Meeting Your Power Budget with the World's Fastest FPGA

Webcast Date February 15, 2005

Presented by:
Anil Telikepalli, Marketing Manager, Virtex Solutions

Matt Klein, Sr. Staff Engineer, Applications Engineering, Advanced Products Division

SLIDE #1
Hi. I'm Peter Alfke here and I'm very happy to introduce this second seminar in the Xilinx 90nm FPGA design series. So this time we are going to talk about power consumption. The world's lowest-power high-performance FPGA. And I might remind you that we are -- Xilinx is #1 in 90nm FPGAs and at last count we had shipped 100 times more 90nm FPGAs than the rest of the industry combined. Whenever we embark on the design of a new family we ask our customers. We check what are your challenges, what are your problems. And the normal answer we get is that design times are getting shorter, designs get obsolete faster, and we understand that because so do our parts. And you find that you have more competition and cost becomes ever more important. And we understand that because the same thing is true for us also. You say the complexity of your designs is getting very demanding and the performance has to go up and up every time. And we understand that because our designers have the same problem. Then you also talk about the signal integrity issues that you have on PC boards, on ground bounce, all caused by the faster I/O. And that's a subject we will cover two weeks from now in detail. And lastly on this list is your concern about power consumption and thermal issues, about heat things and so on, and that is the subject of today's seminar. We're going to talk about power. And I'm going to hand it over to Anil Telikepalli and he will talk about power issues.

SLIDE #2
Hello everyone. In today's seminar we will go through an introduction session followed by the advantages Virtex-4 technologies offer for power reduction. Then we will take a look at how Virtex-4 stacks up versus Stratix-II. And then move into looking at real design examples, the actual measure data. Followed by tools and resources to help you with your power estimation and analysis.

SLIDE #3

SLIDE #4
So why is reduced power important? As you all are designing your latest products, you are very well aware of power budgets that your systems have and the importance of keeping the power low. Reduced power means fewer thermal concerns, simplified heat management. Oftentimes that means eliminating heat sinks and fans altogether. Now why is that important? Fans tend to be in a large system often the most unreliable component. So eliminating fans by reduced power gives you a huge advantage. The next item is also lower cost. Your power supply costs go down, supply circuitry becomes simpler, you use fewer components which means smaller bill of materials to deal with, as well as overall a smaller PCB. Now not only is it beneficial at the

component and PCB level, but this has a huge impact to your entire system, improving system reliability because you don't have higher power, and hence higher temperature.

SLIDE #5
Now let's take a look at what it means in terms of doing 90nm design and the challenges it puts on semiconductor vendors. Your needs for lower cost with every generation of product drives our process geometry shrink. With every new product you desire to get the costs down, which means we have to drive our process geometry down. Now all things being equal, what this means is with shrinking process geometry, leakage current increases. And hence static power. Now Virtex-4 breaks this trend and reduces leakage current as well as static power at 90nm. Matt will go into details later in the session today to tell you how.

The second aspect of your designs is that designs are getting faster. Frequencies are increasing every generation. Again, if everything is equal, it would mean increasing frequency would increase the dynamic power. Virtex-4 has built-in compact embedded IP again to reduce dynamic power.

Finally, densities are increasing, designs are getting larger, which means again everything being equal, static, dynamic, as well as inrush currents increase. Virtex-4 has built-in technologies to address all of these, which I'll hand over to Matt to talk about. OK.

SLIDE #6
Matt Klein: Thanks, Anil. There are three components to power that I'm going to talk about here briefly before getting into some of the details. So there's inrush current, which is important during the startup of the FPGA. There's static power consumption, which is always present in the device, but when you download a blank design you will have static power consumption. Then when you start operating the design in the third phase you have an additional dynamic power on top of the static power. So we'll talk about this now.

SLIDE #7

SLIDE #8
So as far as inrush current goes, as the power supply voltage rises from 0 to its nominal voltage, internal parts of the circuitry start coming on. Transistors start coming on. These transistors can be in contention with each other if you're not careful. That can cause large amounts of current. Now this may not be much when you have one or two transistors, but when you have upwards of several hundred million transistors, as we do in the case of some of the larger devices in Virtex-4, you have to have ways in your system of mitigating inrush current and mitigating these contentions. Virtex family of parts has taken care of this for many years and consequently we have circuitry that allows sequential turn-on of the internal portions of the device. So sections turn on at a time and that mitigates the amount of current during turn-on. So Virtex-4 doesn't have any inrush issues, but these can be issues in various FPGA technologies.

