

109764 - Xilinx 90nm FPGA - Matt Klein & Peter Alfke

(music stopped @ 3:04)

MODERATOR:

Hello, I would like to welcome everyone from around the world. Thank you for joining us today for this Webcast on "Achieve 1 to 5 Watts Lower Power Per FPGA," brought to you by Xilinx and TechOnLine Webcast. Your presenters today are Matt Klein, Senior Staff Engineer, Applications Engineering, and Peter Alfke, Director, Applications Engineering.

This Webcast software allows you to sit back and have the navigation advance automatically. As a user participating in the Webcast, you will be able to enter questions at any time during the presentation by clicking on "Ask-A-Question" button, typing your question in a pop-up window that appears and, then, click "Submit." Matt and Peter will be answering questions at the end of this Webcast but please enter them at any time.

Also included with this Webcast is a survey. Please take the time to open, fill out and submit the presentation survey. You can access this survey at any time in the Print Documents and View Links pull-down menu 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 this Webcast. By submitting this survey, you will be providing Xilinx and TechOnLine Webcast with valuable feedback on the subjects covered on this Webcast and also how we can improve the Webcast product.

And now it gives me great pleasure to introduce you to Peter Alfke.

PETER ALFKE:

SLIDE 1:

Good morning on the Pacific Rim here. Good afternoon in the U.S. and good evening in Europe. This is the 8th in our series of 90 Nanometer Seminars that present the Virtex-4 family. We will get into more specific details on power this time.

SLIDES 2-3:

So before we introduce the Virtex-4 family, we have talked to our customers and asked them what their challenges are and we heard about faster design time, more competition and so on, and one of the major issues was power consumption. Power

consumption in Virtex, in modern FPGAs, is very important, and so this is the second seminar in this series and it will go into more details.

So I'll do the short introduction and, then, Matt Klein will cover the majority of this presentation, and he'll give you a short review of the first seminar and go through extensive lab measurements that he has done. And, then, we will talk about optimization techniques and tools.

SLIDE 4:

Why is power so important? Well, it really isn't so much the power. It's much more the heat generated by the power that you have to get rid of; you have to cool the device, so it doesn't overheat. We specify 85°C for commercial temperature and so increasing power means more heat sinks, more fans and so on. So if we can reduce the power, we reduce the need for heat sinks, we make your whole thermal design easier.

The other aspect is cost, you know. The more power we consume, the more expensive is the power supply and when we can reduce the power, we can save money on power supply. And the third aspect is, power means that the devices run hotter, and hotter temperature always works against reliability. So by reducing power, we can increase system reliability. So there's less cooling need, there's a cheaper power supply and you get higher reliability. So there are many, many reasons to reduce power and to be concerned about power.

So Matt Klein will now go through the details of this presentation, and I hand it over to Matt.

MATT KLEIN:

SLIDE 5:

Thank you a lot, Peter. Good afternoon, good evening, good day, depending on where you are, as Peter mentioned. Today, I'm going to talk about a brief review of the previous seminar that briefly goes through some of our power advantages in Virtex-4. Then, we will go on to lab measurements and on to power optimization techniques and, briefly, into the tools and, then, see what you guys have to say.

SLIDE 6:

So, we believe that Virtex-4 can achieve 1 to 5 watts lower power per FPGA and the broad way that we come to that conclusion

is because we have up to a 73% reduction static power and up to an 86% reduction dynamic power compared to our competitor. This is enabled through a number of industry-leading technologies. One of them is our triple oxide which we introduced in our 90nm Virtex-4 FPGAs and, additionally, through extensive use of embedded IP, which we also have in Virtex-4.

SLIDE 7:

So, broadly speaking, from our previous presentations, we believe that across the difference sized FPGAs, from small to large, and considering static power at 85°C and considering core frequencies which easily reach 200MHz, we believe in designs like the one shown on the right, that we can achieve a benefit of 1 to 5 watts compared to our competitor.

SLIDE 8:

So, now, as a review of the triple oxide, again, normally, in 90nm technology compared to 130nm or larger, leakage current increases with the decrease of channel length and with gate oxide, thickness decrease. Normally, FPGAs and ASICs have two oxide thicknesses which are commonly used. A thin oxide is used in the core, the fast logic, and a thicker oxide is used in the I/O. Virtex-4 adds a third middle thickness oxide which we use for a number of areas of the FPGA that don't compromise performance but dramatically reduce the overall leakage. This is the general way that we reduce static power.

SLIDE 9:

As far as the way that affects difference-sized FPGAs, you can see that we're comparing Virtex-4 FPGAs to the Stratix II FPGAs, and, on the bottom, is the chart showing similar-size, equivalent devices based on logic cell counts. And across the board, we achieved from the high 60's to a 73% reduction in static power, and this is based on a comparison tool that 85°C, the Stratix II static power tools and the Web power tools.

