

BGA Crosstalk

Xilinx[®] Virtex[™]-4 FPGA and
Altera[®] Stratix[®] II FPGA

by Dr. Howard Johnson

Details, measured lab results, and theory of crosstalk involving hundreds of outputs switching simultaneously in a high-speed Virtex-4 FPGA package, as delivered via TechOnLine webcast March 1, 2005.¹

Measurement Setup

Thank you for inviting me to speak at this tech online forum.

This talk is about measurements, so I should begin by showing you how I made them. First, I obtained a test board comprising one Virtex-4 LX60 FPGA in the FF11480 package, plus an Altera Stratix II 2S60 FPGA in their F1120 package. On both these parts I arranged to bang a lot of outputs up and down, making a good deal of noise—hundreds of outputs switching at once.

While making all that noise I measured the crosstalk on one poor little guy that tried to remain “stuck at zero” or “stuck at one” during the test (Figure 1). This technique represents a realistic worst-case appraisal of the crosstalk emanating from any I/O driven out of the package during the noisy events.

It also indicates the crosstalk received by I/O's directed inwards towards the package. To see how that works we need to delve for a moment into the details of inductive crosstalk, because that is what this configuration delivers: inductive crosstalk.

Figure 2 illustrates a pair of PCB power and ground planes, between which I have connected a BGA package. Of course, the BGA package does not reside physically between the planes, but in a schematic view that is a convenient way to represent the circuit. The BGA balls and BGA routing appear in the diagram. The package holds two totem-pole drivers.

Suppose the load at **F** is initially charged HIGH. At time zero, switch **C** drives LOW, creating a huge I/O current transient. As this current flows through the finite inductance, L_{GND} , representing the matrix of ground balls underneath your BGA package, it creates a voltage disturbance on the chip substrate.

¹ © 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.

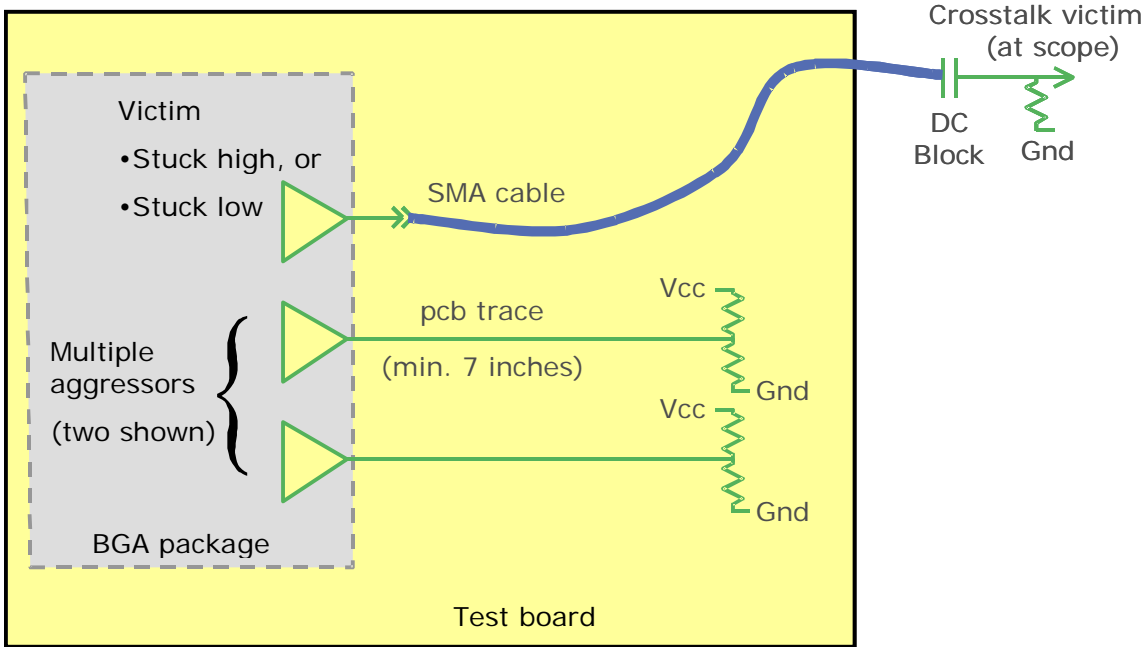


Figure 1—Each FPGA on the test board transmits multiple aggressive signals while I observe noise emanating from a few static outputs.




Figure 2—This simplistic view of ground bounce ascribes inductance only to the ground connection.

The victim **D**, which remains stuck low throughout the experiment, picks up this voltage glitch and transmits it straight out of the package to the scope, where you can see it.

