Achieving Superior Signal Integrity for Breakthrough Performance

Webcast Date:_03-01-2005

Presented by:
Dr. Howard Johnson

Peter Alfke, Director, Applications Engineering

PETER ALFKE: Good morning, on the Pacific Rim here. Good afternoon further East and good evening in Europe.

SLIDES 1-2
This is the fourth in our series of Ninety Nanometer Web Seminars and your participation is steadily increasing. Thank you for signing on -- more than a thousand of you signed on. When we started the Virtex-4 family, we asked our customers about their main concerns and we heard many obvious things: ever-shorter design cycles, faster obsolescence, competitive price pressure, increasing complexity.

SLIDE 3
We, here in Xilinx Engineering, struggle with the same issues. You also told us that you need higher performance and our first Web Seminar addressed that issue. You are concerned about power consumption and we covered that in the second seminar. You're concerned about signal integrity on your PC board and that is the subject for today's Seminar. Chip output transitions are getting faster and the higher the VDT and the IDT [] generate vortex spikes across the inductances in the package and on the PC board. There are many names for it. It's called "ground bounce," "crosstalk" or limitation of the number of "simultaneously-switching outputs" -- SSO. This is a really tricky issue with potentially very nasty effects on system reliability, but it's analog -- and even worse, it is ultra high-speed analogue -- and the average digital designer is not familiar with exotic things like that. Engineers who'd rather work with a logic analyzer than a scope feel uncomfortable with little, and sometimes not so little, analog spikes. Investigating and debugging this requires a specialized knowledge and fancy equipment. We just exchanged our $65,000 scope for a $100,000 scope because we had to. You will see some screen shots from that kind of instrument: Beautiful, crisp pictures, even when they show ugly effects.

SLIDE 4
Most digital designers want to stay away from these details. They want a solution that works and that they can trust. Well, we worked on that and we designed new packages for Virtex-4 parts with radically changed power and ground board assignment. The idea was to minimize the inductions in the path of the output-switching current. To measure our results, we built a board that uses our Virtex-4 and the competing Stratix-II chip, both in exactly the same way. And to increase our confidence and, thus, your confidence, we worked together with a well-known and respected expert in the field of high-speed engineering, Dr. Howard Johnson.

You may know him as the author of "High-Speed Digital Design: A Handbook of Black Magic," considered by many of us the bible of that subject. He's a signal integrity columnist at the "End" (sp?) Magazine" and he is a frequent guest lecturer at Oxford University. That's actually where I met him for the first time.

So, Howard, what can you tell us about our improved packages? Was it worth the effort? Tell us about your evaluation, please.

DR. HOWARD JOHNSON: Thank you so much, Peter, for that kind introduction and for inviting me to speak today to your audience.

SLIDES 5-6
This talk is going to be about measurements of actual crosstalk on BGA packages, so I think I should just begin by diving into the test set-up to show you what we did. The nature of our experiment was simply to take multiple aggressors, which you see in the bottom of this picture, all coming out of one BGA package, bang them up and down, making a lot of noise all at once -- actually exercise hundreds of outputs at the same time, each one going onto PC-board traces that were terminated within the terminations. This draws a lot of current out of the device and should create a great deal of noise. Then, we took a few of the outputs and left them in a condition where they were either stuck high or stuck low. That's what's noted as "the victim" in the top of this chart. We took those and routed them out of PC-trays to a very fine SMA connector and into a semi-hardline SMA cable, through a DC block into our oscilloscope. And at that point, we were able to see the crosstalk that results on one poor victim who tries to stay, well, either high or low, while all the other outputs are banging up and down.

This technique gives us a realistic, worst-case appraisal of the

crosstalk that you would see coming out of an I/O that was stuck high or low. It also, as you'll see in a minute, gives you a good appraisal of the crosstalk you would receive on an input pin of the same chip during similar conditions.

SLIDE 7

In order to understand how that works, I think we should take a moment to look at a little theory. I'll show two slides here: One is a basic crosstalk theory associated with FPGAs. This is the way that we thought of it when we had one ground pin on our old [zip?] packages a long time ago. The effect was known at that point in time as "ground bounce," and the diagram illustrates here -- at the top and bottom of the diagram, the green bars represent a power plane at the top and a ground plane at the bottom. Now, the chip, which is a BGA chip -- it shows four balls and there's a die in the middle -- is between the planes in the drawing. Of course, physically, it's not between the planes; it's on top of the board, but I think you get the idea in this schematic view of what I'm trying to get across.

So I have a chip that has two I/O cells inside it in this drawing. They're both connected to the same VCC ball up to the power plane and through the same ground connection down to the ground plane. And the idea here is when Switch C, on the right-hand side, when it closes pulling Output F down low, you immediately, on that transition, draw a huge surge of current coming from the load in through Ball F, in through Switch C and from that point down to the ground plane, returning back to the load. Remember, a chip is not a vacuum cleaner for electrons. It doesn't just suck them into a big bag and hold them somewhere. No, it's just a switch. And when the switch closes to the extent that there are different potentials between Point F and the ground, it allows current to flow through that pathway coming in F, going out at the switch at C through the ground pin.

The important aspect of this drawing is that when you have a surge of current going through the inductance of the ground pin -- and, of course, any metallic conductor structure always has inductance -- when you surge current through that inductance, you create little voltage differences across the ends of the inductor. That is the voltage on the substrate of the chip temporarily becomes different from the voltage on the PC board ground plane, down underneath the chip. If you think of the PC board ground as being perfect, then you'd say there was going to be noise on the ground substrate. And that ground substrate noise then, if you look at how the switch on the left works,

noise on the substrate goes straight through Switch A and comes out of Output D. That creates a glitch that we think of as "ground bounce."