SLIDE 10:

Additionally, Virtex-4 has a number of hard IP items. Hard IP items reduce both static and dynamic power, if you could take advantage of those functions, because, effectively, you're using ASIC-like processes for those hard IP. So Virtex-4 has a number of those blocks from the ChipSync Serial I/O through Ethernet MACs through BlockRAM FIFO logic, hard logic and through our XtremeDSP Slice. All of these objects reduce dynamic power and also can run at very fast speeds for high performance.

SLIDE 11:

And here is a chart coming up that shows some of these reductions. You'll notice in objects like LUT and FlipFlop, we have similar dynamic power to our competitors. This is LUTs/FlipFlops and routing -- and similar with raw I/O power. But when we get into objects like DSP and Processors and Ethernet MACs, we get a dramatic improvement in power.

SLIDE 12:

Now to move on to some of the lab measurements.

SLIDE 13:

We've constructed a board, which we will discuss, which is a competitive board that allows us to make measurements on both the Virtex-4 and the Stratix II parts. This board accommodates basically the 1148 package on the Virtex-4 and the 1020 package on the Stratix II, and there are a number of measurement points on the board for measuring not only power but for measuring signal integrity. And we basically have a method by which we can measure the power consumption for each of the power supplies feeding into the Virtex-4 or the Stratix II part, and we can provide stimulus sources in the form of controlling the temperature environment or in the form of, at a variable rate LVDS Clock going into each of the two parts.

SLIDE 14:

What we have in the voltage regulator leads of each power supply are a 5mA Kelvin resistor, very accurate, high-power resistor. So we can measure the voltage with headers that we have across each of those resistors and find the current going into each power supply. They are adjusted to be producing exactly the nominal voltage; so when we measure the current across these, we get an accurate representation. When we measure the millivolts across there, we get an accurate representation of current and, hence, power with that power supply.

SLIDE 15:

The basic calculation method is that we download desired designs to the board for either the Stratix II or the Virtex-4, or both. I might want to mention that the board halves of that large board you saw are completely isolated from each other. After downloading the designs and selecting our stimulus sources, whether they're temperature or variable-frequency sources, we then select a power supply to measure, connect the power leads with a special connector that we have across a header. And, for instance, if we were measuring the 1.2 volt supply, which is the

core supply and one that's typically of interest, we are measuring the number of millivolts across the Kelvin Resistor.

So the current from the regulator is going through that Kelvin Resistor with a very accurate and high-quality resistor, very thick leads, and then it drops into the plane after that. So, for instance, in this case, where the volt meter is showing 15.884mV, we simply divide that by the size of the resistor to get currents, the voltage of a current. Our voltage of a resistance gives us current, and, in this case, we derive 3176mA or 3.1 amps approximately, and calculate the power dissipated by that supply, we get 3176 x 1.2 volts, and that gives us about 3.8 watts or 3800mW.

So that's the calculation method.

SLIDE 16:

For design creation, we basically can create a couple of different types of designs. One simple type of design that we create for static power measurements are: We create a blank design with no I/O and no fabric resources that we can download to each of the FPGAs.

Then, later on as I'll describe the method, we can control temperature and measure basically static power consumption by measuring the V_{CCINT} power supply. We don't have really any need to measure the V_{CCAUX} supply, although we have. The V_{CCAUX} is not a leakage effect and we can show from other data that the V_{CCAUX} is processed and temperature-independent or invariant.

On dynamic power case, we would be measuring designs which contain fabric and/or I/O. The designs might also contain some of the hard IP such as FIFO and BlockRAM and/or creating the designs or creating hard IP, we would use our Xilinx CoreGenerator and we'd use the Altera MegaWizard IP Generator, where possible.

So you map the I/O and clocks through the constraints files that are matched up to the schematic and layout of the board. Then, you synthesize, place and route the design and create the bitstreams, using the Xilinx tools or the Quartus Tools, the Altera Quartus Tools. And, then, you download the appropriate designs to each half and you can measure it yourself, your power.

SLIDE 17:

So for the static power case, we basically, again, download a blank design and, in the case of static power, because it has such a high dependence on temperature, we are hooking up to a thermal forcing system and we basically can control the case temperature of the device under test. And if you're doing bench tests and you don't have thermal force, you can optionally use a heat gun or cold spray to elevate or decrease the temperature.

Anyway, now you power the system on, download that design and, through a measurement loop, we can basically measure the leakage current at the different temperatures as we adjust the temperature control system. And we basically dwell at a temperature until the case temperature stabilizes and, then, we can measure the voltage across that resistor for V_{CCINT} and divide by that resistor to get the current and calculate power.

SLIDE 18:

For dynamic power, the technique is similar except in this case our stimulus source is an LVDS Clock Generator. So you power-on the designs, connect your LVDS sources to the Xilinx and/or the Stratix or the Altera side, power on the FPGAs and download your desired designs. Then, in a measurement loop, in this case, you would be adjusting frequency, you wait for the stabilization of power, in case there are any thermal effects, and then you measure the voltage and calculate current and power accordingly, as to the method that I described earlier. You can optionally measure I/O power on different designs if you are testing that and, then, move to the next point and go again.