This simplistic view of the problem is qualitatively good at explained *why* switching events on one pin cause noise on another, but quantitatively poor at predicting noise amplitudes.

To explain the inadequacy of that diagram, I will have to ask you to abandon the concept of inductance as a property of an individual wire or conductive pathway. You must instead perceive inductance as it truly exists: a property of the space between conductors.

In any high-frequency inductive problem, the relevant magnetic field resides in the space between conductors, not in the conductors themselves. This field, not the conductors, causes all inductive effects.

In analog circuits, we are used to thinking of an inductor as a tightly wound coil. The coil concentrates its magnetic field in the space within the body of that inductive component.

The inductances we deal with are different. They occur as parasitic inductances, bound to magnetic fields that exist in the spaces between our signal and return pathways.

Figure 3 shows a better way to think about mutual inductance. This diagram shows the chip package on top, above the balls and vias. The three signal vias **D**, **E**, and **F** each penetrate the ground plane.

Switch **C**, when it closes, initiates the change in current drawn in red.

This action fills the space between vias with an intense magnetic field. I have partitioned that field into three sections labeled L_1 , L_2 , and L_3 . Faraday's law says the crosstalk measured at **D** varies in proportion to the total magnetic flux L_1 lying between via **D** and the nearest return-current position. Crosstalk at **E** varies with both fields L_1+L_2 . Since L_1+L_2 exceeds L_1 alone, the crosstalk at **E** must exceed the crosstalk at **D**. That is, crosstalk varies strongly with the position of the victim via.




Figure 3—Inductance exists not as a property of any individual wire or conductive pathway, but as a property of the space between conductors.

The variation in crosstalk with signal-pin position is not explained by figure 2. By assigning all the package inductance specifically to the ground pathway, that figure ignores the impact of signal positioning on observed crosstalk. In a real problem, ground bounce always varies with your proximity to a good return connection. The signals closest to a good return suffer the least "ground bounce" or "power bounce".

Since in a complex multi-ground-pin package we can no longer think of "ground bounce" using the simplistic model in the previous drawing, I suggest that you simply think of it as a "crosstalk effect". Do not try to partition the effect into discrete portions due to noise on the ground substrate, noise on the chip's internal Vcc rail, or noise accumulated in the ball-and-via field underneath the package.

This type of magnetic coupling varies with the separation between signal pathways and the *length*, or in this case, *height = via + ball + package thickness*.

It also scales with the rise or fall time of the aggressive signal (specifically, the di/dt).

In case you are wondering about capacitive crosstalk, it contributes very little to this setup. You will see that in the measured results in a few minutes. If there were significant capacitive crosstalk, then when pin F switches LOW, you would see pin E also glitch LOW, but it doesn't. It glitches in the HIGH direction, an action only made possible by the inductive coupling explained in this diagram.

If you understand how transformers work, Figure 4 makes it clear that inductive crosstalk affects both drivers and receivers alike. Inductive crosstalk acts like a

low-impedance transformer hooked in series with your circuit. The primary of the transformer is the I/O current path. The various multiple secondary windings correspond to each of the victim vias. The inductive crosstalk voltage induced across each secondary is determined by the proximity of that secondary loop to the primary.




Figure 4—Inductive crosstalk affects both drivers and receivers alike.

If a low-impedance connection (in the top case, a stuck-at-low driver) appears to the left of the transformer then the noise voltages develop to the right. If the low-impedance connection (in the bottom case, an incoming line) appears to the right of the transformer then the noise voltages develop to the left, affecting the receiver. Either way, inductive crosstalk gets you.

Measuring the crosstalk *output* from a stuck-at-zero (or stuck-at-one) I/O cell is a great way to determine the general level of BGA crosstalk *input* to receivers on that same package.

A picture of the test board appears in Figure 5. Mark Alexander did a fine job building this setup. It was designed as much as possible for a direct apples-to-apples comparison of the two parts. On the left you can see the Altera FPGA with 1,020 balls, on the right is the Xilinx FPGA with 1,148 balls. The board has 24 layers, and is 110 mils thick. Three I/O voltage regions are included on the board (1.5, 2.5 and 3.3-volt) with each power plane sandwiched between grounds. To eliminate any concern about differences in power architecture, Mark came up with one common power arrangement that more than met both manufacturer's guidelines, and used precisely the same setup on both parts.

You can see the SMA fittings for viewing particular ports on each chip. Not every I/O is instrumented.




Figure 5—On the left of the test board you can see the Altera Stratix II with 1020 balls, on the right is the Virtex-4 part with 1148 balls.

The Altera and Xilinx parts fit on opposite sides of the layout. All planes are completely isolated side-to-side, with separate power supplies. Not even the grounds touch, except through the instrumentation connections.