Now, this simplistic view of the problem is qualitatively good at explaining why switching events on one pin can cause noise on another, but it is quantitatively poor at predicting the noise amplitudes. To see the flaw in this diagram, I would like you to abandon the concept of inductance as the property of an individual wire or conducted pathway. We need to perceive instead inductance as it truly exists. It is a property of the space between conductors. In any high-frequency inductive problem, the relevant magnetic field resides in the space between the conductors, not in the conductors themselves. It is this field, not the conductors, that causes inductive effects.

You know, in an analog circuit, if you're used to working with an inductor that is a tightly-wound coil, in that coil configuration, you concentrated the magnetic field in one spot where it becomes incredibly intense. Outside the body of that inductive component, the field is much less -- it's small enough to ignore. Therefore, it is correct in that sort of circuit to think of the inductances belonging inside the circuit elements. But the inductances we deal with are different. We don't have coils. We have inductance that occurs as parasitic inductance having to do with straight wires. In that case, the magnetic fields are spilling out all over the place in the space between the wires and it's those fields that determine our inductances.

SLIDE 8

The next diagram shows a better way to think about the crosstalk problem in a BGA package. Now, in this diagram, what I show is, the chip package is on the top; it's blue. That's representing a BGA substrate. Somewhere there's a die in the middle of it. And in green below that is the ground plane of your PC board and I've tried to illustrate four balls coming off the package, going through vias down to the ground plane. Now, I left some things out. For example, there's no dog-bones underneath the balls that you would really have. But I just simplified the drawing in that way.

The essence of how crosstalk really works in this BGA package is if you imagine Switch C, over on the right, that's the low side of a totem pole. When it switches low, it draws a surge of current in through Via F, which is coming from some trace hooked up to something to the right. It draws current in through Via F, up the ball, through Switch C, on to the ground substrate of

the package and out the ground pin of the package, back to the PC board from which that current returns back towards the trace exiting to the right. The loop of current is illustrated as that red dotted line, and everywhere inside that loop you have created a magnetic field.

Now, I've partitioned the magnetic fields into three pieces: I call them L1, L2, and L3, according to the three cavities placed between the sets of vias. And the idea is that crosstalk for Via D -- that's the first one, "D" as in "dog" -- crosstalk for that one only to do with the total magnetic field, L1, penetrating the space between the ground via and Via D. That's what's called Faraday's Law. We have a changing magnetic field that penetrates some conductive loop that creates voltages.

If we look at Via E, the total magnetic field between Via E and the ground is equal to L1 + L2. It's more field; therefore, there's more inductance and there's more crosstalk. The crosstalk at E has to be bigger than the crosstalk at D, and that's the most important thing to remember out of this diagram. Specifically, crosstalk is a function of how close your signal lies to the nearest return pin. In this case, I showed a ground pin as the return pathway for current.

This also highlights the problem with the previous diagram. In the previous diagram, we ascribed all the inductance to just the ground pin itself and, therefore, it didn't matter where a signal pin was located; you'd have the same ground bounce effect. But, in a real, modern FPGA, when you measure crosstalk, you find out that it is a very strong function of how close you are to the nearest power pin or ground pin.

Now, this type of magnetic coupling varies with a couple of things, I should point out: It varies with the height of the BGA package above the ground substrate, and there's also crosstalk of a very similar nature between the vias, themselves, underneath that ground plane. That is, if the board is, say, 100 mills thick, there may be another hundred mills of a space down below there where the vias are interacting creating the same sorts of magnetic crosstalk that you have between the BGA package and the first ground plane. Both effects operate in basically the same way. So the total length of your connection, that is, how deep your signal goes down into the board has a big impact on how much crosstalk that individual signal will receive.

Now, for those of you who may be wondering about capacitive

crosstalk, of course, capacitance can create crosstalk. But I will point out, it contributes very little to this set up. You'll see that in the measured results in a few minutes. If there were significant capacitive crosstalk, for example, then, when Pin F switches low, you'd see a negative glitch on Pin E. And when Pin F went high, you'd see a positive glitch on E. That's the way capacitance crosstalk would have to work.

The inductive crosstalk is opposite that, you'll see, when Switch C closes and the via on the right side, F, goes low, you actually get a positive glitch on E, and that's the proof that we're dealing with an inductive effect. Since the effect is inductive, we can model it as a little transformer, and I show that in the next picture.

SLIDE 9

The correct circuit model for thinking about this form of BGA crosstalk is to imagine a transformer. I'm showing one with three windings that are all connected to the same, well, it's an air cord in this case. And you can imagine if the signal in the middle, the aggressor, were, suddenly, to transmit a signal to the right, headed towards its load, it suddenly has current traversing that inductor, and that creates a little voltage across the ends of the inductance of that pin that's positive on the left side, negative on the right side. And if we consider what that's going to do to a victim -- we show the victim at the top -- the victim is either in a stuck-high or a stuck-low state, either way, as long as it's a low-impedance connection to the return plane, then what you'll see is the secondary of this little transformer, (break in tape) as shown in the top part of the circuit will react. The left hand goes positive, the right hand goes negative. Since the left hand is connected to a fixed voltage, what you get is a negative volt propagating away from the driver. This explains why we expect the voltage to be negative when the aggressor goes high.

All right. Down at the bottom, I show a receiver, and this is where I illustrate the fact that if you have a receiver with a high-impedance input and it's driven by something on the right side of the circuit, the voltage from the driver on the right is a fixed voltage, our transformer, therefore, creates a little positive blip going into the receiver. The crosstalk blip, in either case, is about the same size but it has different polarities.

OK. Enough theory. I think we should go on and look at the lab setup and find out what measurements we did and how they turned

out.