SLIDE 19:

This next slide just shows the basic active set-up with volt meters attached across in the upper-left and upper-right across the Kelvin Resistors for V_{CCINT} and shows the left side, the Virtex-4 side, and the right side, the Stratix II-side, and our little frequency sources there, just below the volt meters or our stimulus. So we've gone through a number of designs that I'd like to cover during this presentation.

SLIDE 20:

We have designs of static power that test equivalent density devices and we have designs for dynamic power which have a number of tests listed there-- fabric test, FIFO test, DSP and I/O tests. And we continue to do additional tests to find out information competitively as well as for customers on other functions and our parts.

SLIDE 21:

So for the static power test, again, you connect the voltage, multi-meter across the Kelvin Resistor for the 1.2 volt supply for the Virtex-4 part or for the Stratix II part. You have a thermal forcing unit which has a hood that covers the part and you sweep temperature. In this case, we've gone from 0 to 100°C and at 0, 25, 50, 75 and 100, and we've listed the measurements for milliamps and milliwatts for the LX 60 and the 2S60 -- and the same for the LX100 and the 2S90. Those are the tables below. Later, we'll cover what some of the results mean, but now we'll move on to some of the dynamic tests.

SLIDE 22:

So, as far as the test description for the fabric test, this is a test meant to test the fabric, as the design implies, which includes logic, registers and interconnects. And we want to measure the power for this type of design across a large design, so that we can get reference points for smaller types of designs for fabric and see how we compare.

So this design has a 36% toggle rate, by design of the stimulus sources, and we used ISE7.1 Service Pack 1 and Quartus 4.2 Service Pack 1. In these designs, the Verilog code is identical, the hierarchy is identical, and we've synthesized both designs.

Now, let's look at some of the utilization. So in this, the designs on the Virtex-4 side have about 21000 LUTs and about 25000 registers. The same design maps out on the Stratix II side, as you can see, with 22127 ALUTs broken up into various-sized function generators, shown there -- and the total of one extra register, 25421. These, again, were identical VHDL codes.

SLIDE 23:

What this image shows is the volt meter connected on the Virtex-4 side and on the Stratix II side for the same design, both, of course, in the LX 60 and 2S60, respectively. And the frequency source below each of the volt meters is set at 50MHz, for the first data point. So we look at the millivolts across that, divide by .005, 5 milliamps, and we calculate the number of milliamps of current and we get the number of milliwatts.

SLIDE 24:

And, for the fabric design, we will show the incrementing data points as the measurements proceeded. So this is what we've yielded at 100MHz and this, by the way, all these tests were

done at room temperature, for the sake of these experiments.
And, move on, if you would?

SLIDE 25:

Now, we're at 150MHz and, of course, you can see that the number of milliwatts is incrementing linearly.

SLIDE 26:

On to 200MHz.

SLIDE 27:

And, then, on to 300MHz.

So I will refrain here from making observations as to the results because we have several tests we're going to go through. I will describe the results of these tables when we go through the rest of the tests.

SLIDE 28:

Now we'll move on to a FIFO test. The purpose of this test is to illustrate the benefits of the Hard IP. As you know, the Xilinx has, the Xilinx Virtex-4 has hard FIFO logic built in which, for a number of communications applications means that you don't need too much extra routing and interconnect to create a FIFO.

And the design done here was 60 16K-bit FIFOs, which were configured as 4K x 4. We used fabric for the stimulus sources, counters and *** RAM the bitstream generators and for output-logic functions, in order to stop from having circuitry removed, we have all the outputs going through combinatorial logic and registers and to a final output destination point, but just one output. So there's an LVDS Clock input and one output.

The implementation method was that we used the Virtex-4 and Xilinx CoreGenerator to make our FIFO 16 hard macros, and we used the Altera Stratix II with Quartus 4.2, service pack 1, as I mentioned, and the memory MegaWizard Generator to generate its 16K FIFOs.

Then we put the identical circuitry around the 60 FIFOs -- the same stimulus, the same receiving circuitry, running in Verilog code, and that's the design I would describe. But one thing you might notice is that, in the Xilinx design for the Virtex-4, we have about 200 LUTs and our 60 18K-bit BlockRAM and about 263 registers. In the Stratix II design, because of the lack of a

hard macro for FIFOs, we have 3056 ALUTs and about 2400 registers, and that's what the MegaWizard generated across the 60 FIFOs. And, then, of course, we have 20, 240 M4K blocks.

SLIDE 29:

Now here is this measurement table of results. So as not to be exhaustive, I'm not going to go through every test at every frequency point, but this is what's being measured at 200 MHz, for this test. You can see for this one, it's about 1.1 watts of total power for the Virtex-4 LX 60 and about 2.4 watts total power for the Stratix II 2S60.