We ran all the tests shown here with only one side of the board active at a time, although we did test for interference between sides and saw no observable changes.

From the prior discussion, you should expect crosstalk to vary strongly with the proximity of the victim to a return pin. Building on that point let's look at the pattern of ground and power balls in the Virtex-4 LX60 FF1148 package.

Figure 6 depicts the array of return pins used in the FF1148 package. The dark dots represent ground balls; the lighter dots power balls.

Assuming perfect decoupling within the BGA package, the power and ground balls are equally effective as conduits for returning signal current. For the rest of this paper I will not differentiate between power and ground balls, but simply call them "return balls". The other marks (X) represent high-speed signals. Blank

spots indicate low-speed programming signals, PLL control signals, and other non-high-speed functions.



Figure 6—The LX60 package includes a well-distributed array of power and ground balls to control crosstalk.

Some of the blank spaces on the left and right sides of the array are not populated in the Virtex-4 LX60 FF1148 package. We picked this configuration because it has about the same total number of active high-speed signal pins (500) as the Altera Stratix II 2S60.

The existence of all those return balls helps tremendously to reduce crosstalk from distant aggressors. If you were to do a “walking ones” test, moving the aggressor further and further away from the victim each time, and observing crosstalk on the victim during the whole test, you would see crosstalk fall off geometrically each time you pass a return ball position.

To illustrate that point, let’s look in detail at the distribution of crosstalk associated with one specific I/O pin. I will pick signal A10, which is up on the top row, about a third of the way in.

By the way, I picked this pin for today’s studies because our calculations predict pins located near the edge of the pattern should be a worst-case location for crosstalk performance.

Because there are so many power and ground balls in this package, and so evenly distributed, the crosstalk contributions around the particular signal A10 lies tightly grouped around that particular location (Figure 7).




Figure 7—Only the nearest aggressors contribute significantly to aggregate crosstalk at position A10.

Figure 8 examines a number of other locations. The story is the same in each case: crosstalk falls off geometrically as you pass beyond each successive ring of power or ground pins.

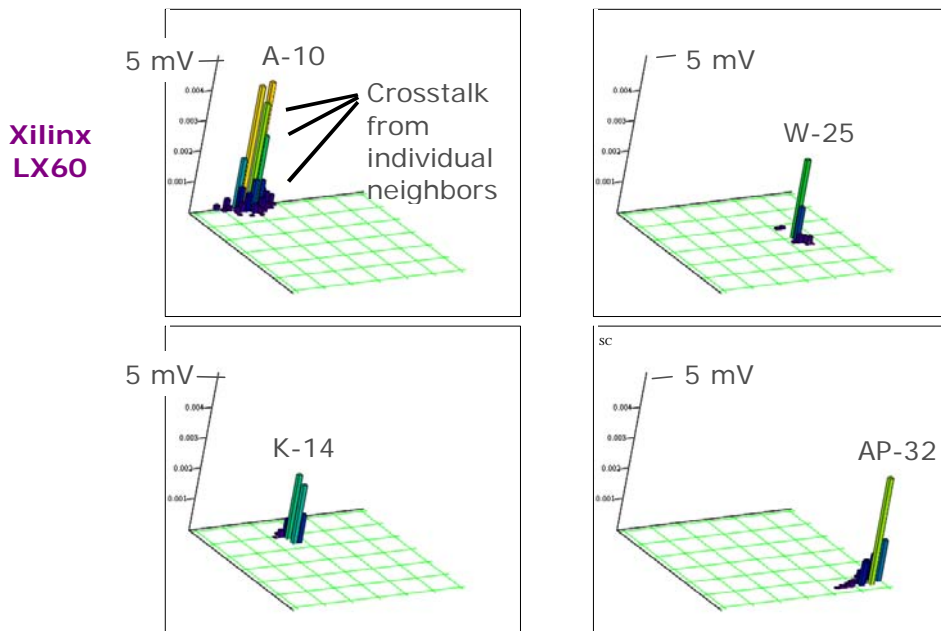


Figure 8—No matter where you go, the main contributors to crosstalk always lie tightly grouped around the victim.

Some details behind the simulation used to generate these pictures are that the drivers are set to a di/dt of $9.05E+06$ A/s, corresponding to a Xilinx 1.5V LVCMOS 4 ma *fast* driver, and all vias are set to an average depth of 0.035 in. below the surface of the PCB.

What contains the distribution of returning signal current, and thus crosstalk, is the existence of lots of densely packed ground and power pins. These power/ground pins are tessellated in a regular array of ten elements called a “sparse chevron” (Figure 9). The overall ratio of signals to grounds to powers in this package is 8:1:1.