SLIDE #9

SLIDE #10
So now we'll talk a little bit about static power consumption. Again, this is the portion of the power consumption due to leakage. And when the devices are powered on and programmed with a design or even not programmed yet, there is leakage through transistors. If you look at this plot from the ITRS, International Technology Roadmap for Semiconductors, it shows over time as devices are getting smaller and time is going on that static power is becoming a larger and larger concern. Dynamic power's rising nominally over time, and that's mainly because system speeds are going up. But static power rises dramatically, and it's getting to be quite an issue when you get to 90nm and smaller technologies.

SLIDE #11
Now one other very important thing about static power is that it increases dramatically with temperature. The static power is made up of two components. There's leakage from the source to drain, and that's when the transistor is turned off by the gate. And this leakage is more and more predominant, again as I mentioned when we go to smaller device technologies. And there's also gate leakage, which is a component from gate to substrate. That's important at room temperature. It's actually dominant at room temperature but doesn't increase dramatically with temperature. However, source to drain increases exponentially with temperature. And you can see at 85 degrees C you have two-and-a-half times the static power consumptions. 85 degrees C for a junction temperature is not at all unusual. As a matter of fact, it's a very reasonable junction temperature to be operating at. Looking at 25 degrees C really doesn't tell the story. 85 degrees C and even beyond is important. When I've been talking to customers and looking at various system designs, ambient temperatures of 60 degrees C, 80 degrees C are not at all uncommon. So when you include raising from the case to the junction and even with heat sinks, you still get high junction temperature. Again 85 C is an important consideration.

SLIDE #12

SLIDE #13
The next component is dynamic power consumption. And dynamic power consumption is on top of static power consumption. So after you program the device, got the configuration in there and your design in there, now you turn on your system clocks and you get that dynamic component. An important factor about this that's probably obvious to most engineers is that the dynamic power is related linearly to capacitance and it's related as a square of the voltage. So when we've gone from -- in 90nm technology we've gone from a 1.5 V core to a 1.2 V core. That 20% drop gives us actually a 36% drop in power just by dropping the core voltage. Because the smaller transistors need to operate at a lower core voltage. Additionally, the smaller devices, shorter interconnect lengths, have dropped the capacitance by about 20%. The net effect is that all things being equal, going from 130nm to 90nm we've dropped dynamic power by 50%. I/O power remains about the same because for a given I/O standard driving into a given load, you still have the same power consumption, same frequency, etc. And I'm going to pass it back now to Anil.

SLIDE #14
Anil Telikepalli: Thanks, Matt, for explaining each of the three components of power when you

bring up a system. At this point we will go and look at how Virtex-4 technologies enable reducing power for each of inrush, static and dynamic.

SLIDE #15
So Virtex-4 provides the lowest power in any 90nm high-performance FPGA. Up to 94% reduction in inrush current, up to 78% reduction in static power, and up to 86% reduction in dynamic power. So as you can see there is phenomenal reduction in each of the power components, even with the 90nm technology. Now how were we able to make this possible? Is really with technologies such as power-saving configuration circuitry that we specifically put into Virtex-4 FPGA and industry's first 90nm triple-oxide technology. Different oxide thicknesses to keep the static power low. And finally handcrafted embedded IP blocks customized and put into the FPGA fabric to make sure the dynamic power is low. Let me hand it off now to Matt to go over how each of these Virtex-4 technologies reduce the different components of power. Matt.

SLIDE #16
Matt Klein: Thanks again, Anil. As I alluded to a little bit earlier, with regard to inrush current, this is basically gone in Virtex-4 and it's been gone in many of the Virtex-II parts for quite a while. Xilinx has invested quite a lot of time and design effort into making the configuration circuitry turn the parts on in such a way that the contentions don't occur and that they come on in a logical sequential manner. So there's internal logic that releases different portions of the circuitry at different times. This inrush current had been a problem for SRAM-based FPGAs for a long time, but this is something again for several years that we've had a handle on. Virtex-4 is no different than that. And so we've handled that in this case.