SLIDE 10
Here's the picture of our test board. It shows us in my lab here in Washington state. I live about 200 miles northeast of Seattle, way up in the Cascade Mountains near Canada. Mark Alexander, the board designer, dragged himself all the way up here with that board to do some measurements. You can see in the close-up view of the board, on the left-hand side of the board, there's an Altera Stratix-II part, the Model 2S60. It comes in an F1020 package, which is the 32 x 32 array of balls on the bottom of the package. The Xilinx Virtex-4 part is on the right. It's a Model LX60. It comes in an FF1148 package, which is a 34 x 34 BGA package, just a little bit larger package size.

The layouts for these two chips, in order to do our crosstalk testing, were as similar as we could possibly make them. The board, overall, has 24 layers. It's 110 mills thick. It has three I/O voltage regions -- 1.5, 2.5 and 3.3, with each of those power planes sandwiched between ground layers. To eliminate any concern about differences in the power architectures for the two chips, what we did was, we came up with one common power arrangement, that is, one pattern of bypass capacitors that more than met both manufacturers' guidelines for power quality and, then, use precisely the same setup on both parts.

You can also see in this drawing the SMA fittings that are used for viewing some of the particular ports on the chip, also for injecting test signals and programming with the chip and doing other operations with it.

We ran all of our crosstalk tests using only one side of the chip at a time. Although we did test for interference between the sides and found practically no observable crosstalk from one chip to the other because the board is completely partitioned from stem to stern. It's got the two halves. There's no copper touching from one side to the other; not even the grounds touch except through the instrumentation cables. So the isolation between sections is pretty good, if we wanted to do experiments showing both simultaneously in the future.

Now, from our prior discussion, what we expect to see in this sort of configuration is we should see crosstalk vary as a strong function of proximity to return pins. And also it varies as a strong function of the distance between signals to each

other.

SLIDE 11
I think it would be a good idea if we take a look at one of the power ground-pin layouts, just to see what we'll be talking about. This shows the Virtex-4 part, and I'm showing the complete array of 34 x 34 pins. In this drawing, the grounds are marked as dark circles; the power pins are marked as the lighter pink circles; and the signals, all the high-speed signals, are marked as "x's." Assuming that you had perfect d-coupling inside the BGA package, then, the power and ground pins would be equally effective as return conduits, and we would not really have to differentiate between them. We just think of them as return vias. And, actually, in both chips, the bypassing worked pretty well; so that's a fairly good assumption, I think, when thinking about how the crosstalk problem works.

Now, you will notice that the Virtex part has a huge number of power and ground pins spread out all over the package. That is, they are not all grouped together in the core. There's about the same number as you would have if you had done a traditional layout with all the powers and the grounds grouped in the core. But, in this case, we've taken those pins, or Xilinx has taken those pins and distributed them out in an even regular array, all over the surface of the pattern.

Now, what I'd like to do is, look at crosstalk, into some of the victims, and I think one particular victim, you'd edge the pattern. If you do a lot of crosstalk testing, you'll discover that the edges are almost always the worst places. So I picked what looked to me like the worst spot on this chip -- it's marked in red up at the top -- and we're going to observe that point while we're exercising various aggressors near that point. Now, this first picture I'm going to do with a simulator, and I'm going to exercise one at a time each of the pins around that victim location and see how much crosstalk it generates and, then, we'll draw a plot to show how the crosstalk works.

Before I leave this drawing, though, I should say one more thing about it and that is, you'll notice on the left and right side, there are some triangle shapes that are not populated with high-speed signals. That's because this is the LX60 version, which we selected because it has roughly the same number of I/Os, that is about 500 high-speed I/Os, the same number as the comparable Altera Stratix part.

SLIDE 12
All right. So let's flip to the next slide and I will show you a distribution of crosstalk for our position A Pin, which we're using as the victim. And what you can see is on this scale, crosstalk goes up to five millivolts. The assumption, here, is that we're using 1.5 volt, MOS drivers, width set to 4 milli-amp driver strength, and that's how much crosstalk we get from nearby signals. One of the things that stands out in this drawing is out of the entire array of 34 x 34 balls, only a few nearest aggressors contribute significantly to the aggregate crosstalk at this position.

SLIDE 13
Let's look at some other positions. I show four different places that I picked and, in each case, the story is the same. Crosstalk falls off geometrically as you move away from the A victim pin. Now, the principle here is that if you were to move the aggressor in a straight line away from the victim, getting further and further and further and further away each time, as you pass by each successive return-ball position -- that is, either a power or ground pin -- if you pass by those locations, the crosstalk suddenly falls off dramatically every time you pass one of those. And so if you have a lot of power and ground pins close by every signal, by the time you get three or four rows away from that signal, crosstalk has dropped off dramatically, and that's going to help this chip to achieve a low, overall, aggregate crosstalk rate.

SLIDE 14
Let's take a close look at the power and ground pattern that's used on the Vertix-4. This is what contains that distribution of return current, so it's worth taking a look. The pink dots are power pins, the dark dots are ground pins, and you can see they're at a regular array that is composed of little cells of 10 elements, each. This is a regular tiling pattern. It's called "a pin-row's tiling pattern." The name for the overall combination of power-and-ground signals that Xilinx has selected, they've called a "sparse chevron" power and ground pattern.

Looking at these little tessellating cells, you can see that out of each cell of 10 things is eight signals, one power, one ground. So we've got a ratio of signal-to-ground, a power of 8:1:1, which means that if you think of the power and grounds for purposes of crosstalk as just being return conduits, then what we really have is a 4:1 ratio between signals and returns.

9

SLIDE 15
But if you have a package for [forums?], now this is a measured result. In the measured result, I show at the bottom a nearby aggressor, on a scale of 200 mV per division; that's one of the four milli-amp drivers driving out a 50-Ohm line with a 50-Ohm termination. A 50-Ohm termination, in this case, is actually in the scope, so that we can see the aggressor.