Since both power and ground pins act equally as conduits for returning signal current in this package, what we really have here is a well-distributed 4:1 ratio of signals to returns.





Figure 9—The LX60 FF1148 package is tessellated with a regular array of power and ground pins, called a “sparse chevron” pattern.

Figure 10 plots the first of our measured results. This data was acquired using a Tektronix TDS6804B Digital Storage Oscilloscope, 8 GHz bandwidth, 20 Gs/s, using direct inputs with 40-inch low-loss SMA cables. The bottom portion of the figure depicts a nearby aggressor, going first high and then low. The resulting crosstalk waveforms just above represent the crosstalk picked up with victim A10 stuck high (red) or stuck low (blue).



Tek TDS6804B Digital Storage Oscilloscope.
8 GHz bandwidth, 20 Gs/s
Direct inputs with 40 inch semi hard-line SMA cables

Figure 10—The crosstalk waveforms move in a direction opposite the motion of the aggressive signal—proof that crosstalk in this experiment results from inductance, not capacitance.

Note that the crosstalk waveforms move in a direction opposite the motion of the aggressive signal. This movement proves conclusively that the crosstalk results from inductance, not capacitance. Capacitive crosstalk would have made the victim move the same direction as the aggressor, which is not happening here. Certain tricks that work well to reduce capacitive crosstalk, like driving an input ball using a very low-impedance source, do nothing to improve the inductive crosstalk generated within the BGA ball field and vias underlying this FPGA package.

Crosstalk from this aggressor peaks at only six millivolts (about one percent). I made heavy use of averaging in the TDS5804 to improve the noise floor for these measurements, making possible the clarity of view represented in the diagrams.

The fact that the stuck-high and stuck-low signals present nearly the same crosstalk indicates near-perfect performance of the internal bypassing within the FF1148 package.

Figure 11 shows a similar view, but with more aggressors. The aggressors turn on and off in sequential fashion, first one, and then the next, so you can see

clearly the crosstalk resulting from each. Crosstalk falls off quickly as you make your way further and further from the victim.




Figure 11—Crosstalk falls off quickly as you move the aggressor further and further from the victim.

Three Measures of Crosstalk

Now that you are comfortable with the format of my measurements, let me present the three comparative tests used to evaluate the Xilinx and Altera packages. The purpose of this work is to quantify the practical improvement in crosstalk made possible by the "sparse chevron" ground pin tessellation. These tests are called:

- Spiral Test
100 nearest outputs, exercised individually
- Accumulating Spiral Test
100 nearest outputs, aggregating in larger and larger groups
- Hammer Test
500 outputs all together

I ran all three tests run on both packages. These packages each have solid internal power and ground planes. In that sense, they are very similar. They differ in their BGA pin assignments (Figure 12).

The two packages have nearly the same number of high-speed signal balls: 504 for Xilinx and 505 for Altera. The Xilinx package, being slightly larger, accommodates more power and ground balls (185 ground and 144 power for

Xilinx, versus 114 and 103, respectively, for Altera). The biggest difference, though, lies not in the sheer number of ground and power balls but in their distribution. The Altera package has large areas devoid of returns. The Xilinx pattern is more evenly peppered with returns.

