



WP221 (v1.0) March 23, 2005

Static Power and the Importance of Realistic Junction Temperature Analysis

By: Matt Klein

Total power consumption of a board or system is important; each FPGA or ASIC in a system is beginning to be forced to meet a power budget. With this concern and the trend of increasing static power with use of high performance 90 nm FPGAs, Xilinx has put considerable effort into reducing static power in the Virtex™-4 FPGAs. To this end, it is important to consider a realistic operating temperature for the FPGAs, which can easily have junction temperature up to and in excess of 85°C. As junction temperature rises, static power rises exponentially, fueling this concern.

Introduction

Customers expect FPGA and ASIC vendors to reduce cost and increase performance. Typically, this was achieved in the past by reducing transistor sizes, hence increasing the performance and decreasing the die area and cost.

Reducing transistor sizes increases transistor leakage current, and therefore static power. With 90 nm technology in FPGAs or ASICs, there is a particular challenge to reduce static power. A further problem is that leakage rises dramatically with junction temperature. This white paper describes where static power comes from and its variation with temperature, providing insight into how Virtex-4 FPGAs address decreasing static power. Xilinx leverages its vast experience as the industry leader, producing millions of 90 nm FPGAs, to research this issue and consequently reduce power in 90 nm Virtex-4 FPGAs.

Even though system speeds are increasing, core voltage is dropping, which reduces the rate of increase of dynamic power; however, static power is growing exponentially over time as we move to smaller and smaller technology nodes because of increasing transistor leakage. **Figure 1** from the International Technology Roadmap for Semiconductors (ITRS) shows a cross-over point as the industry arrives at 90 nm and smaller technology nodes, where static power is beginning to eclipse dynamic power for many applications.




Figure 1: **Static and Dynamic Power vs. Technology Node**

Decreased Power Lessens Other System Design Issues

FPGAs are being used increasingly in many applications, so reducing power consumption in FPGAs provides huge benefits to the system design. Some of the key benefits are shown below:

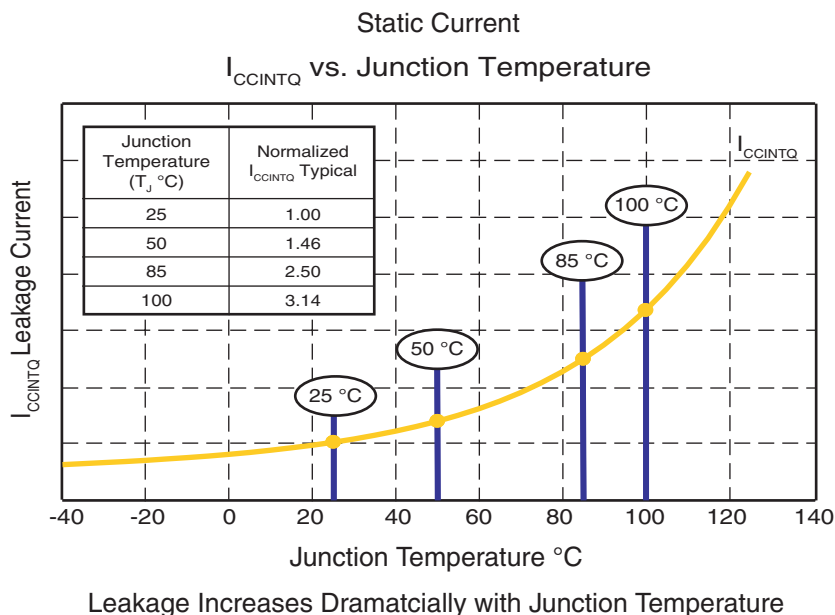
- Fewer thermal concerns - lower power causes smaller rise in junction temperature, preventing thermal runaway:
 - ◆ Use smaller heat sinks or no heat sinks at all.
 - ◆ Allow for a simpler system thermal design requiring less airflow and smaller or fewer fans.
- Lower cost power system - reduced power requires smaller power supplies:
 - ◆ Power supplies cost from \$0.50-\$1.00/Watt.
 - ◆ Additionally, there are some *total power* system considerations where it is important to stay below a step function jump, such as < 500 W, < 1 KW, etc. If the design is at the border, every watt counts.
 - ◆ A smaller and simpler power circuitry means:
 - Fewer components.
 - Smaller PCB.
- Higher system reliability - running with reduced power lowers the junction temperature, and lower junction temperature improves reliability.