SLIDE 30:

Now, we're going to move on to a DSP test. Now the DSP is another item where we have a significant advantage in power because we have a lot of Hard IP built into the DSP block. This test, the goal was to construct an asymmetric filter with 63 taps and we used an 18-bit data input and 18-bit coefficients. The implementation method was that, in Virtex-4, we used all DSP 48 blocks in a single column. In the Stratix II, the DSP blocks allow you to have four taps in a DSP block. And for additional reconciling of the adders, for summation, we used all three-input adders in the Stratix II part, in the VHDL code. And those three input-adders were synthesized using the special three-input adders or the Altera parts.

SLIDE 31:

The control logic in stimulus-and-receiving circuitry were identical, again. Again, the goal is to test the DSP and some of the surrounding logic and so there's just the LVDS as a stimulus and one output line to eliminate removal of any of the circuitry.

So here is the result of DSP. Again, we're showing the result at 200MHz listed in this table with the other measurements that we did and, in this case, for the 63-tap filter with its stimulus-and-receive source and targets, in the Virtex-4, we see about 642mW of total power and, in the Stratix II, 2S60, we see about 1179, almost 1200mW.

SLIDES 32:

The next test that we have is a very stressful I/O test. We have 500 I/Os that are moving; their output set to LVCMOS; they're moving in DDR mode, 8mA drive strength and all output are toggling at 100%. So we're basically feeding the clock into the DDR output flip-flops in both the Stratix II and in the

Virtex-4. The amount of internal logic is insignificant for that. Basically, again, we're just using the DDR output flip-flops.

Before I move on to this test, I might mention, if we could go back, for a moment, to the DSP, actually to the DSP description. Thank you. I might mention that because we utilized the columns in the DSP48, there is very little extra logic aside -- It's, basically, exclusively the DSP48. In the case of the Stratix II, we had to utilize 187 extra ALMs and for implementing the adder reconciliation.

Now, if you'd go back to the I/O again. Sorry for the interrupt there. But, again, we have 500 LVCMOS outputs DDR, 8mA strength, all toggling at 100%. Implementation method again, were the ISE and the Quartus Tools, and identical stimulus source. Again, the stimulus source being just the LVCMOS clock going to the DDR inputs to the flip-flops and the 500 outputs. By the way, all of the output have a feminine load, 2BTT.

SLIDE 33:

Now let's look at some of the results. For this test, it's very important to, actually, look at core power, which is power mostly consumed by the clock buffers internally and clocknets, as well as power consumed by the I/O flip-flops. Additionally, $V_{CC_{AUX}}$ power is important because AUX and $V_{CC_{pd}}$ in the Stratix II part are used pre-drivers for the I/O. And so that does actually consume power. So, in this case, we use two volt meters, simultaneously, on each side -- one across the $V_{CC_{int}}$ and $V_{CC_{AUX}}$ supplies, across the Kelvin Resistor, and, on the Stratix II side, we use one across the $V_{CC_{int}}$ core and one across the $V_{CC_{pd}}$, 3.3 volt pre-driver supply, with the Stratix II.

And you can see the measurement result at 200MHz here and, additionally, the other table items. And as we'll see later, there's a significant difference in the power consumption for both of these supplies, which is a little surprising but it is, in fact, the case. And we also made I/O, $V_{CC_{IO}}$ measurements as well.

SLIDE 34:

All right. So now we're going to move on to the test-result comparison for the different tests which takes the tabular data, puts it into some graphic form. So we're going to talk about the static power results and the dynamic power results. So let's move on.

SLIDE 35:

For the static power results, we see a measured chart here, one of our pieces of measured data for the 2S90 and for the LX100. And this shows the dramatic difference in static power versus temperature between the Stratix 2S90 and the Virtex-4 LX100.

You might also note that the 2S90's data is actually quite substantially above the prediction. It's about 50% above the prediction at higher temperatures. Now, in previous presentations by Altera, they mentioned that while we hand-selected a part for the Virtex-4 and hand-selected a part for the Stratix II, we buy these parts in very small quantity. We have no ability to really hand-select them. They're actually quite expensive parts in the case of the Stratix II, to us. These parts -- While it was mentioned in the Altera presentation that parts can be above the typical or below the typical, this is true. However, the distribution of data is typically clustered very close to the prediction; not 50% above the prediction.

In the case of our Virtex-4 part, admittedly, it is below the prediction -- not very far below the prediction -- but it goes to the argument that it's clustered [above/about] the prediction. Anyway, at 85°C, this is a difference of 2.2 watts, and this is an important difference. First of all, the 2S90 is a much smaller device than the LX100. A more appropriate comparison would have been the 2S130 and the LX100. However, we couldn't get a hold of a 2S130. But we did use a production 2S90. 2.2 watts for a device this size represents a significant portion of most people's power budgets.