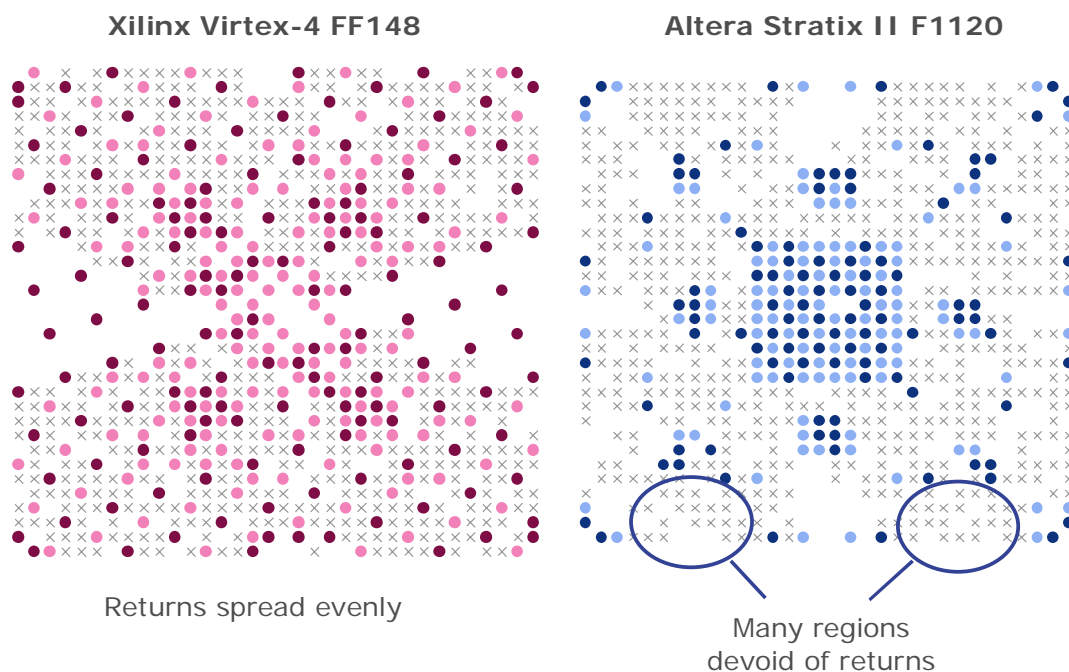


Figure 12—Both packages have solid internal power and ground planes, but very different power/ground pin assignments.

The Spiral Test (Figure 13) exercises the 100 nearest outputs, one at a time, each going up and then down just like in Figure 11. This figure shows the region of the Xilinx Virtex-4 LX60 FF1148 package used in this test. The victim location, A10, appears on the top row, in the middle of the aggressor region. Mark's spiral exercises the nearby aggressors in succession, working its way around and around location A10 to increasingly distant aggressors. During this test, all outputs are equipped with 1.5-volt LVCMOS 4ma *fast* drivers.

The Altera Spiral Test (Figure 14) exercises slightly different pins due to differences in the pin assignments, but encompasses the same number of I/Os at similar average distances. During this test, all outputs are equipped with 1.5-volt LVCMOS 4ma drivers, but there is no speed selection capability in the Stratix II.

Spiral Test Pattern: Xilinx




Figure 13—The Xilinx “spiral test” exercises 100 nearest-neighbor aggressors in succession, working around and around location A10 to increasingly distant locations.

Spiral Test Pattern: Altera

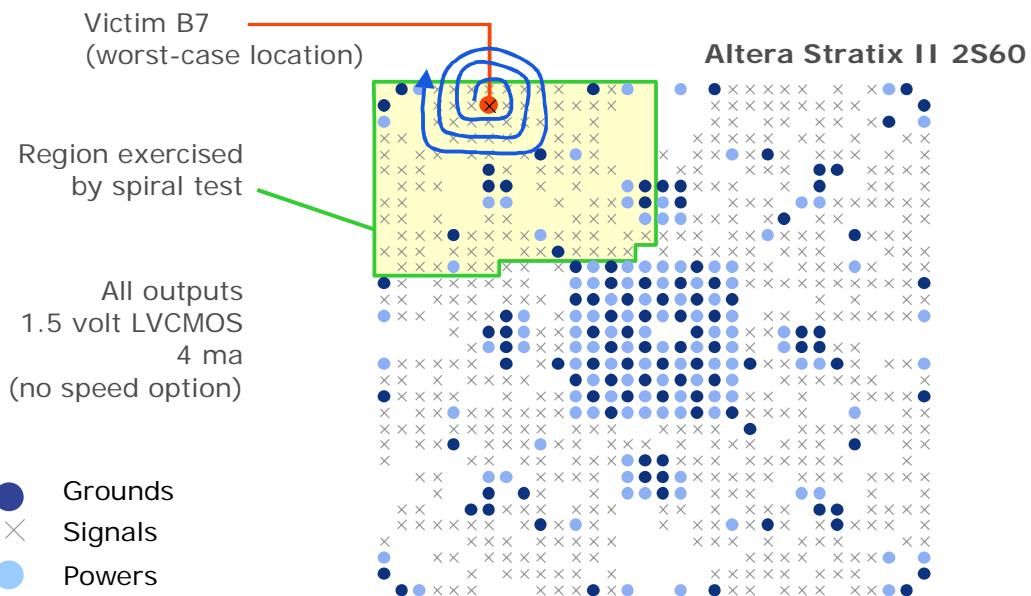
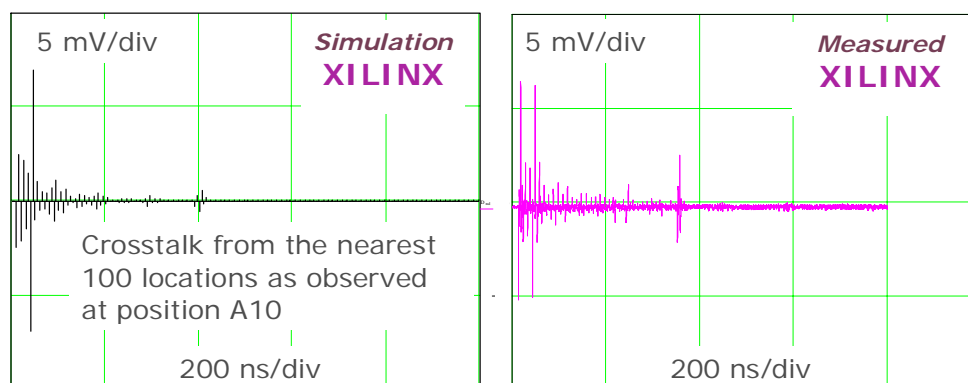


Figure 14—The Altera spiral test exercises slightly different pins due to differences in the pin assignments, but encompasses the same number of I/Os at similar average distances.