Next door, I have another signal which I'm observing as the victim. I can either leave the signal stuck high or stuck low and, on a scale of 2 mV per division, you can see that crosstalk is as much as -- on the right-hand side, it looks like about six milli-amps where the crosstalk, when the aggressor is going down. A little less than that when the aggressor is going up; that's a hint that tells you the fall time is a little bit faster than the rise time on this particular part.

You can also see in this diagram that the stuck-at-high, stuck-at-low outputs are very comparable, which tells us that there's not a lot of noise difference between the power-and-ground rails inside the chip at this particular moment; that is, the bypassing inside the chip is doing a pretty good job at the moment.

Now what I'm going to do next is just show you a few more aggressors because as long as we're looking at one victim position, we can exercise a number of nearby aggressors -- 1, 2, 3, 4 in succession -- and what you see when you do that kind of test is you see the crosstalk from each individual one.

SLIDE 16
And that's in this next diagram. The first aggressor was close by and generated a lot of crosstalk and as we got further and further away, the crosstalk drops off dramatically. Now, you may be wondering how in the world did I make this picture because we're down to things that are less than 1 mV and any of you who have worked at high-speed scopes know that they have way more noise than that in the front end. And what we did was -- I'm using a Tektronic scope. It's a TDS6804B. It has a 8 GHz bandwidth and a 20 Gs/s sampling head, and it has a fantastic averaging feature. So we're reliably triggering the scope on the beginning of our crosstalk-generating pattern and we repeated the pattern about a thousand times, collected all that data, averaged it all together, and the random noise averages out leaving you with a crystal clear picture of these little tiny crosstalk effects. That's the way to measure small crosstalk amounts.

SLIDE 17
OK. Let's get on next to our measures of crosstalk. I'm going to try three different experiments now, and we'll go through them pretty quickly. The first one is going to be called a "spiral test." It's going to involve 100 outputs. The next is the "accumulating test," using the same 100 outputs but exercising them in a different pattern. And, third, what I call "the hammer test," where we went all out to see how much crosstalk we could make.

SLIDE 18
Before I show you the results, we need to compare the two packages to each other. I've been talking about the Xilinx package for a while. Here's the Altera Stratix-II package and if I remember right, I believe there's a little booboo on the slide where it says, "The Stratix-II F1120," it should be, "The Stratix-II F1020," is the actual number of the package that's used with this particular product.

Anyway, what you can see is that the pattern to power and ground (inaudible) different. On the left in pink, the Xilinx chip spreads the returns evenly over the whole bottom of the package, and the Altera package tends to group them together more in the core. And, therefore, there are many areas in the Stratix-II package that are devoid of returns. What we expect that to cause is increased crosstalk in those areas.

SLIDE 19
The specific pattern of pins that I'm going to exercise for this next test, in the case of the Xilinx chip, the region exercised is illustrated in the yellow box at the top left of the diagram. That shows the nearest 100th aggressors to our victim position, which is Position A-10 on the layout. And I'll be exercising those in a spiral pattern. First, I'll do all the outputs, one at a time -- bing, bing, bing, bing, bing -- everything right close to the victim, and then the second row and the third row and the fourth row getting successively further and further away from that victim as we go.

SLIDE 20
If we look at the Altera part, it has a similar region, a slightly different shape because we've chosen a slightly different location. That's just due to the pin assignment. Some of these parts are not exactly the same. So, we picked on both parts what we thought would be a worst-case location. It looks to me like B7 would be pretty bad on this part, so I

selected it, and this shows the spiral pattern spiraling around Location B7, going out to the nearest 100 output. In this test, all the outputs will be 1.5 volt, LVCMOS standards, width set to 4 milli-amp driver strength.

SLIDE 21

Now, in the case of the Xilinx chip, there is a speed option, and we set them to the fast speed option. In the case of the Altera FBGA, there was no speed option, so we just set them to 4 milli-amps. And here's what we're going to get. Now, the first thing I'd like to show you before we get to comparing the chips to each other is I just wanted to compare theory and practice to see whether our theory of crosstalk matched up in any way with the actual measurement. And here you can see, for the Xilinx chip, doing the spiral test with 100 nearest aggressors, on the left, that's our simulation; on the right, that's our measurement. And the act of doing this comparison, the purpose, really, of simulation, is to achieve what a friend of mine, Terry Morris at HP, calls a "triumvirate of understanding between theory and measurement simulation." That is, if you have a theory about how something works and you make a simulator to predict what should happen and if the simulator matches the measurement, then you probably have a good understanding of what's going on.

And, in this case, I'd say the simulation and measurement match up fairly well. There are a few anomalies. I've been looking into those to find out why. There are a couple of places where there was some extra crosstalk inside the package, itself, we weren't aware of. And, in one case, near the edge of the part, there was a via that was not laid out on the regular BGA grid, that was moved out away from the grid; therefore, further from the power and ground pins and, therefore, experienced substantially more crosstalk than it should have if it had been in the correct position. Those are interesting things to learn when you start looking at the details of how a crosstalk works. But even though there's not an exact match in these lay forms pin by pin, the overall aggregate trend shines through. The crosstalk falls off really fast as a function of distance with this particular part.

So let's put the simulator away now, but we're going to bring it back later to predict aggregate crosstalk across the whole surface of the chip. But for now, let's put the simulator away and go look at some measurements.

SLIDE 22
The first one here shows a comparison between the Xilinx Virtex-4 part and the Altera Stratix-II part. Now I've changed the scale on this. The last drawing was at 5 mV/division, so you could see the waveforms clearly. Here, I've had to change it to 10 mV/division, so I could fit all the crosstalk waveforms on the screen. And you can see the Xilinx part exhibits two features: First, the crosstalk is generally smaller and it falls off more rapidly. The Altera part has more crosstalk generally and it persists for longer because the crosstalk doesn't fall off very much until you get to a signal point that's passed one of the return via positions, as the return vias are so sparse in that package you have to go a long way before crosstalk really starts dropping geometrically. So I think you can imagine that that persistence in crosstalk with distance is going to cause substantially more aggregate crosstalk for the Stratix part than for the Xilinx part.