A power budget for a device of a size of an LX100 might be 5-6 watts. So when we're talking about a 2.2 watt-difference, that represents 40-50% of the entire budget. So, again, because of our triple oxide, we have a dramatic difference in the amount of static power that we consume. And this is true across all of our devices, in comparison to the Stratix II devices. And as I mentioned in our previous presentations, we've actually made a dramatic reduction in static power, even compared to our own Virtex-2 PRO devices through the triple oxide.

SLIDE 36:

Now let's move on a little bit to some of the dynamic comparisons. Now as we stated in our previous presentations and as we've stated again, fabric power, the raw fabric power -- at

least, the dynamic component -- is very similar for both the Stratix II and the Virtex-4 devices for similar designs.

This is shown in the graph in that the slopes are nearly identical. The offset is due to the static power difference at room temperature, which is not a small room temperature. Of course, we've illustrated previously that it's important to look at elevated junction temperature for representative cases.

Despite the fact that Stratix II uses ALMs and low-K which are touted as measures of reducing dynamic power, we see a very similar dynamic power. And, in fact, in this particular measurement, the slope of the Virtex-4 device is slightly lower than that of the Stratix II device.

SLIDE 37:

Now let's move on to some of the Hard IP items. So in the FIFO BlockRAM texts, if we look at the slopes here, we have greater than a 2 to 1 advantage in dynamic power and, of course, the static-power differences are evident as well. In this case, the Hard IP FIFOs, as well as the fact that we have 18K-bit BlockRAM, makes the Virtex-4 much more efficient for dynamic power. And you can see, at something like 200MHz, in designs which rely on a lot of FIFOs or a lot of BlockRAM, this difference is over a watt, if you're doing very high utilization.

So this can be a significant part of a power budget. And what we're looking at when we've talked to customers, what we're seeing is that power budgets for devices the size of the LX 60 are in the order of 3-4 watts. So a difference of 1 watt is significant.

SLIDE 38:

Now let's look a little bit at the DSP. Again, in the case of the DSP, the ultra-cascadable DSP48 blocks have given us a big advantage because we don't use any fabric or interconnect or any external adder structures like the Stratix II needs to use and so, again, this is another power-efficient item that saves you logic, routing and interconnect and gives you the benefit of lower power and potentially much higher performance.

Again, we see a difference in slopes of greater than 2 to 1, with the Stratix II having a much higher dynamic power slope. And, again, if it wasn't obvious, these measurements are the measurements that came from the tests shown earlier.

SLIDE 39:

Now let's look a little bit at the I/O power test results. So we did measure the different V_{CCIO} supplies and we can see that, for lower frequencies, both the Virtex-4 and the Stratix II parts have similar I/O consumption. What we're seeing here is different, of the V_{CCIO} supplies, and, at higher frequencies, the Stratix II I/O power actually rises for reasons which I can't entirely explain. Similar results were actually shown. If you look at the data points, carefully, similar results were shown for the Stratix II in their Power TechOnLine. The general idea, though, is roughly they have similar power in the raw I/O supplies.

SLIDE 40:

But where we get a more dramatic result is actually in the next picture, which shows the core and the pre-driver supply for Stratix II and the AUX supply for Virtex-4. Now with 500 I/Os toggling at DDR, we look at points such as at the 200MHz point - - you have to go to the bottom graph to see where the 200MHz is -- but if you go up on to the V_{CCINT} , there is actually a difference of close to a watt of power at 200MHz. This is just power that's being used for the clock [tree] internally and to drive the DDR output flip-flops. For 500 I/Os at this very high, at the 200MHz with the high-toggle rate, that's a big difference.

But another dramatic difference occurs in the V_{CCPD} compared to the AUX-supply. If we look at that, the Stratix II parts' being in the blue curve as they were in all of the other charts, the Stratix II parts, at 200MHz, are consuming more than two watts different power, just for the V_{CCPD} supply. And, again, this is something that contributes dramatically to power budget and being able to meet, or not meet, the power budget.

You'll notice that while the V_{CCPD} supply for the Stratix II does intersect at zero, the AUX-supply intersects probably at a little over 150-200mW. So this is a dramatic difference and, in the operating rate where you're using the I/O and you're moving at high speed, this looks like, with Stratix II, it would consume quite a bit of power.

SLIDE 41:

Now we're going to move on to some of the power optimization, and what we're trying to look at there.

SLIDE 42:

So basically, we have two areas where we can potentially reduce power: One is in static power due to adjustments to the operating environment and one is in dynamic power due to some adjustments in the operating environment -- we'll see later what those are -- and optimizations to code in the design.

SLIDE 43:

So let's move on to the static power. We're showing a chart here that shows some of the items that vary, and I'll mention at the start here, there's an error in the formula for the static-power variation with voltage, which I'll explain momentarily.

Anyway, as it has been made very clear, we have process variation that can be dramatic for static power. The voltage variation actually will vary static-power consumption in a linear fashion. That equation is wrong but we'll see in another picture how it varies. And, of course, due to leakage, we have a big variation in ICC INTQ which affect static power in the same way, exponentially.

