner**

Altera® Stratix® II FPGA

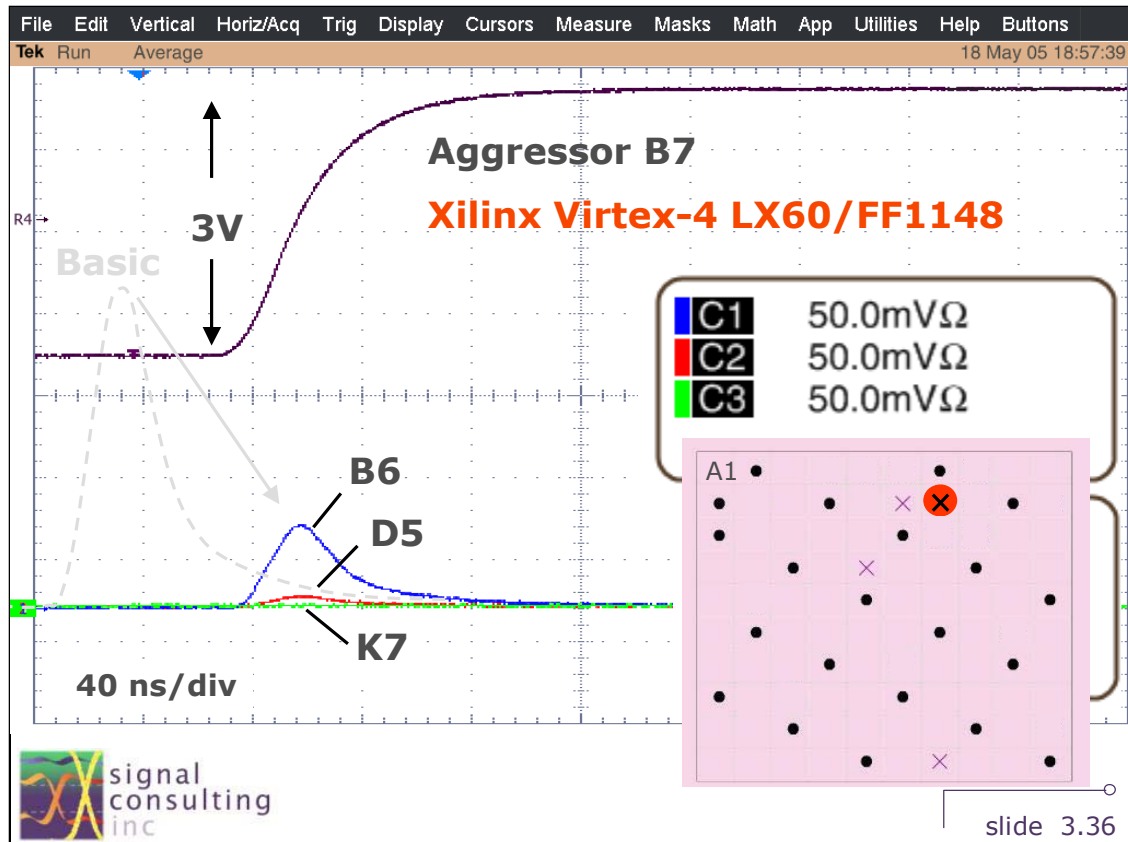
- **100 balls at top left corner**

In this figure, I have added a total of sixteen extra return connections, representing a combination of both power and ground pins.

For comparison, the figure shows crosstalk to B6, taken from the previous view representing the basic 6-ground pattern, as a very light dotted line on the left side of the figure. [NOTE: the dotted line is visible on-screen, but you may have difficulty seeing it on a B&W hardcopy].

As you can see, compared to the basic ground pattern the crosstalk has abated substantially.

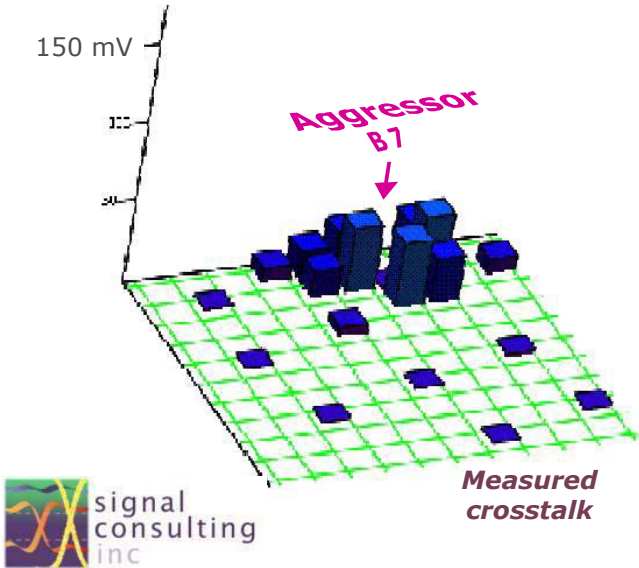
This pattern of returns is a wonderfully regular array, called a sparse chevron.