And that's what we'll see in the accumulating test, which comes up next.

SLIDES 23-24
The accumulating test exercises the same exact pins as what we used in the spiral test -- it's the nearest 100 aggressors -- in the same order starting with the newest ones and spiraling out to further and further pins, but the difference is that, this time, the first cycle, only the nearest aggressor fires off -- up and down. Then, in the next cycle, we get the nearest one and the second one up and down. Then, three, up and down; then, four, up and down; five, up and down More and more and more are going together until, finally, at the end of the test, they're blasting 100 balls up and down simultaneously, all together. And what this shows us is, it shows us how crosstalk aggregates as you move away from the victim location.

Now, in this case, we're showing aggregate crosstalk, which is much larger than what you get from an individual pin; so we've gone to a scale of 100 mV/division. And we show, in the top, the Xilinx part and, in the bottom in blue, the crosstalk that we measured at Location B7 in the Altera part. Two things stand out in this wave form: First, well, there's more crosstalk in one chip than the other, obviously. And, second, one of the reasons why there's a lot more in the second chip is because it keeps aggregating at large distances away from the victim. In this chart, for every division horizontally in the chart, every horizontal division represents about 20 aggressors and so if we go from the beginning of the Altera wave form, you go 20

aggressors away, you've still only got about half your crosstalk. You got way more than that to accumulate whereas, in the Xilinx part, by the time you get 20 aggressors away from the victim, you've accumulated almost all the crosstalk you're ever going to get. That is the difference that we get due to the assignment of return balls in these two packages.

SLIDE 25

The final test is what I call the "Hammer Test." Mark programmed the part to exercise all 500 high-speed signals up and down, banging up and down, all together. We also turned up the drive strength. We went to 8 milli-amps. There were some considerations about exactly what settings to pick, and here's what we picked: We went for a 2.5 volt LVCMOS, 8 milli-amp driver. That was the selection we made for both parts everywhere, although due to the architecture in this board, they were not all powered by 2.5 volts; some were powered by 1.5 and some by 3.3 volts. So there was a -- But they were done the same way, the same (inaudible) pins and the same banks of all three, in both parts. And here's what we got. The aggregate crosstalk, the most we could make with our setup for the Xilinx part was 123 mV peak to peak. And the most crosstalk we were able to make in the Altera part, as measured at Location B7, was 572 mV peak to peak. Now that's in a region which is a 1.5 volt logic region and you can imagine that 572 mV peak to peak of crosstalk in a 1.5 volt system is probably more than you would want to have.

SLIDE 26

Now in order to understand precisely why we got this ratio -- it's a 4.5:1 ratio of crosstalk wave forms -- I'd like to break things down into two components because the rise time of the signals and also the packaged geometry have an effect on crosstalk. And, in this case, I made a careful measurement of the rise-and-fall time for both chips and found that the Altera rise-and-fall time -- now this was measuring on a pin adjacent to our victim -- was about a factor of 2 faster than the Xilinx rise-and-fall time; that is, the Altera part generated a DIDT change in current per unit time that was double as much as the Xilinx part, which means that that particular aggressor is generating twice as much noise. And that accounts for a factor of 2 difference in crosstalk between the two packages. But we measured a factor of 4.5:1 which means that you've got another, more than doubling due to the package. So more than half of the effect we've seen is due to the package, about a factor of 2, seems to be due to the rise time in this particular test.

SLIDES 27-28

Now the way my board was architected, there were only a few locations that we could measure crosstalk at. We pre-selected the ones that we thought were going to be the worst based on our simulations and common sense. But what about all those other locations? I'd like to show you that. I demonstrated earlier that we had pretty good correspondence generally between the aggregate crosstalk predicted by our simulator and by actual measurement. And so what I'm going to do is set up my simulator now so that it makes these three simplifying assumptions in order to get a feeling just for what should be the difference between these two power-and-ground pin layouts: First, I assumed that all traces went to the same layer. You know, in a practical system, in order to do your breakout, you had different pins going down to different layers and the exact pattern of what traces go to which disc has an impact on crosstalk. Well, let's ignore that for a minute. We'll assume they all go half-way through a 63 ml, pick 4, just to pick something. Next, I'll ignore the differences in rise-and-fall time.

Let's assume both parts did the same thing, that they both switched a 1-volt signal into a 50-Ohm load with a 1ns rise time, that gives you 2×10^7 amps per second. And let's ignore details of the dog-bone offset and exactly how we did our dog-bones. I'll just assume that the balls go straight to the vias, the vias go straight down into the planes, down to a depth of .035" inside the border.

With those simplifying assumptions, I shall now compute the worst-case aggregate crosstalk that you could ever see for each pin, all over the pattern for each device. And that's the number that theory tells us should happen. It shows a number of features that I find interesting. Of course, the fact that one chip generates a lot more crosstalk than the other just due to the power-and-ground layout is one thing. Also, you'll notice that the crosstalk in the Xilinx chip is more uniformly distributed among pins, that is the variations aren't as great between the peaks and valleys, and that's simply because, with that pin element tessellation, every signal pin is adjacent to some ground pin and to some power pin. They're all next door to something good, and that tends to reduce the crosstalk for everybody.

In the Altera part, predictably, the worst crosstalk places are in the middle of those regions that are devoid of power-and-ground pin.

SLIDE 29
Well, I hope this talk has been informative for you in terms of your understanding of crosstalk -- why it happens, how to measure it, what some of the measurements turned out like and, also, what factors in the package design influence crosstalk.

If I have stimulated your interest in researching this problem further, I can suggest a few related articles that may be good for you to read. They are located at my Website, which is at www.sigcon.com. And the five articles I picked out here are specifically related to crosstalk effects in large chips. Also, on my site, you'll find my schedule of public courses, my newsletter, if you're interested in that. I write once a month about some issue of interest to high-speed digital designers, and I hope you find that entertaining.