SLIDE 44:

Just briefly, here is a chart that shows typical LX 60 devices and this is across a large number of dies. If we look at the center there and we see that typical wafers have about a +/-30% variation. However, if we get wafers that are significantly skewed from the process, those give us a much higher variation which needn't be considered when looking at worst case. And we'll also notice that devices which have higher static-power consumption also are faster devices. So we have a distribution from fast-to-slow that also increases power or is proportional in that way.

SLIDE 45:

Another very interesting thing which needs to be considered is that variation in VCC from nominal, either below-nominal or above-nominal, does cause a dramatic increase in static power and leakage current. So if you increase above the 1.2-volt nominal, you'll see in the second graph, on the right, that at 1.26 volts -- which, in our cases, are maximum guaranteed functional spec -- we have about a 20% increase in static power. And by the time you get to 1.3 volts, you have about a 30% increase in static power, from 1.2 volts.

All FPGAs have this variation in leakage current and, hence, static power, with V_{CCINT}, so for items which are due to static

power or due to leakage, we need to consider this. And the V_{CCINT} -core supply is the one which is influenced by leakage.

When we quote our worst-case numbers in our tables and data sheets and in other items where we quote worst case, where we've always quoted for items like Virtex-2 PRO and Virtex-2, we always quote at worst-case voltage, which is actually above our functional worst case. So we actually quote at 1.3 volts. Our competitors quote their worst case at 1.2 volts and, as you can see from the graph, there's a dramatic difference in that.

SLIDE 46:

Again, as we've covered in previous presentations, leakage varies dramatically with temperature. We're waiting for the slide to come up for all of you. But, again, going to 85°C, you have a 2.5x increase just because of an increase in junction temperature on leakage and, hence, static power. At 100°C, industrial-grade parts, you have a 3 to 1 increase. So, again, our tools allow you to predict static power with temperature, and it's important to look at that.

SLIDE 47:

So just to summarize what degrees of freedom you might have with static power, if you can keep the junction temperature as low as possible, either through heat sinks or other methods or controlling the environment, that helps a lot. Use the smallest parts you can, if you have that freedom. So smaller parts will have lower leakage. For instance, an LX 60 has about a 40% lower leakage than an LX 100.

Also, keep your V_{CCINT} close to nominal. The V_{CCINT} , for instance, at 1.26 volts, which is the high-end of our spec, is 20% higher, then, hence, higher static power, 20% higher static power than at 1.2 volts. So try to keep it close to nominal, if you can -- or even slightly below.

And, again, keep in mind the variations of static power with worst-case process. If you only have one part on your board, you certainly need to consider worst-case power. If you have a number of devices on your board, which are all the same device, you may not need to consider that all of them would simultaneously get worst case. It would, statistically, be much closer to typical, generally speaking.

SLIDE 48:

With dynamic power, we need to consider internally the number of

nodes, the capacitance, voltage swing and frequency. And moving on to the next table --

SLIDE 49:

-- we sort of see what kinds of things affect dynamic power. So, for dynamic power, there's a very minor variation with process and temperature. However, dynamic power varies as the square of the ratio of the voltage to nominal 1.2. And so we do need to consider dynamic power voltage with respect to that.

SLIDE 50:

So, again, the number of node-switching into capacitive load is influenced by the number of levels of logic you have. So if you can try to pack the logic tightly, that's better than not. You can use what we have, which are called "RPMs" -- relationally-placed macros, we've had them around for a number of years. Those can allow you to do tighter packing of the design, when possible. You can also use our plan-ahead tools which can help you pack a design much closer and give you a benefit in both performance and potentially power as well.

Additionally, and a really important method is, if you have clocks driving loads and your design can take advantage of the BUFGMUX, you can use this to turn off groups of flip-flops, and that will reduce clock power to your target flip-flops and on that clock net. This is a common technique that's used in ASIC design because, typically, they're trying to isolate sections when they're not using them -- and you can, as well.

SLIDE 51:

Some other methods that you can do in order to reduce amount of routing because routing and going through interconnect hops gives you more capacitance and more power consumption.

Another thing you can do is you can bump up the performance target in the Xilinx Tools. So with our XST Router, you can bump up the performance target by a little bit. This can sometimes have an improvement of 5-10% on power consumption, if you can minimize the lengths of paths. And, additionally, using the RPMs will give you not only a tighter placement but will give you better routing and, hence, lower power.

SLIDE 52:

This just shows the variation in dynamic power with V_{CCINT} . As we expect, it varies with the square of the voltage relative to nominal, and, on the right side, the prediction lines up well

with the measured. It's difficult to see those two lines because they're basically on top of each other but, again, as a designer, you all need to look at what the ratio between where what you're going to see on the board is and nominal because you will get a variation. So when you get dynamic power results from our tools, you have to look at how far above normal you are in order to get where the worst case for dynamic power is going to be.