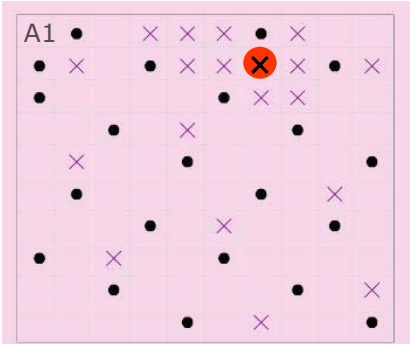
The sparse Chevron is the pattern of return connections in use on a Xilinx Virtex-4 LX60 FPGA in the FF1148 package. As you can see, this pattern keeps crosstalk down to a consistent level everywhere across the chip.

Next let's try a different pattern of return connections. This time I'll set the switches to emulate the pattern of power and ground connections beneath an Altera Stratix-II FPGA.

Xilinx Crosstalk



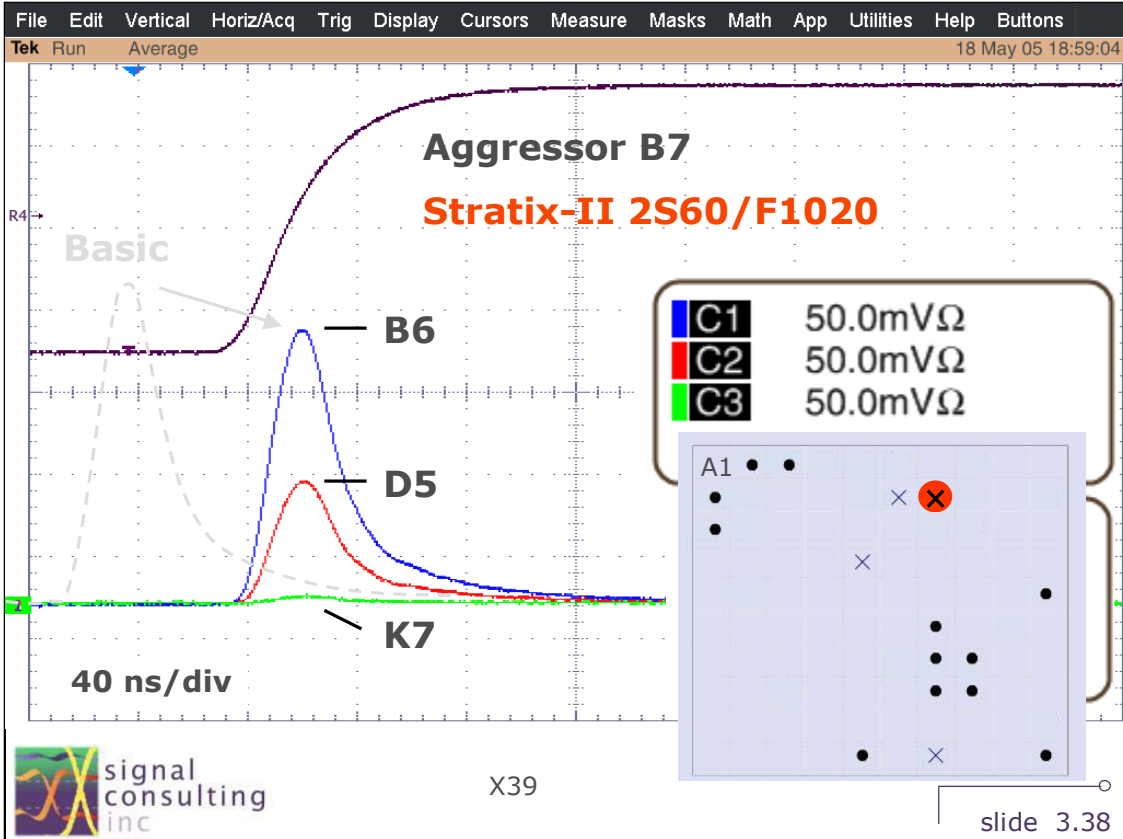
Virtex-4 LX60/FF11480



At this point we have a total of twelve return connections (representing a combination of both power and ground pins) . That's all the Altera part has in this 100-pin region. As you can see the crosstalk on our first three victims is considerably larger than in the previous case.