SLIDE #17
If we take a look at the static power, a big innovation that we've made here is that we've introduced this triple-oxide technology. Virtex-4 is the first 90nm FPGA that uses triple-oxide technology. So the problem that we're trying to address is that static power is going up as transistors are getting smaller. This is a basic law of physics. As the transistor's gate length gets smaller, it's easier to conduct a current from source to drain. And even when the gate is off you can't completely turn it off enough and so you get leakage from the source to drain. Normally, FPGAs and other ASICs have two oxide thicknesses that are used. The thin oxide is used in the core of the FPGA or the core of the ASIC and these are used on transistors, which have to be very fast, and they tend to be very tiny transistors as well. The thick oxide is used in the I/O. These I/O have to drive much higher currents and much larger voltages. So they require those. You use the thin oxide when you want the speed. You have combinatorial logic and various things like that. But what we've done in Virtex-4 is we've added a middle thickness oxide. Now one important thing to understand about oxides is that semiconductor industry has been producing oxides of various thicknesses for years and years and years. This is a process that is extremely easy to control very very accurately, because it's based on growing an oxide. You do this by exposing silicon to a certain concentration of oxygen at a certain temperature for a certain amount of time, and you can grow an oxide very accurately to a fraction of an Angstrom thickness. By introducing a third thickness of oxide, which we've done in a number of areas of the FPGA, these small transistors can still have very very low leakage and we've put them in places where we don't compromise performance. By doing this, if we for instance have a million

or a billion transistors and only 600 million of ours are these very very leaky thin oxide transistors and we compare ourselves to someone else who has a billion of the leaky transistors, we're going to have a lot lower total leakage. This triple-oxide technology is endorsed and supported by both of our FPGA manufacturing vendors, UMC and Toshiba. And it's been well-known to the industry that introducing various oxide thicknesses can be a benefit.

Now you may be asking yourself why don't people like Intel and IBM who make ASICs for a living do this type of thing? Well, there's a very good reason why. In an ASIC, every single gate needs to be fast. Every single gate counts. There aren't configuration circuits to worry about. There aren't pass transistors that are programmable, various things like this. So you can't reduce the performance of various items in that. We can take advantage of that as an FPGA vendor.

SLIDE #18
So now let's take a look at static current versus temperature. We have several different curves here. One curve is for Virtex-II Pro. One is for Virtex-4. And one is for other high-performance 90nm FPGAs. You may notice, and we've said in our literature for several months now, that between Virtex-4 and Virtex-II Pro, we've dropped static power by 50%. Now note, this is a plot of static current due to leakage. So you have to compensate for the voltage of 1.5 versus 1.2. And as I mentioned earlier, gate leakage is a major component at 25 degrees C in 90nm parts. That's why you have actually the Virtex-4 having higher leakage at room temperature, at 25 C. When you get to 85 C, you begin to see the difference. And again if you compensate for the voltage difference, which is 25% higher in a Virtex-II Pro, you'll see our 50% reduction. This is based on measured data and predicted data. And you'll notice though that the other 90nm FPGAs are three to four times the static leakage. And the reason why this is the case is because triple-oxide wasn't introduced into that type part. Once you don't have the triple-oxide or you don't have a substantial way of reducing leakage, without compromising performance, there's basically nothing you can do with it after the fact. The Xilinx IC designers had a mandate two-and-a-half, three years ago to design the Virtex-4 with low leakage because we predicted that power consumption in 90nm would be dramatic.

SLIDE #19
The next component that I want to talk about is again this is just a review of what happened with static power going from Virtex-II Pro to Virtex-4 and this is at 85 degrees C worst case. I can't emphasize enough that the appropriate comparison points for measuring static power are at 85 degrees C, not 25 degrees C ambient. At 85 degrees C junction temperature, as I've mentioned many times, is very important and very realistic. So you can see in Virtex-4 for a similar number of logic cells we've again as I stated reduced static power by 50%. And this is dramatic in designs, as we'll see later on.