Before I compare the packages to each other, let's compare theory with measurement (Figure 15). The combination of theory, simulation, and measurement is what Terry Morris at HP calls his "triumvirate of understanding". Terry points out that if you have a theory about how something works, and can make a simulator to predict what *should* happen, and the simulation matches your measurements, then you probably have a good understanding of what is going on.

Spiral Test Results: Simulation vs. Measurement



MathCad

- Signal positions
- Power/ground positions
- Signal amplitude and rise/fall time (di/dt)
- Trace depths

Tek TDS6804B Digital Storage Oscilloscope.

8 GHz bandwidth, 20 Gs/s
Direct inputs with 40 inch semi hard-line SMA cables

Figure 15—If you have a theory about how something works and can make a simulator to predict what should happen, and the simulation matches measurement, then you probably have a good understanding of what is going on.

In this case, I would say the simulation and measurement match fairly well. Both are telling us that the crosstalk pulses fall off rapidly as a function of distance from the victim A10, which remains stuck at zero during this measurement.

Even though there is not an exact match in the waveform pin-by-pin, the aggregate trend shines through. I conclude that this simulation technique realistically captures the crosstalk effect in these packages. Let's put the simulator away for now, but bring it back later to predict aggregate crosstalk at other locations around the package.

In the first comparative test result (Figure 16), I had to turn down the scale to fit both waveforms on the screen. The top trace shows the Xilinx crosstalk at A10, the bottom shows Altera crosstalk at position B7. The Altera part displays two artifacts: more crosstalk on a pin-by-pin basis, and a pattern of crosstalk that does not fall off as quickly as crosstalk in the Xilinx package. If you add up all the individual crosstalk pulses in each waveform, you can see that the Altera crosstalk aggregates to a much higher level.

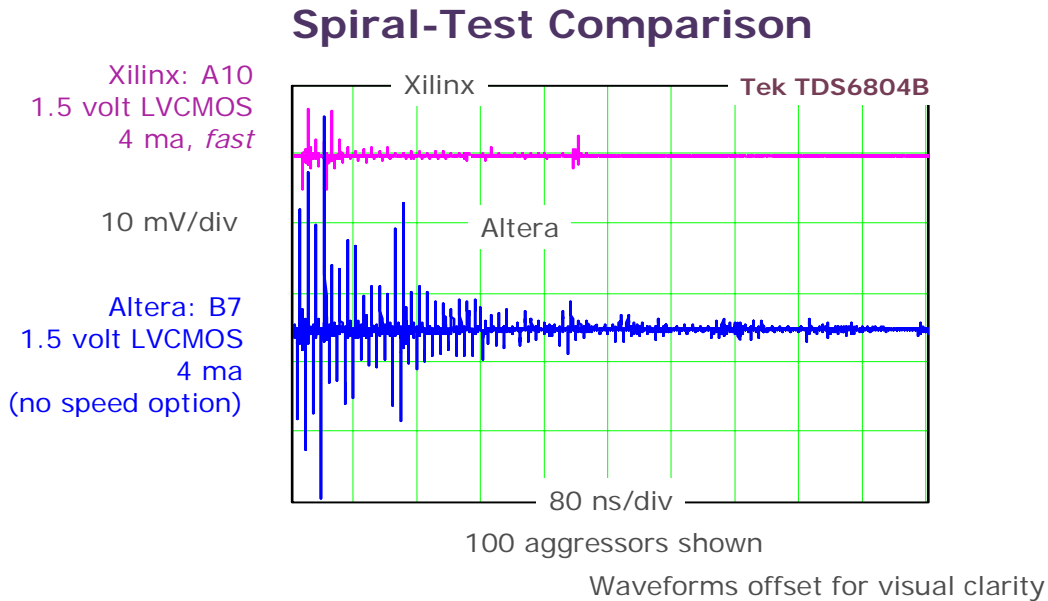


Figure 16—In this spiral test you can clearly see the rapid fall-off of crosstalk on the Xilinx part.