Lastly, down, here, at Xilinx, there are two specific places you might want to check on their site: First, www.xilinx.com/signalintegrity. That's the main resource page for high-speed designers where Xilinx is starting to pack in a lot of useful information. And, specifically, at their store, I have done a tutorial for the Xilinx Rocket I/O Serial Transceivers, which is now available on DVD, if you'd like to check that out. That's where you go to find it.

Well, at this time, I would be pleased to consider any questions or comments you may have but before I'm able to do that, I need to turn things over for just a moment to the TechOnLine Studio. Peter Dobisz has a few words he'd like to say to you about how to input questions. Peter?

PETER DOBISZ: Thank you very much Howard and Peter. At this time, we will move into the Question and Answer section of this presentation. If you would like our speakers to answer your questions live, please submit them now. Please take the time to open, fill out and submit the presentation survey. You can access the server at any time in the Print Documents and View Links, pull down on the left-hand side of your interface. This survey will also pop open when you choose to close your viewer window or when the viewer window closes automatically at the end of the Web cast.

And, now, let's go back to our speakers for the Q&A.

SLIDE 30
Q & A

HOWARD JOHNSON: I've got a host of questions in front of me, and I'm going to just pick some and begin answering as much as I can. I'll read each one as it comes up. The first question is: "Hi. Is the current loop the same for switching to VCC? Is that what you call VCC bounce?"
And I think what I'd like to say is, if you go back to the simplistic view of cross talk that I showed at the beginning of this talk. I showed ground bounce. Well, in the simplistic view, if you have an output that switches high and if you consider the inductance of the power pin, you'd see VCC bounce; that is, when you switch high, you get a glitch on the power rail. Now what complicates things, in a package like Virtex-4 package, is that they have a lot of bypass caps internal to the BGA package itself, which help a lot with your power noise. And what that does is it clamps the power-and-ground rails together, so that as far as high-frequency noise is concerned, to first order, you could consider those two planes to be one and the same thing; that is, all the power-and-ground pins act together in concert as they come in return pathway for signals traversing through the BGA ball field, so we don't have to separately consider them. You get current traversing power-and-ground pins -- whichever one's closest; that's where most of the returning current goes. And, then, once it gets into the BGA package, it goes through a bypass cap, gets on the correct plane, returns back into either the high side or low side of your totem pole driver, depending on which way you were driving.

Well, let's see. What else can I address here? What am I defining as a high-speed signal?
Well, heavens, if you get to 1ns rise times and you don't think about the pattern of power-and-ground pins on the bottom of a big BGA, you'll have so much crosstalk, it doesn't work. So I'd say 1 nanosecond is certainly high-speed enough to create a lot of problems, as many of you, I'm sure, know.

OK. Next. G. Cowart Cisco asks: "Why was the victim pin chosen on the edge of the package, not somewhere in the middle where it could be surrounded by a lot of toggling I/O pins?"
Well, if you try in the middle, you're correct, you have more things around you generating crosstalk. But you also have more ground pins than power pins surrounding you on all sides. And it's not intuitive really but it does turn out that the worst performance is almost always near the edges. How else could I

illustrate this to you? If you were to imagine a sandwich made of two flexible planes, separated by, you know, sponge or foam, and you imagine yourself embedded between these plans with a rope, standing on the bottom plane and pulling on the top one with a rope and wiggling it, trying to make noise between the planes, when you get to the edge, it's easier to make noise. That's one visualization that I found a little helpful in seeing why noise problems are almost always worst near the edges of arrays.
If you wanted to fix that, by the way -- and I'm beginning to toy with this idea -- there's crosstalk that happens in the BGA ball pattern, itself, which about 20 mills tall and there's equivalent crosstalk that happens in the via field underneath the BGA package, which can be a lot taller, depending on how thick the board is. That is, you can get more crosstalk in your vias than you get in the BGA balls, themselves. Well, if that's the case, then, maybe we should be thinking about going to the edge of our array and on the PC board, adding a few more power-and-ground pins connecting together all the power-and-ground planes in your board. Just one more row of power-and-ground pins outside the edge of the BGA package. That is, if you look at the Virtex-4 power-and-ground layout and you see that sparse chevron array of where the power-and-ground pins go, extend just the ground pins even, ground and powers, extend them out one more row of power and ground. So it'll be one out of five pins will, then, be a new via that you'll have to add out there, and that would reduce the crosstalk in the worst cases near the edges and make them more like the ones in the middle.

OK. I have another inductance question here. Someone says, "Why did you spread the grounds? Doesn't that impact loop inductance and create more problems?" Well, the idea is that where current is flowing is actually -- If you think of an I/O, it's going out the I/O pin and comes back on the nearest power- or ground-pin and if you put that power- or ground-pin closer to the I/O, you've reduced the size of the loop over which current flows and, therefore, you've reduced the total amount of magnetic fields you've generated under the package and reduced crosstalk. So it's a good thing that we or that dialects spread them out in that package.
And by the way, I should point out, this is not a unique idea. Many, many chip designs, if you look in high-science servers and so forth, and you look at the custom chips that are designed for that application, I have seen designs where, out of a thousand pins at the bottom of a package, 650 of them are power-and-ground. If you want to go faster, the faster you go, the more power-and-ground pins you need more evenly spread out over the

entire surface of the package.