SLIDE #20
The other area is in dynamic power, and I alluded to this earlier in the presentation. You need to reduce the core voltages when you have smaller transistors. That in itself gives you a 20% reduction in power going from 130nm to 90nm. Additionally, as you can see from this picture with the interconnect geometries and transistor geometries, parasitic capacitances have been reduced and these items both together result in a 50% reduction in dynamic power at the same operating frequency.

SLIDE #21
This is illustrated in this table. So what we show here is that how each component changed as a percentage. So in Virtex-4 you went from 1.5 V to 1.2 V going from 130 to 90nm, that's about a 20% reduction. But it has an effect of reducing power by 36% because of the square law that I mentioned before. Capacitance in total, lumped capacitance used at which these voltages are switching against inside the device have dropped by about 20%. There's a linear relationship between capacitance reduction and power reduction. So for a fixed frequency we have a 50% reduction in dynamic power. Even at higher operating frequencies, we have about a 25% reduction in dynamic power if we increase the frequency by 50%. We plot this because our customers are telling us, you're telling us, that you are going to operate these parts at higher frequency.

SLIDE #22
Another tremendously important thing to illustrate about dynamic power is that hard IP dramatically reduces dynamic power. Why is this? Well, hard IP you can form logic functions of as many variables as you want. You're not limited to what a lookup table can do. Additionally you don't have interconnect pass transistors. And you can minimize the length of the interconnects. Those factors together make hard IP have sometimes upwards of 80-90% reduction compared to their fabric counterparts. So Xilinx from the Virtex-II Pro and now into the Virtex-4 offers an extremely rich set of hard IP. We offer FIFOs, PowerPCs, Ethernet MACs, DSP blocks, source synchronous I/O items that help you with those sorts of situations. All of these things together in your design, if you don't have to use fabric logic, reduce dynamic power. We're going to show that in the next set here.

SLIDE #23
So the PowerPC just in comparing the Virtex-II Pro to the Virtex-4, in the Virtex-4 we've dropped the power consumption by 50%. That's the normal dynamic power reduction that we're getting. And of course compared to any fabric implementation, this is a dramatic reduction, way way more than 50%, on the order of ten to one or more. The XtremeDSP slice is an extremely rich multiply, accumulate block which has other functions also built in such as rounding, max-min, ceiling-floor. All these sorts of things reduce the amount of dynamic power consumed. Additionally, you can cascade extremely long chains of DSPs without using any fabric resources at all. Again another power reducer. ChipSync technology has built in in the I/O serialization, deserialization, timing match, input delay circuitry that's programmable. These items reduce dramatically the amount of power you need compared to fabric resources. The FX family has the embedded tri-mode Ethernet MAC. This is a savings of close to 3,000 logic cells. Our Smart RAM blocks build on our previous block RAM technology and add FIFO controllers. Many of the designs that our customers are telling us they have to build their own FIFOs. This takes a few hundred logic blocks per FIFO. A few hundred logic cells. If you put this across a whole device, it's a big reduction in power and a big reduction in logic. Any time you can use an embedded block, you can use less fabric logic, and this may mean that you can use a smaller FPGA in the first place, which reduces static power and dynamic power. I'm now going to pass it on to Anil to introduce the next section, and here he is.

SLIDE #24
Thanks, Matt. Now we have looked at the key technologies in Virtex-4 that enable lower power. Now when designing the Virtex-4 FPGA, we knew that the investment we made into these technologies and R&D would have a huge impact in reducing power and in turn help you, our customers. Now what we did not know at that time was how much of an advantage we had versus our competitor, Stratix-II FPGA.

SLIDE #25
So we did an in-depth detailed analysis using our FPGAs and our competition Altera's Stratix-II FPGA. We used power estimator tools that both the companies provide you to estimate power for your designs. We used Xilinx Web Power Tool v4.0 and Stratix II PowerPlay tool v2.0. Now again, Virtex-4 data, which we are very familiar with, is based on extensive characterization and again available to you as tools. Similarly, Stratix-II data is based on what Altera provides for you for your power estimation. As Matt has alluded to before, we used 85 degrees junction temperature, a more realistic temperature for your designs, for our analysis.