The figure also shows crosstalk to B6, taken from the previous view representing the basic 6-ground pattern, as a very light dotted line on the left side of the figure.

This part, in the upper right-hand 100 balls, has twice as many returns as the basic pattern. Did you expect that to halve the crosstalk? It didn't. Compared to the basic 6-ground pattern, crosstalk is only slightly improved.



This chart tells the whole story on crosstalk from aggressor location B7.

The pattern of power and ground connections in use now is an irregularly planned array. The returns appear in clumps, which is known to be an ineffective use of return positions.

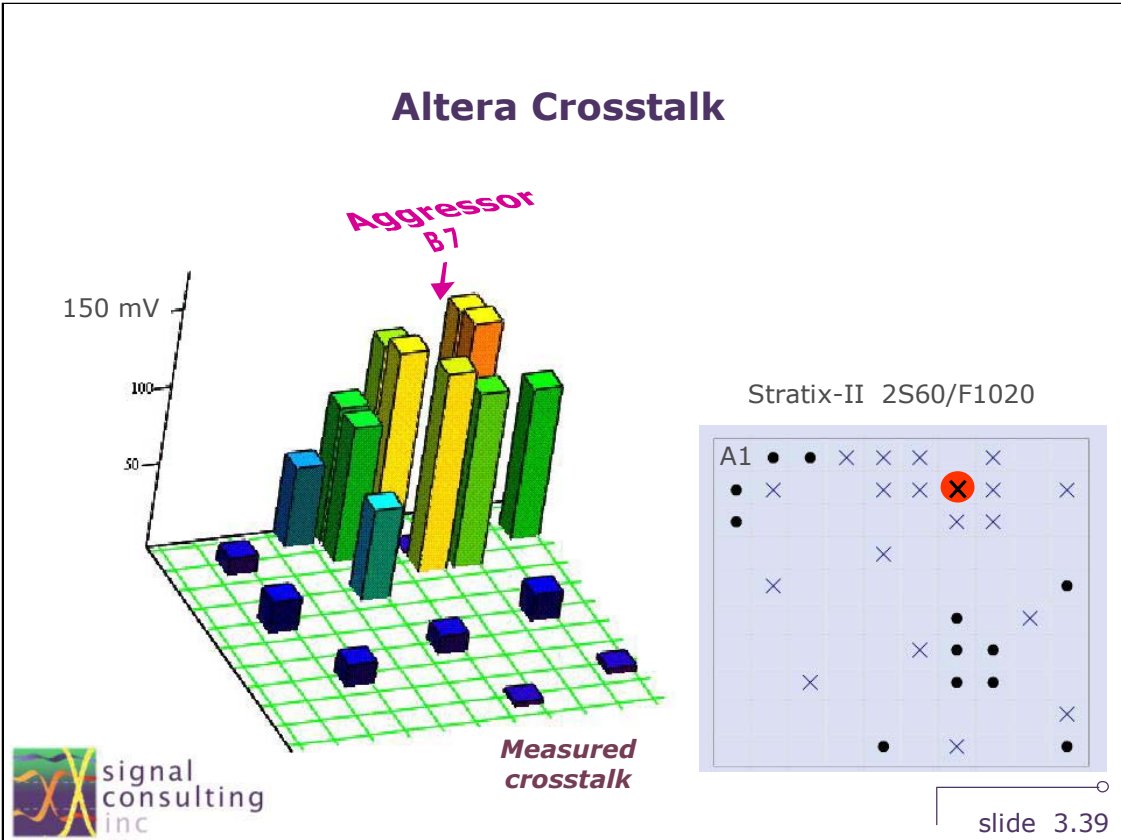
A clumped pattern is ineffective because each return has only a localized, regional effect. The general rule for returns is that adding more returns always diminishes crosstalk. To first order, that's certainly correct, but clumping all your returns in one localized area, while it may press the crosstalk in that local region down to impressively small levels, doesn't help way over on the other side of the chip.

Where this package needs help is not in the bottom right corner. It needs help near the aggressor, near the positions where there is a lot of crosstalk.

To reduce crosstalk near the aggressor, you must disperse the returns over to that part of the package.

Interspersing returns among your signal pathways reduces crosstalk.

We used to be at a level of speed where this stuff didn't matter. Those days are long gone.




Have you ever tried grounding unused I/O pins as a way to mitigate crosstalk? The theory behind this is that an I/O pin, which switched low, acts somewhat like an ordinary ground ball.