SLIDE 53:

So, again, as a re-cap, if you can, tighten up the VCC to run at the center, the nominal voltage or slightly below because running there will keep your dynamic power down. And if you can run a little below nominal, you'll keep it down as well.

SLIDE 54:

Frequency Reduction. If you can run non-critical parts of your circuitry at lower speed rather than taking the one arbitrarily-high clock you have, you can also dramatically reduce power in those non-critical sections. If we move on --

SLIDE 55:

If you can take advantage of the Hard IP, this is one of the largest areas where you can get dynamic-power reduction because some of these Hard IP objects will give you between a 5-20 to 1 reduction over traditional FPGA fabric and traditional programmable logic.

SLIDE 56-58:

So, now if we move on to some of the Virtex-4 advantages again.

We have actually demo kits available that allow you to make these power measurements yourself or have FAEs come in and demonstrate these. We have a transit case, which contains basically our ML481 board that has the Stratix II and the Virtex-4, 2S60 and LX 60s on them. We have test-and-measurement connectors that allow you to measure both power consumption as well as signal integrity. We have LVDS sources available at our rocket labs. The kit shows some of the other equipment -- multi-meters, etc. -- and we certainly encourage you to come and look at this kit in our rocket labs or have someone visit.

We have a number of locations around the world that have these demo kits or will have them very soon, and you should contact your Xilinx salesperson or FAEs and look at what we have.

SLIDE 59:

Again, back to the tools. And we've talked about this in the past. We have some very good and innovative tools: XPower which is part of ISE for accurate power prediction which is integrated into the design flow. And we have the Web Power Tools for early power estimation to give you predictions of where you might be. And we have Power Management Solutions. We have a lot of application notes and articles and we continue to produce those and to do more work on power because it's something that's very important.

SLIDE 60:

So, again, to wrap up, if you could take advantage of a number of the benefits that the Virtex-4 devices have over our competitor, you can save 1-5 watts per FPGA compared to the Stratix II device, and our measurements show these positive results. And, again, we encourage you to try it and look at what we have.

SLIDE 61:

So just to mention what the next presentation is before we move on to the Q & A Section, the next presentation will be on May 17th at 11:00 a.m. Pacific time. It's on the MicroBlaze 32-bit soft processor cores and how Xilinx helps you achieve performance in those embedded systems.

Now let's look at some of the questions, and it looks like we have a few minutes to answer some.

SLIDES 62-65:

Reference Material.

Q & A

MODERATOR:

Thank you, Matt. At this time we'll move into the questions-and-answer section of this presentation. If you have a question for Matt, please submit it now by clicking on the Ask-A-Question button, typing your question in the pop-up window that appears and, then, click "Submit." Please take the time to open, fill out and submit the presentation survey. You can access this survey at any time in the Print Document 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 this Webcast.

Now let's go back to Matt for the Q & A.

MATT KLEIN

OK. This is Matt. I'm back again.

Q: The very first question that came in is one that may have been a little bit confusing during the presentation. It says: "What was the reason for using 2.5 volts on Xilinx V_{CCIO} while using 3.3 volts -- or Altera V_{CCIO} ?"

A: In fact, we didn't use 2.5 and 3.3 for the I/O. We used 2.5 volts for our V_{CCAUX} supply which is the recommended voltage for that and we used 3.3 volts for the V_{CCPD} supply, which is their pre-driver supply in the Altera Stratix II. For the I/O, we used the same I/O voltages and we used 1.5 volts, 2.5 volts and 3.3 volts.

Q: Let's see. The question asks: "Does Altera's Web Power Tool accurately report V_{CCINT} and V_{CCPD} for DDR I/O example?"

A: It looks like it is reporting that there is significant V_{CCPD} supply. I don't know how accurate it is, but it is reporting that there is V_{CCPD} used when driving I/O.

Q: I'm trying to look at the next question. Someone has asked: "Which of the two devices needed the most cooling?"

A: In our observance, the Stratix II devices feel much hotter even when they're consuming similar power in designs.

Q: I'm reading through these questions as I'm processing what I'm going to say. This questions says: "Your V-4 FPGA seems to be better than Stratix II in power. What about hard copy?"

A: I presume they mean "Hard Copy II." One interesting thing to note is that because our static power is so significantly better than the Stratix II regular devices, we believe that, even in the case of Hard Copy which will reduce static power, that the static component will be similar. Now if we look at some of the dynamic elements, of course, the Hard Copy II is going to have an advantage on routing and interconnect and more dynamic power. However, if you can take advantage of a number of the Hard IP items that we have, those are going to move over from Stratix II to Hard Copy, without much power reduction. So if you can take advantage of those Hard IP items, which we give you in our Virtex-4 parts, you'll still get a pretty strong power message there. The net-power for a Hard Copy II device may be slightly lower but the benefit may not be large enough to justify using it. And, in many cases, it may be appropriate to use the Virtex-4 device and you could consider using Easy Path on Virtex-4, if it's simply a matter of price.