SLIDE #26
So first let's take a look at inrush current. Matt talked about the three components of power. We'll start with inrush and go to static and dynamic. What we did here is on X-axis you see the devices. Virtex-4 LX devices, LX15 through LX160. And Stratix-II devices 2S15 through 2S180. Here we used devices with equivalent logic densities to compare, and on the Y-axis you see the inrush current. And what you see is up to 94% lower inrush current in Virtex-4 devices. Now across the board, even if you look at every device, there's a huge reduction in inrush current. And we believe that's primarily due to our investment into power-saving configuration circuitry.

SLIDE #27
Now let's take a look at static power. Again, X-axis shows the devices for both Virtex-4 and Stratix-II and Y-axis shows the core static power. Again, on a device-by-device basis, we get up to 79% lower static power. Matt has gone over triple-oxide technology and how this unique technology reduces static power. And again you can see the data.

SLIDE #28
OK. Let's take a look at dynamic power. Here again X-axis shows the functions, like lookup tables, flip-flops, I/Os and so on. And Y-axis shows dynamic power. Now what you see is from a logic or I/O standpoint, both devices get similar benefits moving to 90nm. However, our investment in key hard IP blocks reduces dynamic power significantly for those applications especially where you need these hard IP blocks.

SLIDE #29
So moving on, we're looking at each of the individual components of power, static, inrush or dynamic. Now this gives you an understanding of what it means for each of them. However, we wanted to go all the way and tell you what it means for your total power. Plus, we also recognize that tools and simulation data can be changed as more characterization data comes in. Companies revise their estimator tools. So we wanted to go ahead and actually put a device in the lab and take measurements to see how that compares. So Matt has actually taken both

Virtex-4 and Stratix-II devices on the bench and took measurements at different temperatures. OK. At this point I'll hand it off to Matt to go over design examples that delve into details looking at the total power to operate within the power budget and looking at the measurements that we have obtained in the lab.

SLIDE #30

Thanks again, Anil. As you might have noticed from my bio, I have 18 years of system design experience and hardware design experience. I've designed 20 or 30 boards that have hundreds and thousands of parts on the boards. Meeting a total power budget is something that's very important. Recently I had a design that we wanted to double the amount of functionality in the design. So we had functionality on one whole board that had a certain amount of power consumption. We then wanted to take that entire functionality and double it and put it on the same board. Our goal was to keep the total power the same, which meant that we had to lower the power consumption of each section of the circuitry. So it's quite realistic to have a power budget on the FPGA design. And as a matter of fact, some customers have given us power budgets and said at 85 degrees C worst case power consumption, what will your design do, what will your Xilinx parts do, your Virtex-4, and how will you address my power needs. So in this comparison we've taken a pretty reasonable power budget for a part the size of an LX60 or a Stratix-II 2S60. A power budget of 6 W in a part this size isn't that unusual. In fact, I've had customers who I've worked with from a system design point of view at Xilinx who have said they even want lower power than this for design that's basically 75% full in a part the size of a 2S60. So with a power budget here of 6 W, let's look first at what the static power consumption is. This set of data is from the prediction tools. And in looking at the prediction tools for worst case power at 85 degrees C, worst case over process, that is, the Stratix-II is already starting at a tremendous disadvantage because it's taking 3 W of static power consumption before we've even begun to operate the design. That's 3 W out of a 6 W power budget, 50%. The Virtex-4 part on the other hand, worst case, is taking less than 25% of the power budget. As we begin to increase frequency, the benefit of some of the embedded blocks that we have and the lower power per block on some of our blocks comes into effect. Because the slope of the Stratix-II line is higher than the slope of the Virtex-4 line. You'll notice that the power budget in this particular case was exceeded in the Stratix-II device at a frequency substantially below the target of 200 MHz. The Virtex-4 part on the other hand at 200 MHz still has 1.1 W of power that it's not using below the 6 W. It's only at 4.9. At that same operating frequency of 200 MHz, the Stratix-II part would have been at 8.5 W, which is close to 50% over the power budget. So the Virtex-4 part in this case has given 50 MHz of headroom in frequency and about 1.1 W of power headroom. And this can be important as you're trying to get more performance into your parts. Sometimes when you can do operations in parallel with the parts running at twice the rate, you may consume more power, but if you have that available you can take advantage of that in the Virtex-4. You may not be able to in the competitor's part. So static power offset worst case 85 degrees C very important, and you need to meet your power budget. So now let's move on to a couple actual real measured design examples on the bench.