Soft Grounds

Also known as "Virtual Grounds"

The difference is that a real ground ball connects straight to the ground plane inside the BGA package. A grounded I/O pin, on the other hand, leads to an I/O cell, which may have a substantial output resistance. The output impedance of the I/O cell acts in series with your ground connection. The question is whether the grounded I/O can possibly serve as a good ground connection when burdened by this additional series impedance.

Soft Ground Concept




Rather than dwell on the theory, let's just try it out. I have equipped each ground position with two switches. The first connects a ball straight to the ground plane. The second switch connects the same ball in question to ground through a series resistance of twenty ohms (representing the output resistance of a 2.5-volt 8-ma driver).

Starting with the stock Altera ground pattern, let's add in some soft grounds. I will switch in a pattern of 14 additional soft grounds. These additional soft grounds are positioned in the same locations as the Xilinx Sparse Chevron pattern, as if you were trying to duplicate the crosstalk performance of the Virtex-4 Sparse Chevron power-and-ground array using soft grounds.

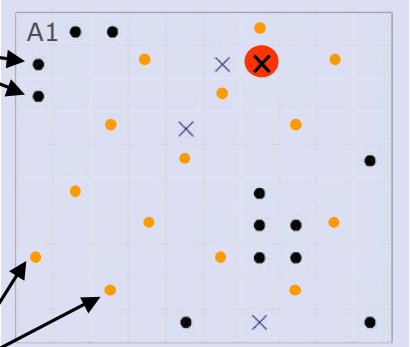
This test will indicate whether 20-ohm soft grounds in those locations can provide the same crosstalk protection as a real Sparse Chevron matrix of hard grounds.

Testing Soft Grounds

Soft ground switches (20-Ω)



Hard grounds



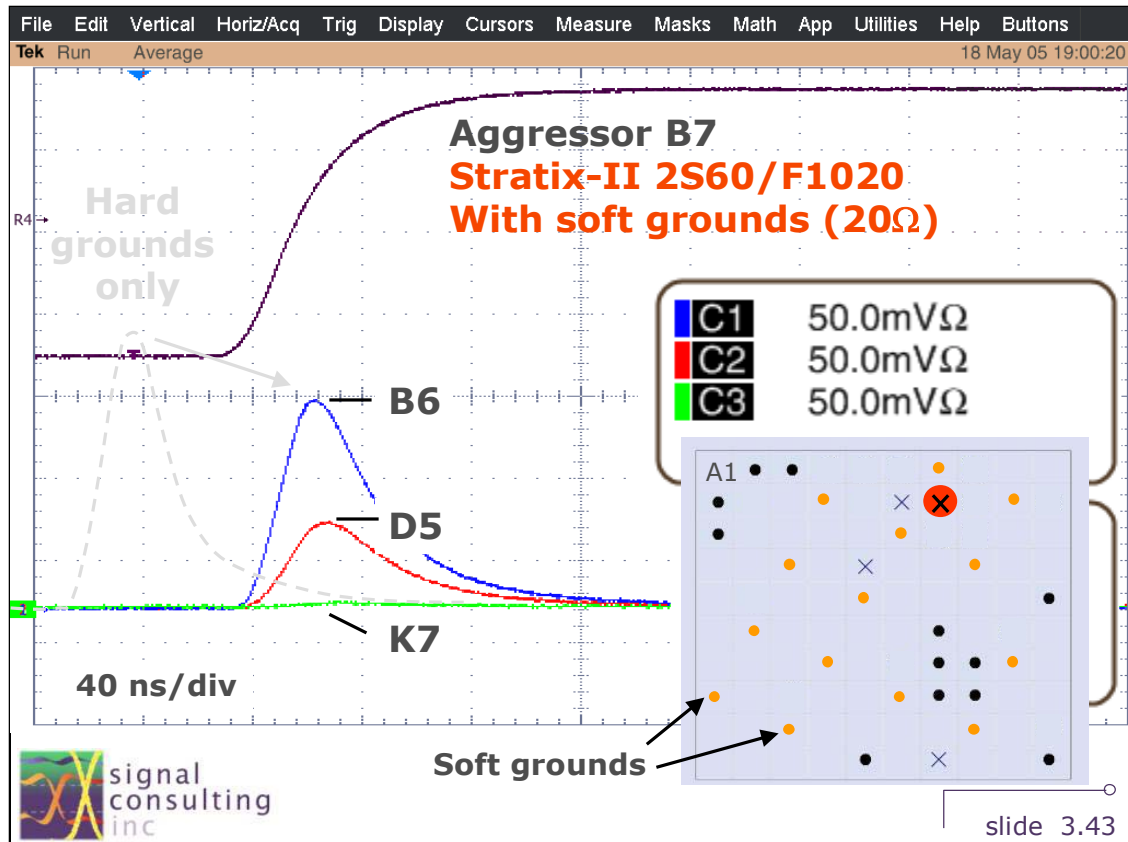
Soft ground connections (20-Ω)

Stratix-II 2S60/F1020

On the left (dotted line) I show crosstalk using the stock Altera hard-ground pattern (from our previous measurement). The solid lines depict the crosstalk measured when the 14 additional soft grounds are added.