Since total power consumption of a board or system is important, each FPGA or ASIC in a system must meet a power budget. Users need to design within the power budget while being pushed to higher performance and lower cost - a rather difficult challenge at 90 nm. With this in mind, Xilinx has put considerable research and development into reducing static power in the Virtex-4 FPGAs and actually reversed the trend for high static power in a high performance 90 nm FPGA.

Where Does Leakage and Static Power Come From?

A well-known rule of semiconductor physics is that when transistor length decreases, leakage current increases. Smaller physical distances make it easier for current to leak. Both source-to-drain leakage and gate leakage are inversely proportional to channel length and gate oxide thickness, respectively, and show a dramatic increase in leakage. Static power is the power consumed due to leakage in the transistors even when the transistors are not switching. Therefore, power is consumed in the FPGA even when it isn't performing a task.

The leakage is dramatically influenced by junction temperature. This is why it is important to look at realistic junction temperatures for the FPGAs in the system to properly account for static power. Figure 2 shows the dramatic increase in I_{CCINTQ} , the leakage or static current for the core transistors in the FPGA.



wp221_02_032305

Figure 2: Leakage Current vs. Temperature for 90 nm Virtex-4 Devices

Two components contribute to leakage:

- Source-to-Drain Leakage:** This current flows from source to drain of the transistor, even when the transistor gate is off. When the transistor gets small, it is harder to prevent this current from flowing, and therefore, 90 nm transistors tend to exhibit source-to-drainage leakage with much greater magnitude than larger transistors, all parameters being equal. The other problem is the thickness of the gate oxide. A thinner oxide allows the transistor to be switched on and off faster, but it also increases leakage. The amount of leakage is also influenced by the threshold voltage of the transistor. The threshold voltage, or V_T of the transistor, is the voltage between the gate and the source at which the channel conducts current to a certain extent. Small high-speed transistors need a lower threshold voltage (influenced by oxide thickness and doping) to maintain the speed with which the transistor can be turned on and off via gate control, but this also tends to increase the leakage because the channel of the transistor can not be turned off completely. The goal is to make fast transistors, but unfortunately, physics is against this. Another important note is that source-to-drain leakage increases exponentially with increasing temperature; in going from a junction temperature (T_J) of 25°C to 85°C, it goes up by a factor of 5x (see IS →D, Figure 3, page 5).

- Gate Leakage: This current flows from gate to substrate. This component of leakage is now more substantial as transistor gate oxide thickness has decreased at the 90 nm technology node. At room temperature, the leakage from gate to substrate is larger than the source-to-drain leakage in the 90 nm fast thin oxide transistors; however, unlike source-to-drain leakage, this gate leakage only increases marginally with increased temperature (see I_{GATE} , Figure 3).




Figure 3: Leakage Current Components versus Temperature for the 90 nm Virtex-4 FPGA

The net effect is that total leakage $I_{CCINTQ} = I_{S \rightarrow D} + I_{GATE}$ goes up by 2x - 3x between 25°C and 85°C.

FPGAs and ASICs are getting larger and larger due to requirements of hardware and system designers. This means that with parts the size of the new Virtex-4 LX200, there are upwards of one billion transistors to worry about. A leakage of 10 nA per transistor amounts to a leakage of 10 A!!


There is a silver lining for FPGA manufacturers that is not applicable to ASIC manufacturers: Xilinx has solved the issue of high static power in the Virtex-4 device, a high performance 90 nm FPGA, by using triple-oxide.