The Accumulating Crosstalk Test (Figure 17) exercises the same balls used in the one-at-a-time spiral test, but with different patterns.


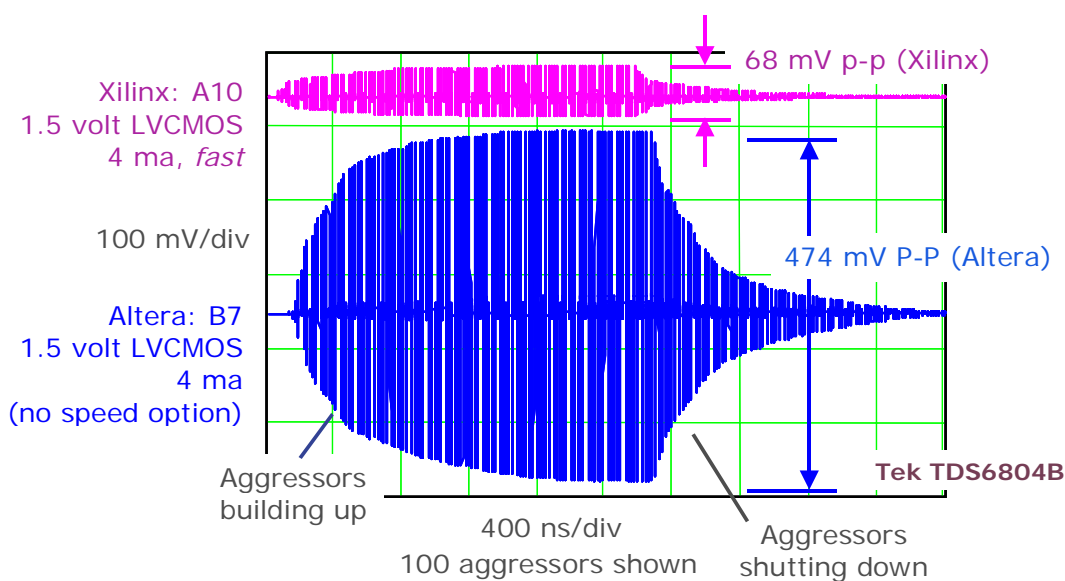


Figure 17—The accumulating crosstalk test exercises the same balls used in the one-at-a-time spiral test, but with different patterns.

First, only the nearest ball fires off, going up and then down. Then it fires off a second time in conjunction with the next nearest ball. Then it does three at a time, then four, and so forth, until a large number of balls (100) are blasting up and down together.

The resulting crosstalk waveforms (Figure 18) clearly show the relative importance of remote balls in the overall aggregate crosstalk waveform. All aggressors in this shot are set to a 1.5-volt LVCMOS 4 ma driver (*fast*, for Xilinx, no speed option for Altera). As before, the Xilinx victim A10 remains stuck low. Similarly, the Altera victim B7 remains stuck low. The Xilinx component tops out at 68 mV p-p of crosstalk, the Altera component generates 474 mV p-p.

Accumulating Test Comparison



Waveforms offset for visual clarity

Figure 18—This accumulating measurement shows aggregate crosstalk building up and then ramping down as you exercise different numbers of pins surrounding the victim.

In the final test, Mark and I went all out to see how much crosstalk we could make. The Hammer Test (Figure 19) exercises 500 simultaneous outputs on each part. We wanted all the outputs running at the same voltage level for this test, but that wasn't possible with the test board architecture, so here's what we did. We configured each aggressor as a 2.5 volt LVCMOS 8 ma driver (*fast*, for Xilinx, but no speed option for Altera). The outputs are physically powered by a mix of voltages including 1.5V (for at least 100 aggressors nearest the victim location), 2.5V, and 3.3V. Between the two parts, each voltage rail powered the same numbers of outputs, in approximately the same positions.

Final Test: *The Hammer*




Figure 19—The final test exercises 500 outputs on each part to maximize crosstalk.²

The victims are, again, A10 and B7, this time stuck-at-high. The stuck-at-high setting displays BGA crosstalk along with some of the noise extant in the power system.

The stuck-low waveforms show the same high-frequency (short term) spike of crosstalk, but lack the resonant behavior visible in the power system after the main pulse. This resonance provides some clues about the efficacy of the overall power system bypassing network.

In this test the Altera part produced crosstalk waveforms 4.5 times larger than the Xilinx part.

The ratio of crosstalk voltages we measured (4.5:1) is not all due to differences in packaging. Both packaging and signal risetime contribute to this factor.