Apparently, in this geometry, it is slightly helpful to have 20-ohm soft grounds, but, if you remember the level of crosstalk in the Xilinx package, these soft grounds are not nearly as good as a real, directly-connected ground balls.

The next slide shows the complete record of comparison.




The complete record of measured results illustrates the advantages of the hard grounds provided by the Xilinx sparse Chevron power and ground matrix.


Comparative Crosstalk



Altera



**Altera with
20-Ω Grounds**



Xilinx



*Measured
crosstalk*

Grounding unused I/O pins used to be an effective way to mitigate crosstalk. An older package, for example, might have a per-pin ground inductance on the order of 5 nH. The impedance of a 5-nH ground connection, at 1 GHz works out to 30 ohms. In comparison with the 30-ohms of reactive impedance you already have, 20 more ohms of real impedance hardly matters. Thus, in older packages a grounded IO might be nearly as good as a real, directly-connected ground ball. In today's world, however, the Xilinx Virtex-4 FPGA package provides solid power and ground planes, good internal bypassing, and a ground ball adjacent to every signal pin. The effective inductance per ball associated with that geometry is very low, perhaps less than 1/2 nH, or 3 ohms at 1 GHz. How are you going to create an impedance comparable to that using a grounded IO structure that, even for the most powerful types of IO cells, has at least 6-10 ohms resistance? It isn't going to happen.

Lets use our model to see how low an impedance you need in a grounded-IO configuration to accomplish the same reduction in crosstalk as one direct, hard-grounded ball.

Why Soft Grounds Didn't Work

Old Package	Modern Package
5 nH $\times 2\pi (1 \text{ GHz})$	$1/2 \text{ nH}$ $\times 2\pi (1 \text{ GHz})$
<hr/>	<hr/>
30 ohms	3 ohms

**Required soft-ground resistance
(closest ball)**



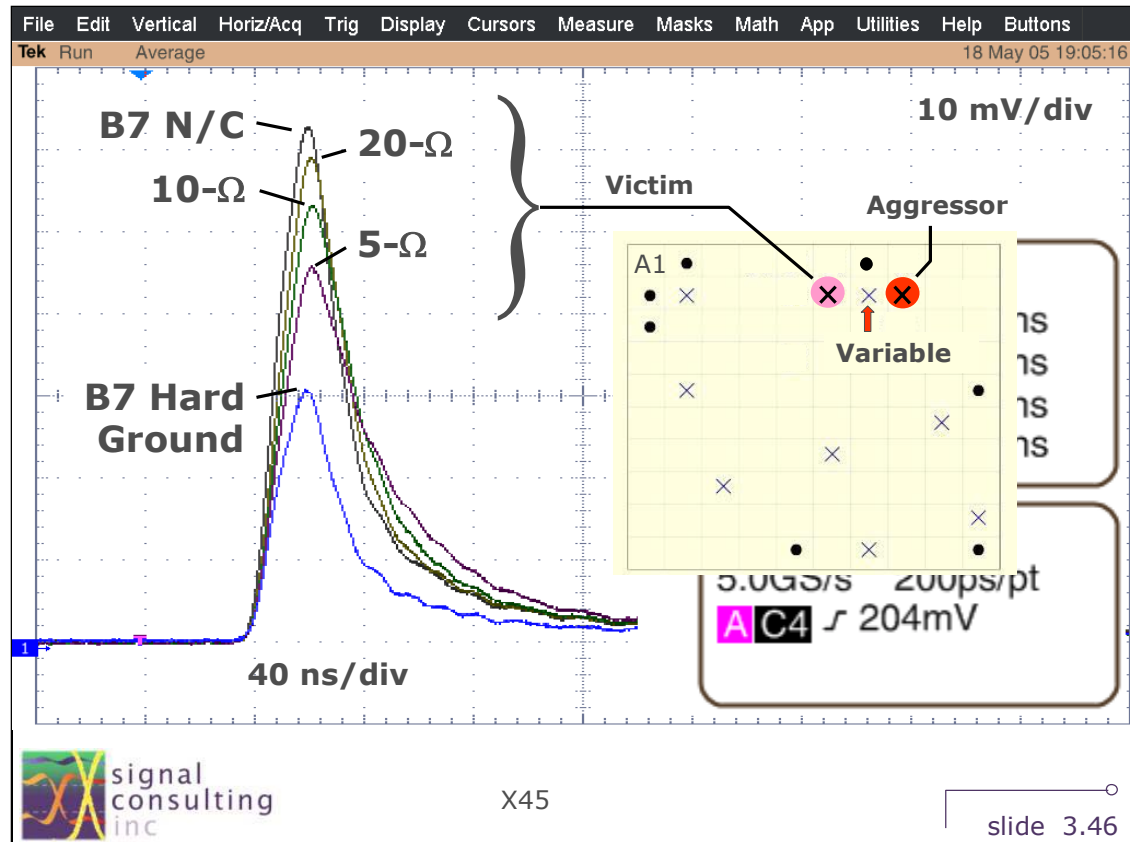
To perform this test I will use three positions on the model. I'm sending a signal into B8 (aggressor), and receiving on B6 (victim). I have the basic pattern of six fixed grounds engaged around the periphery of the model, plus location A7, right above the aggressor/victim pair, is grounded, as it would be in a Sparse-Chevron pattern.