SLIDE #31

So as Anil mentioned earlier, I did these measurements myself personally. I have quite a lot of design experience and measurement experience from years of this work. We used as stimulus devices a Thermonics temperature-forcing inducer. This is the type of device that we can put

down on top of an IC. We can attach a thermocouple to the case of the IC and use that temperature as a feedback to control the exact case temperature. And as you know from looking at data on case to junction temperature, there's a very very small rise in junction temperature compared to case temperature. We also used as a stimulus source a variable-frequency signal generator. Xilinx has made actually a little box that allows you to do this, and we can sweep frequency with this sort of stimulus device. The devices under test, not unexpectedly, are the Virtex-4 LX60 in the 1148 package and the Stratix-II EP2S60 in the 1020 package. So these parts are parts of similar logic density. And as measurement equipment, we used the Tek digitizing oscilloscope, DMMs and current probes. So now let's look at the measurement of static power.

SLIDE #32
So as you notice from our descriptions before, static power increases pretty dramatically in both devices, Virtex-4 and Stratix-II, when you sweep temperature from a low of -40 up to 85 C and even beyond. Now it's important to note that the commercial operating range for these devices are 0-85 C. However, in many applications, customers choose an industrial range, which goes wider, -40-100 C. So I made this range of measurements. And the basic method is that you start at the low temperature, -40, you get it to reach a case temperature specified and you get the part to dwell at that temperature for a few minutes. So you stabilize. Basically then you just move on to the next temperature, do the same sort of measurement. As you can see, when we get to higher and higher temperatures and particularly a normal junction temperature of 85 C, the Stratix-II part has substantially more power consumption than the Virtex-4 part. This difference in power consumption, particularly if we think about -- this is typical, this is a bench measurement of one part, if we think about worst case, these numbers are higher for both of our companies, but a couple watts of difference in power consumption on your board when you may have a number of FPGAs can be very very substantial and can amount to having a very different power supply design. Particularly if you're replicating a number of boards in a system.

SLIDE #33
Now let's take a look at an example of a dynamic power measurement. That will be coming up shortly. In this particular measurement, what we wanted to look at was a 48 tap FIR filter. This particular filter has nonsymmetric coefficients that are 18 bits and it has data input that came from a linear feedback shift register. That's also 18 bits. And that's going basically into the input of the first stage of the FIR filter and of course based on what an FIR filter is it gets cascaded to subsequent multiply blocks. And this plot shows you the dynamic power with the static power subtracted off. So it's only the dynamic power. And we've swept the measurement from 0 to 250 MHz in 25 MHz steps and measured the dynamic power consumption. As you'll notice, the Virtex-4 is lower than the Stratix-II in this case as well. So this is just another measured example. Again, I made these myself personally. And I plan on making additional measurements as time goes on. Bench measurements are very important to get an idea of what's really going on. Changing data in prediction tools doesn't change the part. If I go downstairs and measure my part tomorrow, it's not going to change just because someone changed their datasheet. It's important for customers to measure real parts on the bench. And I think you'll find that in both the area of static and dynamic power, the Virtex-4 has some substantial advantages. I'm going to pass it back to Anil because he's going to tell you a few more things about -- a little bit more on design examples and some of our tools.

SLIDE #34
Thanks, Matt. As you can see, we have made extensive analysis, both at the tools level and then taken the designs into the lab to make actual measurements. And we see reduction in static, dynamic as well as inrush currents. Now in addition to the examples that Matt has shown, we have been rolling out this huge power advantage message over the last several months to customers, and customers have been asking us how can I see this advantage for my unique design. So one customer in particular in Japan, designing with SPI4.2, wanted to do a comparison and see how this power reduction benefits him. The example on this slide is showing that SPI4.2 example, the data rate at 800 Mbps, with 1 port, and on the user interface side a 128-bit interface. Again, measurements are done at 85 degrees junction temperature. And just as the customer has indicated, we see a two to one reduction. Customer has asked for data on LX25 and LX60, which we provided to them, and they obtained data for equivalent devices from our competition and told us that they got a two to one reduction. So we went ahead and wanted to check it out for ourselves, do the comparison with all the setup we had. We made the comparison and we find exactly 50% reduction as the customer has stated. So again by no means is this example going to be one that your design has. But we believe using all these function examples and the data for each of static and dynamic power, you can see for yourself, for your design, what the power reduction would be.