Let's see. Can I explain in depth how it is that both the power-and-ground pins can act as return currents or as return conduits in the package? And, boy, I wish I had a good picture that I could show you for that but I'm going to have to refer you to a description because to see the whole thing and draw it out is kind of complicated. But the best place I know of to go is on my Web site at <sigcon.com>, if you look under publications, I have three indexes there: There's one by name of article, one by chronology and one by keyword. And if you look under the keyword "return current," you'll find several articles that deal with return-current situations on multilayer boards of various, different types. And it addresses this idea of how it is that return current could be flowing on either the power plan or the ground plan or coming through a power pin or a ground pin, and the effect of bypassing on that situation. In particular, I wrote some articles for EDN recently on "The Effectiveness of Bypass Caps Inside Chips for Both End-terminated and Source-terminated Configurations" that I think you might find interesting.

Let's see. Next question. "Does the chip-package ground plane affect the return (inaudible) pass by the amount of crosstalk?" Oh, yeah. Yeah, I'm assuming that in order for the BGA balls and the vias to be the dominant source of crosstalk, we assume that, inside the BGA package, you've got a solid ground plane, a solid power plane, really close, to each other and will bypass inside the BGA package. The Virtex-4 certainly has that, and that's what allows us to model crosstalk in the way that we have and the same general situation for the Stratix part, if I understand correctly.

How many I/O banks are involved in the Virtex-4 and the Altera part?
Interesting question, and my recollection is the Virtex architecture accounts for eight, but I better ask Ponch, Chandra Sekker at Xilinx, who would have possibly a better answer to that. Ponch, are you on the line?

__: Ponch is not but I can get him on the line. "How many layers of --
Q: "How many I/O banks are involved in the Virtex-4 and the Altera part?"
__: "How many I/O banks --
Q: "How many I/O banks?" Yeah.
__: -- are involved?"

And Mark Alexander, if you're there, you may know the answer to that. When we get an answer, we'll let you know, but I'm going to go on to another question for the moment. This is a really good one. OK. "Is there a reason that Altera grouped the power and ground near the core? Is there some trade-off or advantage to this that you have not discussed?" It seems like Altera could have distributed the power and ground but did not and they're wondering why.
And my, you know -- I cannot speak to someone's motivations for why they did something but I can observe that some designers feel that grouping the power and grounds in the middle would make it somehow easier for you to do your breakouts in terms of breaking the traces out, with the fewest number of layers, going to all the places that you have to go, and that's certainly an important consideration. And what's happened here is that some of the power and grounds balls instead of being exactly in the middle had been moved over a few rows -- or over a lot of rows in some cases. But if you really look at how that works, each time you popped a ground ball over one position, you could have taken the signal that went to that signal position and routed it past that ball and gone one deeper; that is, it doesn't look -- It's not apparent to me that it takes extra layers in order to break out from this Virtex-4 pattern as compared to the Altera. I think there may be an assumption, in some people's mind, that it does but that did not appear to me to be the case -- at least, not in the layout that we had on our test card. They both took the same number of layers to do the complete breakouts, and we broke out every pin, by the way; they all were going to terminators.

Let's see. Next question. "Are the power pins you show are just VCCO? Is that correct? Not core power but just I/O power?"
Oh, good question. On the drawings that I showed that illustrated the power-and-ground pin locations, I took a very simple approach. I said, you know, "All the grounds are the same. In fact, ground is ground on both packages, so I've not noted all those," and I said, "Any kind of power pin -- I don't care what it is -- I'm going to market as a power location. And in my simulator, I'm going to simulate it as a return pack, that is, I'm going to assume that, inside the package, it's well bypassed the ground and, therefore, would be an effect of return path." And those are the assumptions I made for simulation purposes. Now, for measurement purposes, the power pins are what they are. I mean, you can't change them, and it is true that you get slightly different results depending on whether you're next to, you know, a 3.3 volt power or a 1.5 volt power -

- or whatever -- depending on how well those particular power pins are bypassed to the ground plane. What I discovered in the process of doing this project was that both parts had really good bypassing inside the part and as far as I can tell, the differences between -- in terms of differences between measurement and simulation -- I did not think that I needed to take into account the difference between those various different types of power balls in order to explain everything we needed to know about what was working. Also, you may be concerned that, "Well, I had core power involved. And, yeah, I did, but the core power pins were almost always located near the center of the device layout and they were remote from the victims and aggressors that we were exercising during our tests. So, you know, whether you assume that they are good returns or not, it doesn't really make any difference in your calculations because they're so far away, they weren't having much of an impact.

Let's see, my screen is blinking here. It's taking me a minute to get back to -- Let's see. My "Hammer test shows a horizontal scale of 500 ns/division. Is that an error? Should it have been 500 peak/seconds per division?"
OK. Let's go look at the picture. If I flip back to that one for a minute, to the hammer test and -- It should come up here in a second. And that is -- 500, six divisions. I'm thinking about this for a second. Nope. It's 500 ns per division. What's happening is the little short spike -- the little -- pptt! -- just lasts a few nanoseconds. I used a really thick line in the drawing, so you could see what it did. If I'd used a line commensurate with its actual width, it wouldn't even have shown up in the scale graph. And, then, after that little spike, which is the BGA crosstalk, then, what you see -- the booyyyaaaaang -- so the boing part, that is the reaction of the power plane. This particular picture is illustrating noise for signals which are set high. And so you see the crosstalk first, and, then, you see how the power plane is wiggling. And the power plane noise is partly noise on the board and partly noise internal to the BGA substrate, itself. And I wanted to document that. I also have the wave forms for stuck-at-low signals which show the same big -- pptt! -- spike, but don't have the boingy power-system response following. It looks to me like, in the Stratix part, when everything switches off, there's a pretty high key resonance in the power system the way it is. And the Xilinx one shows a similar effect but at a more muted amplitude. OK?