None of the other ground switches are turned on. Right now the crosstalk reads 97 mV.

Now I'm going to short out ball B7, right between aggressor and victim, as if it were a grounded-IO structure. If the soft ground is any good, you should expect a 2:1 reduction in crosstalk, because you would then have two good grounds (A7 and B7) in parallel.

To vary the impedance at B7, I have an SMA connector with a 20-ohm resistor soldered inside. We don't expect this to do much, and as you can see, compared to a no-connect (N/C) at B7, it doesn't. OK, how about a 10-ohm resistor? Not much better. How about 5 ohms? At five ohms we are starting to get a significant effect (crosstalk is reduced to 69 mV), but let's compare that to a real, direct ground. When I plug in an SMA shorting block, the result is 40 mV, substantially better than anything else.

Conclusion: In our model, even a 5-ohm impedance in series with a ground ball cripples its performance.



I don't want to be misleading, because soft grounds can provide benefits, but only when at least one of several conditions is met:

- 1) The IO output resistance is strikingly low, lower than commonly available in FPGA implementations,
- 2) The BGA power/ground patterns are particularly poor, even worse than either of the alternatives in my study, in which case anything would be an improvement,
- 3) The vias are particularly long. This exaggerates the inductance of the naturally grounded balls, making the resistance of the IO drivers appear less significant.

If you have implemented grounded IO pins in an advanced package as a means of crosstalk mitigation, I suggest you try your system without them (or program them off) to see if they matter. In many cases, they will make little or no difference.

The insertion of a large number of soft grounds does have one clear benefit. It separates the various signal pins to such a degree that crosstalk is naturally mitigated by the separation. This benefit accrues whether the soft grounds are truly switched to zero or simply left in the tri-state condition.

Grounded IOs help when you have...

- Strikingly low IO resistance
- Very poor pwr/gnd ball pattern
- Particularly long vias



I have modeled some fairly long vias in this demonstration, partly because I wanted to remind you that crosstalk occurs not only within the BGA package, and between the BGA balls, but also among the vias buried within your printed circuit board. The vias in my model are 3 inches long; at 57:1 scale that works out to 53 mils, almost all the way through a standard pcb, or about halfway through a thick one.

Using simulation, I can show the expected results with other via lengths. Shorter vias produce generally less crosstalk, but the same sort of soft-ground ineffectiveness.

In the next slide I will change to a via length of 6 mils, representing the distance between the surface of a typical board and the first ground (or power) plane. I'll suppose you have routed your traces only on the first couple of layers, so the current-carrying region through which your signals pass is only the height of the balls (20 mils) plus the 6-mil via.

I will try various patterns of ground connections to see how the "soft-ground" idea works at this reduced via height.

Extend Results with Simulation

Prediction

- Soft grounds work poorly at all via lengths

Simulator

- Quasi-static 2-D field solver (constant B-field in z-axis)
- 1mm ball pitch + first plane 6 mils below surface
- Very short vias to minimize crosstalk

Try

- Altera pattern of hard grounds
- *PLUS* Xilinx pattern, implemented as soft grounds
- All possible signal locations