SLIDE #35
Let's move on to the next section and look at tools and resources we provide to help you with your power analysis.

SLIDE #36
Xilinx provides Xpower power estimation tool in ISE software. And this tool is going to help you take your entire design parameters and estimate power. One thing to keep in mind is our current Xpower tool in ISE does not support temperature variation yet. Our version 7.1i to be released shortly does provide temperature variation, so you can make estimates. We also provide a Web power estimator today that you can use to do estimates for your design. Now in addition to the power analysis and estimator tools, we also provide power management solutions both from Xilinx as well as partners, power supply recommendations, heat sink recommendations, as well as application notes, articles and white papers.

SLIDE #37
So what does this all mean? Whether it is static power you are concerned about, dynamic power, or inrush, or the total power, Virtex-4 reduces your total power compared to any 90nm FPGA. This is based on our unique innovations and technologies embedded within this FPGA. Now system architects doing designs realize every watt saved simplifies your design. Simplifies it in terms of bill of materials, in terms of thermal management, as well as reduced cap ex and op ex, capital and operational expenditures. So we would like you to take advantage of this power reduction, and you will see for yourself the advantages that it gives for your design. At this point, let me hand it off to Peter to give his closing comments and wrap it up.

SLIDE #38
Thanks, Anil. So we come to the end and I want to mention that Electronic Products magazine

called Virtex-4 the Product of the Year. We mentioned that in our last seminar. We'd like to mention it again.

SLIDE #39
And if you want to take advantage of the Virtex-4 power reduction, then you can do that. You can design with the headroom in both speed and power that was mentioned in the seminar. You can download free Web power tools from Xilinx, and you can get evaluation software for free also. And the best thing would be to contact your local Xilinx sales or FAE for additional information.

SLIDE #40
So now a short commercial for the next seminar in two weeks, on March 1st, where we will talk about power and ground distribution and package technology and how you can avoid problems on your PC board, how you can reduce the ground bounce on your PC board, and in general comments about signal integrity, which has become more and more important as the outputs got faster and faster. So we hope to see you again in two weeks from now. Now to some questions. Matt, you took some questions here. But first it goes back to the organizer.

QUESTION AND ANSWER
This is Matt Klein back again. I have several questions that have come in. I'm not taking these in any particular order. Anyway, one question that came in was why don't the graphs show multiple lines, not just typical measurements? We actually have measured many thousands of parts and we measured a smaller sample of those parts over very broad range of temperatures at individual temperature points to see how these curves change with process. Even in an hour-long presentation like this, we didn't have really the time to show the distribution of points versus process. However, there is a distribution. Some of them are in fact below the typical. Some of them are above the typical. And it's those very very rare few devices which will dictate the actual max numbers. And so yes, we didn't show those in our plot. But we weren't trying to hide them, we just didn't have time to show them. In some other presentations that you may be able to get from your FAEs we actually have shown some of the plots of multiple samples versus temperature for static power. And shown where they fall about the typical.

I'll move on to another question here. Question came in, did you mean 60 degrees ambient, that's very high. What I've seen is that with an outside air temperature surrounding your product of 30-40 degrees C, it's not unusual that the ambient in the vicinity of the part itself can be 55-60 degrees C. Even if it's 55 you can get a substantial rise in temperature between the air and the junction. And it's not very difficult to get an 85 C degree temperature. I've had cases where customers have asked, I want to operate actually with an ambient temperature of 85 C, and there have been cases where you have a part that's in a very small enclosure on the side of a building, it gets sun exposure, in Arizona, you can be at 80 degrees C and beyond even for the ambient temperature in that case. Again, it's not difficult to hit a junction temperature of 85 C.