Figure 20 details the aggressive waveforms produced during the Hammer Test (measured at locations A11 and B6, respectively, on the Xilinx and Altera

² Errata: horizontal axis corrected to read: 50 ns/div. The test repetition rate is 9 MHz.

packages). Given the same driver settings the Altera part generates a current slope (di/dt) twice as large as the Xilinx part.

Rise/Fall Times




Figure 20—The ratio of crosstalk voltages we measured (4.5:1) is not all due to differences in packaging; both packaging and signal risetime contribute to this factor.

This 2:1 ratio of di/dt translates directly into more crosstalk for the Altera FPGA. The variation in power/ground pin outs accounts for the remaining portion of Altera's 4.5x increase in crosstalk.

Final Simulations

Now let's leave our measurements behind and go back to the simulator.

What I would like to do is address any concerns you may have about our selection of pin locations for testing, or other factors in the layout that may have skewed the results.

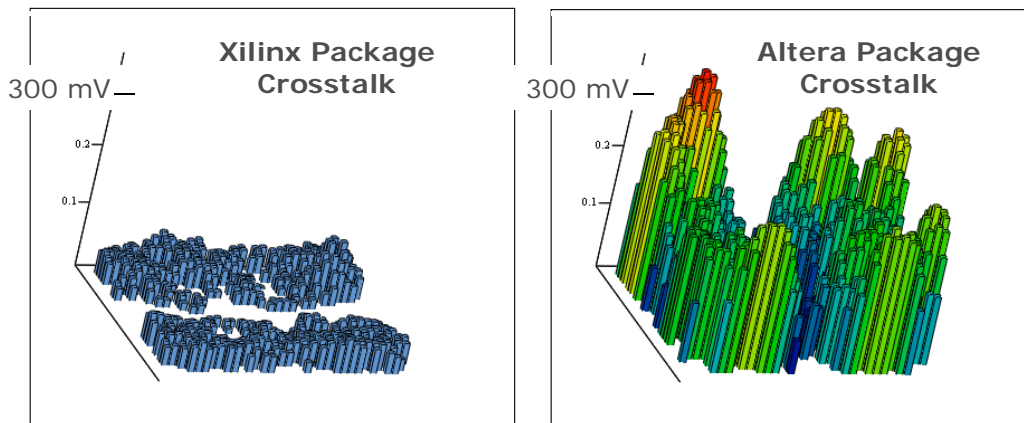
I will set up an ideal crosstalk simulation, making the following selections:

- Ignore exact pattern of trace layers, assuming an average trace depth of 0.035 in.
- Ignore differences in rise/fall time, assuming both parts produce $di/dt = 2E+07$ A/s.
- Ignore details of dog-bone offset, assuming a via-in-pad geometry for simplicity.

Under these conditions, Figure 21 displays the worst-case aggregate crosstalk for *every pin on both devices*, not just the few pins instrumented with SMA jacks.

This chart pinpoints only the differences you would expect to see due to variations in the power/ground ball distribution between the two packages – not taking into account the inherent differences in di/dt due to different rise/fall times. It assumes all outputs are switching 1v p-p into 50 ohms with a rise/fall time of 1 ns.

Simulated Package Performance



- All outputs generate $di/dt=2E+07$ A/s
- Consistent via depth of 0.035 in.
- Simplistic via-in-pad geometry

Figure 21—This chart pinpoints only the differences you would expect to see due to variations in the power/ground ball distribution between the two packages – not taking into account the inherent differences in di/dt due to different rise/fall times.

Taken together with the BGA crosstalk theory from the beginning of this presentation and my actual measurements corroborating the crosstalk effects, Figure 21 paints a clear picture of the differences between the two packages under study.

I hope this presentation has been as interesting and informative for you to read as it has been fascinating (and challenging) for me to produce. If I have stimulated your interest in researching the problem further, I can suggest these related articles:

www.sigcon.com/Pubs/edn/DataCodingLowNoise.htm

www.sigcon.com/Pubs/edn/TimeforAllThings.htm

www.sigcon.com/Pubs/edn/assymnoisemargins.htm

www.sigcon.com/Pubs/news/3_9.htm “Crosstalk and SSO Noise”

www.sigcon.com/Pubs/news/7_10.htm “Scrambled Bus”

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, seminar course outlines, SiLab films, newsletters, Article indexes, and much more.



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.

Conclusion

The Altera component used in this test displayed crosstalk 4.5 times higher than the Xilinx component.

The Altera package suffers from two issues:

- Excessively fast signal rise/fall time
- Over-concentration of power/ground balls in core region

Together, these two effects combine to produce a 4.5:1 ratio of observed crosstalk in the Hammer Test (Figure 19).

Howard Johnson, PhD, author of High-Speed Digital Design and High-Speed Signal Propagation, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. www.sigcon.com, howie03@sigcon.com.