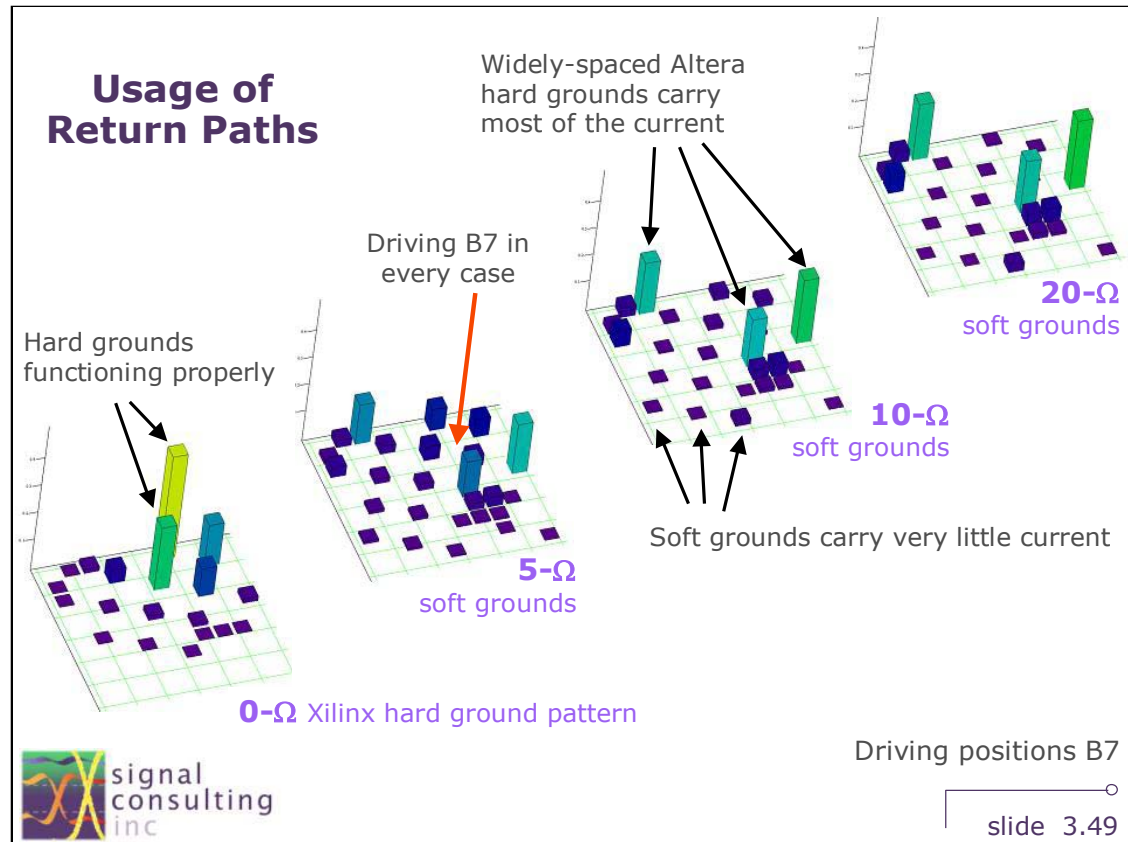
This set of simulation results shows how much current flows in each return connection, as a function of the soft-ground I/O resistance. The aggressor in each case is B7.

Each simulation incorporates the stock Altera ground pattern. In addition to that pattern, I have switched in in a pattern of 14 additional soft grounds. These additional soft grounds are positioned in the same locations as the Xilinx Sparse Chevron pattern.

I tried soft-ground IO impedances of 0, 5, 10 and 20 ohms. I don't think you will ever see as little as 5 ohms in a soft-ground connection, but I ran the simulation anyway. The zero-ohm pattern represents a hardware implementation, like Xilinx, that uses actual hard grounds everywhere. The simulator drives position B7 and calculates the distribution of return currents at all the ground positions. The assumed rise/fall time is 1 ns.

The bottom left diagram shows how the Xilinx hard ground pattern naturally concentrates return current in the three return positions nearest the aggressor, B7. This configuration keeps signal and return current close together, limiting crosstalk.