Let's see, next question is what functions were we talking about that were measured for dynamic

power tests, and how did you compare the V-4 to the S-II? Same functions, etc. In the case of the charts that Anil showed, we compared just individual blocks. In the case of the PowerPC versus the Nios II, we took the Nios II core from the cores tools and we dropped it into the power prediction spreadsheet and compared the total to the PowerPC. In the case of smaller, simpler items like the DSP blocks, we just did them in line in both of our power tools. And as I might note that our Web power measurements on dynamic power are in fact based on measurements. The other thing is that when we actually did the bench test of the DSP filter, we basically took the megawizard to construct a portion of the FIR filter to construct 4 tap FIR, and then we put those 4 tap FIRs together with the appropriate extra adders that you need. In the case of the Xilinx Virtex-4, we made a 48 tap FIR filter, which you can do in line with the -- by basically just cascading multiple of our DSP blocks and the software automatically says oh, these get linked together, and they basically go into a column. So it doesn't require any fabric resources for the interconnect.

See, next question here. It asks are the embedded circuits built in or can they be chosen? For various of our subfamilies of parts, there are various numbers of these circuits built in. So the LX and the SX family -- well, in fact all of the families have some amount of DSP blocks, some amount of block RAM FIFO items, some amount of source synchronous I/O, in fact every I/O pin in every family has our source synchronous I/O blocks, our ChipSync technology. The FX family has the PowerPCs in every single one of the FX devices. And has Ethernet MACs. Most of the FX family additionally has MGT blocks for the multigigabit transceivers. They can't be chosen by the customer other than choosing what part you want to use, like an FX or an SX or an LX.

Let's see if we have some more questions. Getting past a few questions here. Ah. Question came in about inrush with Spartan 3. Spartan 3 doesn't have inrush either. That's been eliminated in Virtex-II, Virtex-II Pro, and Spartan 3. And Virtex-4 of course. I might want to mention another thing about Spartan 3, because this sometimes comes up with customers. They say well it looks like the Virtex-4 static power is a little higher than the Spartan 3 static power. Why is this? You know, you say you have the lowest 90nm FPGAs. Well, one thing is we're careful to state that the Virtex-4 is the lowest 90nm high-performance FPGA. The reason why that distinction is important is because the Spartan 3 doesn't have as high system speeds generally used in the device. So even though they're both 90nm between the Spartan 3 and the Virtex-4, the transistors don't need to be optimized for as fast an operation. This allows the leakage to be reduced. When we're comparing devices like the Virtex-4 and the Stratix-II, these are both considered to be high-performance FPGAs, and that's where the leakage gets you is where you want to optimize for high performance and low leakage.

There are a few more questions. Someone has asked if fans reduce the temperature below 80 degrees C, would that be less static power and less leakage? Yes, that's definitely the case. I mean if you're doing the bench measurements, if you reduce the temperature of the environment, that will reduce the temperature of the ambient, which will reduce the temperature that the heat sink or the case is seeing, which will reduce the junction temperature. Fans help some, but when you're talking about parts that have 6 W of power consumption, you have to have a pretty good heat sink. Because the heat sink, even with a good heat sink, may be a few degrees C per watt. So when you have a few degrees C, let's say you had 4 degrees C per watt, and you have 6 W of

power consumption, you're going to be raising the ambient to case by 20 degrees C. And then you have a little bit more rise, a nominal rise, between case and junction. So even if you drop the air temperature, when you have a couple watts of power consumption, that produces heat.

OK. Someone has asked another question. It says I thought dynamic power was always higher than static power. Are you saying this isn't true in 90nm? Well, the dynamic power that a given design takes is highly dependent on what the design is. Some designs may have the dynamic power representing 90% of the design. Some designs may have the dynamic power representing 20% of the design. Some may have it be 50-50. It really depends on an individual design. But one thing that's true is that when you go to 90nm technology, proper design of the 90nm device lowers static power, which is always a benefit to you if you're not compromising performance, which is the case in the Virtex-4.

That about it for questions? Or do we have any more? The entire Webcast is available, and the slides, are they available separately? They're part of the archive. TechOnLine will post this and there'll be an archive on the seminar in a few days. Because we're running out of time a little bit, we may come back and send some of you personal emails based on the questions that you've asked if we didn't have enough time to get to your questions. Any more? OK, I'm going to pass it back to who? I'm going to pass it back to the administrator, Peter. You want to take it, Peter?