Let's get back to my questions. The purple buttons to hit here, and I think what I'm going to do is take about two or three more

questions and then Peter Elski (sp?) has a little wrap-up that he'd like to do. Let's see. Next question. "If the worst-case noise occurs at the edge of the device, then why in the world are the rocket I/O MGTs placed near the edge?" (laughter) That's a pretty good question. I like that. OK. So here's the deal. The Rocket I/Os are differential outputs and you get the signal and its mate right next door. And so the crosstalk they pick up has to do just with the spacing between those two. The other power-and-ground pins are not as related to that problem. They help a little bit, but it's mostly just the space in between those two which is as close as you can get them, no matter where they're located on the layout. So that's one factor. And the second factor that you'll notice if you look at the Rocket I/O Guidelines -- I went through this in my movie recently when we were doing a tutorial on how to lay that out -- the recommendation is you lay those out on the micro strip layer, which means the via length is zero; that is, you've got a ball that goes down to the board but then there's no via penetrating down further in the board. That means you don't pick up any additional crosstalk. So that's about the minimum crosstalk configuration you could have is a pair of things next to each other going straight down to Layer 1. Also, putting them on Layer 1, at that point, the crosstalk was low enough that everything worked and by having them on the edge, then, you don't have to worry about different trace widths in the breakout region and on the board itself. For example, if you want big fat traces to go a long way on a thick board, you could do that straight out of the edge of the rocket I/O transceiver, if you wanted to.

Let's see. "What software package did you use to perform your simulations and how are the simulations set up?"
The simulation of crosstalk, in this case, is a fairly straightforward calculation of mutual inductance, assuming that you had pathways for current which are vertical, with a fixed radius -- that is they're cylindrical and are going down to solid planes, they're encapsulated between solid planes, top and bottom. And that's something where an experienced field theorist can write down the equations and constraints for how that works in any mathematical spreadsheet. You could use MATHEMATICA, you could use MATHCAD, you would use MATTLAB. Those are three good examples of how that could be done. For simple configurations, if it's only two or three vias, you can do it by hand. But for 1,000 vias, that was a little more than I could do by hand, so I wrote the equations in MATHCAD and did the simulations with it. I chose that particular tool not because it's any better than the others mathematically but

because it makes really cool graphs and pictures, and I like that for presentation work. So I used MATHCAD. Let's see. What else do I want to get to? We're going to do one more question here.

"If you placed additional ground of VCC vias just outside the BGA packages, will this limit the noise on the Altera parts?" Well, I think it's a pretty fair question. I mentioned that extending the ground, you know, if you've got a lot of ground planes in your board -- and we're talking about crosstalk in the via array, penetrating down through the board -- wherever you have signal pins going through the board, you need returns to connect the planes together. Just like if you had a connector. If you had a connector between two boards, you know that you've got to have a lot of ground pins on that connector to make it work. Well, actually, the vias are nothing more than a connector that goes between the top layer and the inside or bottom layers of the board. And they need a signal-to-ground ratio, too. If you don't have enough grounds, you get crosstalk in that region. So for any layout that you ever do -- wherever you are -- adding ground vias in your layout [] intersperse the signals will reduce the crosstalk you pick up between those vias. Now, how much crosstalk are you going to get if you're going with sub-nanosecond rise times and putting big busses through the board, that's when you're really going to notice that kind of crosstalk, and you've got to think about it. So --

"Would that improve?"
It would improve any chip having -- the specific thing that would happen if you add ground vias around the edge is, then, the edges would not be that much worse than the middle of the array. So if we were to go back to the comparison between parts that I cited here in Slide 28, "Simulated Package Performance," and we look at where the worst cases, for both parts, they're around the edges. And if you had more ground vias around the edge of the board, that would help somewhat. At least, it would reduce the crosstalk in the via region; it wouldn't affect the crosstalk in the BGA balls. But vias being the larger part of the problem, it would help to reduce some of those peaks around the edge and bring them down to be more like the things in the middle at which point you'd still see a difference between the two packages. Yeah, yeah.

All right. That wraps up what I wanted to say today. I would like to thank you very much for your attention and for all the good questions you were able to submit. I'm sorry I was not able to address every question. If you have something that you want

to know about, I do my best to answer questions by e-mail. I can't guarantee a timely answer to everyone, but I try my best: At www.sigcon.com. You're welcome to go there. There's a way you can contact me and send in a question and I'll try to answer it for you. I look forward to hearing from you on the Web.

Now, Peter Elski has a few comments he'd like to make, to wrap up today's presentation. Peter, are you there?

PETER ELSKI: Yes, I'm here. Thank you, Howard. I think this would qualify as a technical presentation, not a marketing presentation. When I advertised this on my favorite stomping ground [CompArcFPGA?], I said that Howard Johnson is not only a well-known and respected expert, he's also a very lively speaker and a good teacher. And I think those who listened to you for this, almost an hour, would agree with that.

SLIDE 31
So let me just do a faint little summary here at the end. So we show that Virtex-4 has vastly improved signal integrity and reliability. The sparse chevron pin minimizes the inductance in the current loops. Simultaneous-switching outputs cause now less ground bounce and crosstalk on adjacent pins -- Howard went through that in very much detail -- which makes it an ideal solution for wide data busses.

Now, at the 30,000' level, you know, what do you really get from all this? Why is it important? You get less noise in your system. "Less noise" means "less jitter." "Less jitter" means you have better timing margins. You have potentially increased the performance of your system and you definitely made it more reliable.

So we did that with the packages, as Howard explained. So we say, "We did the job for you," you know. You can get the benefit, and you don't need to buy an 8 GHz oscilloscope and poke around the way Howard did because the package is actually much better.

SLIDE 32
Now the next slide shows, our advertising slide that shows that we're "The Product of the Year," and Virtex-4 is.

SLIDE 33
And the next one shows you that we would like to invite you to come two weeks from now to the next seminar in the Series that talks about Virtex-4 memory interfaces. And you can also visit

the previous and this seminar by going to the indicated URL. And also, you can get Howard Johnson's report on exactly this seminar on www.xilinx.com/virtex4/si.

So, that's it from here. Thank you

End of Xilinx 03-01-2005