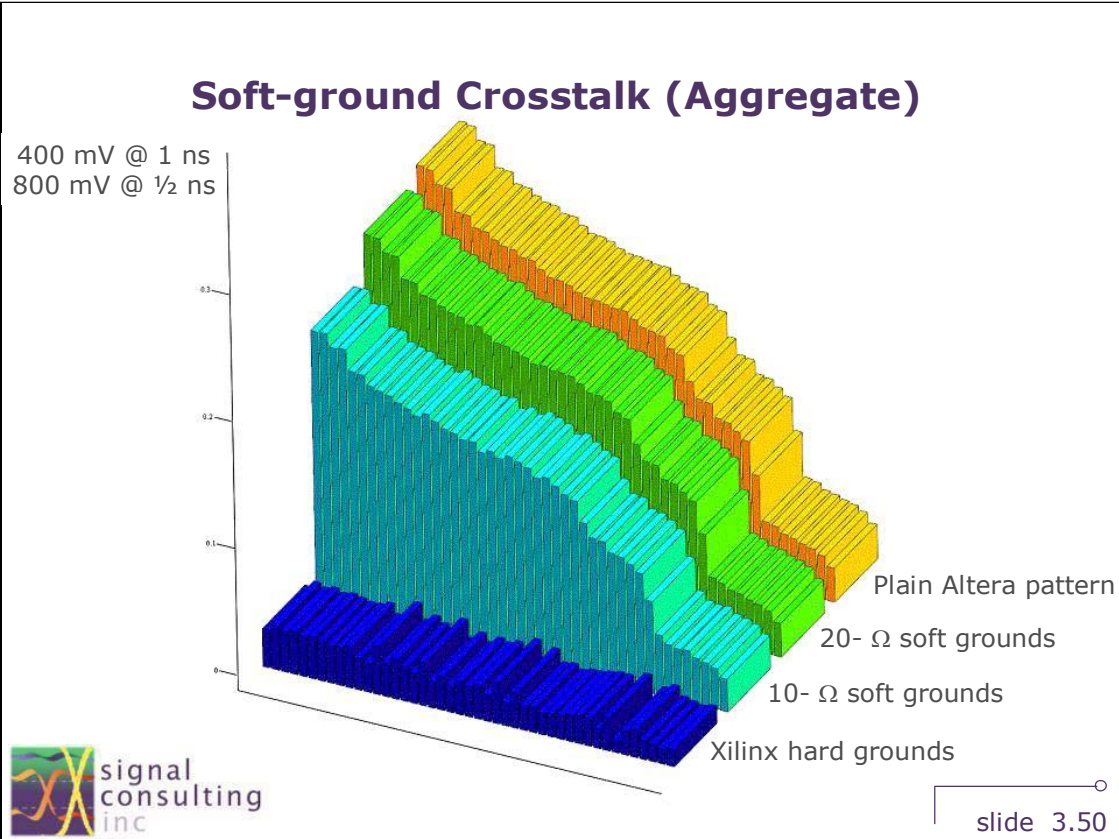
The other three diagrams illustrate the return patterns for various soft-ground implementations. Never is the impedance of the soft-ground pins low enough to attract much significant return current. Most of the return current at 350 MHz flows on the hard grounds. As a result, the soft grounds have limited influence on crosstalk.



This figure wraps up our talk with a picture showing worst-case aggregate crosstalk for our 100-ball portion of a BGA package. It assumes you drive each and every ball with a 2-volt signal swing at 1-ns rise/fall time (or 1/2-ns rise/fall time). The crosstalk is sorted in each row from biggest to smallest. These are peak crosstalk magnitude numbers, not peak-to-peak.

This figure takes into account only the crosstalk generated between the BGA substrate and the first ground plane, assumed to be 6 mils below the surface of your board. It does not take into account any additional via crosstalk, which occurs in proportion to the overall length of your vias.

From this picture I conclude that Altera power and ground pin pattern, when supplemented with soft grounds in the Sparse Chevron pattern, does not achieve the same level of crosstalk as an actual hard-ground Sparse Chevron pattern.



That's all the time we have today. I would like to thank Mark Alexander for his help with this project, and thanks also to Tektronix for providing the equipment that made these measurements possible.

If our program has stimulated your interest in further research, the Xilinx signal integrity site www.xilinx.com/signalintegrity holds a number of resources useful for high-speed designers, including information about my new SI tutorial for RocketIO™ serial transceivers, now available on DVD at www.xilinx.com/store/dvd.

At my web site www.sigcon.com you will find a treasure trove of additional publications (282 at last count), plus a full schedule of my High-Speed Digital Design seminars, complete course outlines, other films, newsletters, an article index, and much more.

Thanks also to all of you for your constant stream of fascinating letters and emails, and don't forget tell a friend what you learned.

For Further Study

Xilinx

- www.xilinx.com/signalintegrity Resources useful for high-speed designers
- www.xilinx.com/store/dvd My new SI tutorial for Xilinx RocketIO™ serial transceivers—now available on DVD
- <http://www.xilinx.com/products/virtex4/advantage/sipi.htm>

Website: www.sigcon.com

- Full schedule of seminars,
- Seminar course outlines,
- SiLab films,
- Newsletters,
- Article archives, and much more.



At this time I would be pleased to consider any questions or comments you may have.

Dr. Howard Johnson

BGA Crosstalk in Depth Ground Ball Placement

I'll be doing courses
open to the public at

Oxford University

June 20-24, 2005

www.sigcon.com

Questions?