- Q: It says: "Are either of the devices used engineering samples?"
- A: In the case of the LX 60 and 2S60 cases, both of those are engineering samples. In the case of the 2S90 and LX 100, the 2S90 is a production device; the LX 100 is an engineering sample.
- Q: It asks if we've performed -- one of the questions here -- similar tests between Virtex-4 and Virtex-2 PRO.
- A: We have done some similar measurements, and what we're finding is, basically, what we've predicted in some of our previous seminars and presentations. We find that in the static-power case, we're about 40-50% lower in Virtex-4 than Stratix II and, in the dynamic-power case, again because of the reduction in internal voltage and reduction capacitance, we're also about 50% lower than Virtex-2 PRO. What this means is that many designs which have used Virtex-2 PRO in the past, when they go to Virtex-4 will get an automatic benefit. And if you can use some additional Hard IP, you'll even get a tremendous benefit above and beyond the 50%.
- Q: This asks: "Can I take my design to rocket labs and measure power?"
- A: In the case of these particular boards that we have, we do have the test connectors for measuring power on the different voltages. The board may not be able to allow you to put in the stimulus sources and get out results in exactly the same way, from a physical point of view, that your own real design would have. But if you're willing to test the internal-core portions in a design that you've altered a little bit, then you could potentially do that or borrow a board for a certain amount of time. So you could either do it in rocket labs or through the other method.
- Q: It asks: "Why are 3-input adders needed for Stratix II? Couldn't you use DSP blocks?"
- A: The problem is that in an FIR filter, we need to reconcile the adder terms in the output of each section of 4-tap DSPs. So the Stratix II does a fine job for 4-taps' worth in compensating for the delay, of delaying the data as well as doing the multiply and the add. But outside of that, we use the 3-input adders and we believe that -- although I don't know for certain -- we believe that the filter tools from Altera also use 3-input adders when they can. So you still would have interconnect routing going to the DSPs from the output of each block, even if you chose to use a DSP block as the adder. And that routing power does take up some power and, then, you would, you would, then, be

comparing basically the DSP block to the adder. And the adder is also a high-speed device, so I don't know what the comparison would be there. But we used adders not to try to cheat things. We used adders because they seemed like the most logical source -- to use those 3-input adders.

Q: Let's see. It says: "Are your comparisons between Stratix II and Virtex-4 worst-case or typical silicon?"

A: As far as we know, they're typical silicon in both cases. An interesting thing about "worst case" that it's worth noting here is that in order to get a worst-case part, you may have to go through 5-10,000 parts, and we don't have access to 5-10,000 parts at the moment for this kind of testing with either Virtex-4 or Stratix II. The way that we yield our worst-case statistical data is by looking through a large number of parts, but we don't do all of our designs in those large numbers of parts. We look through a large number of parts at leakage current because that's something that's tested during wafer sort and during final test and, therefore, we have more statistics on that. But, in some cases, in earlier stages of production, we actually, purposely, do die skews on parts in order to push them toward worst case. So we do die skews where we shorten gate lengths or where we alter doping and then we can see what would come up as worst case, from a statistical point of view.

Q: Let's see. It says: "When will XPower support Virtex-4?"

A: XPower currently supports Virtex-4. It began being supported in the ISE 7.1 Service Pack 1 Tools.

Q: Let's see what else we have? "Were timing constraints used on either of the implementations?"

A: Generally speaking, we try to target the designs for about 200-250 MHz, so we used a global constraint.

Q: Let's see what else we have? "Can I get a copy of the slides somewhere?"

A: The entire presentation will be put up on the TechOnline, as well as you can have access to it from Xilinx, once it gets processed through the system.

And I think that's about all we have time for questions now. We're getting very close to twelve o'clock. I thank you all for the great questions you had and for listening to the presentation. And come see others in the future. And, again, the next one being the Microblaze 32-bit soft processors. Thank you very much.

MODERATOR:

I would like to thank everyone for attending today's presentation of "Achieve 1-5 Watts Lower Power Per FPGA," brought to you today by Xilinx and TechOnLine Webcast.

I would like to remind you to please fill out and submit the survey. This survey will open when you choose to close your viewer window or when the viewer window closes automatically at the end of this Webcast. By submitting this survey, you will be providing Xilinx and TechOnLine with valuable feedback on the subject covered in this Webcast and also how we can improve the Webcast product.

This presentation will be available to all registered users in an on-demand format. You will receive an e-mail with information on how you can access the on-demand version of this Webcast.

Thank you, again, for attending. We hope to have you join us for future on-line Webcasts. For a current schedule of live and on-demand events, please go to www.TechOnLine.com for a complete listing. Thank you, and have a good day.

End of 109764 - Xilinx 90nm FPGA - Matt Klein & Peter Alfke