What is Triple-Oxide and How Does It Affect Static Power

Triple-oxide simply means that we use a third thickness of gate oxide in making some of the transistors in the FPGA (two oxide thicknesses are used in devices like the Virtex-II Pro FPGA). Most transistors in the past had a thin oxide layer, some with NMOS or PMOS transistors that operate at either low V_T or regular V_T . Thick-oxide transistors are mostly used for I/O drivers.

Oxide deposition thickness is a very stable and controllable process in the semiconductor industry because it depends on temperature, concentration, and exposure time. **Figure 4** shows the Virtex-4 transistor with the middle oxide thickness used in the triple-oxide process. The oxide thickness is still very, very thin, but this thicker oxide transistor has much lower leakage than the standard thin-oxide low V_T and regular V_T transistors used in Virtex-II Pro FPGAs and in various parts of Virtex-4 FPGAs.

Industry's First 90 nm Triple-Oxide Technology



Only Xilinx FPGAs Benefit From This Technology

wp221_04_032305

Figure 4: Middle Oxide Thickness Virtex-4 Transistor Used in Triple-Oxide Process and with Highlighted Portions of the Transistors

Why Doesn't Everyone Use Triple-Oxide?

If triple-oxide is such a great process, why don't other companies like Intel or IBM use it in their own ASICs? They probably would if it benefited them. The reason they don't is that all of their transistors need to run at speed; hence, they must use the low V_T , leakier transistors for everything.

Unique to SRAM-based FPGAs and in contrast to ASICs, all transistors don't need to be the smallest, fastest, high-leakage transistors. FPGAs use a large number of transistors (up to 1/3 of those in the FPGA) to hold the design-specific settings for the interconnect transistors and the logic functions. These are the configuration memory cells. Additionally, the interconnect pass transistors, while they need to be fast at passing a signal from source to drain, don't need to be switched on and off rapidly (controlled by changing the gate voltage). The connections in the FPGA are known in advance. A transistor that doesn't need to have its gate switched rapidly doesn't need to be a thin oxide leaky transistor at all, low V_T or otherwise.

Triple-Oxide and What It Does to Reduce Static Power

Xilinx already has had experience with 90 nm FPGAs for a few years through the deployment of the Spartan™-3 FPGAs. Consequently, a few important things were learned. The Spartan-3 is a medium performance FPGA, primarily optimized for the lowest possible cost. In contrast, the Virtex-4 FPGAs are primarily optimized for very high performance along with decent cost reduction. When using 90 nm transistors of the highest performance, the leakage described above becomes very large; so what wasn't a problem in the Spartan-3 FPGAs would have been a problem in Virtex-4 devices had it not been for the early realization of this leakage issue by the Xilinx IC design team.

Triple-oxide means that in addition to thin-oxide, small, fast, transistors for the FPGA core and thick-oxide, higher voltage, swing transistors for the I/O, we have introduced a third middle-thickness transistor, which can be used for configuration memory cells and interconnect pass transistors. These transistors operate at a higher threshold voltage than the thin oxide transistors.

So, the bottom line is that if a design uses hundreds of millions of transistors in a high-performance 90 nm FPGA, the transistors can be lower leakage middle-thickness oxide, which greatly reduces leakage without compromising performance in the FPGA. The net effect is that the total FPGA leakage is substantially lower than that of a similar-sized FPGA that doesn't use triple-oxide.

If a high-performance FPGA is made without using triple-oxide to substantially lower the total leakage, not much can be done after the fact; as mentioned earlier, raising threshold voltage via doping changes or increasing oxide thickness can't reduce leakage affecting performance.

Xilinx made smart choices at the beginning of the design phase. In addition to triple-oxide, Xilinx IC designers balanced gate lengths and gate widths of transistors, voltage thresholds, and circuit design to create a high-performance 90 nm FPGA with low leakage.

Figure 5 shows leakage amounts for the the Virtex-4 device compared to the 130 nm Virtex-II Pro device and other 90 nm high performance FPGAs at 85°C. This data is based on prediction tools and data sheets from Xilinx and prediction tools from the competing 90 nm FPGA. The Virtex-4 device has greater than 60-70 % less leakage than its competitor, which translates to 60-70 % less static power.

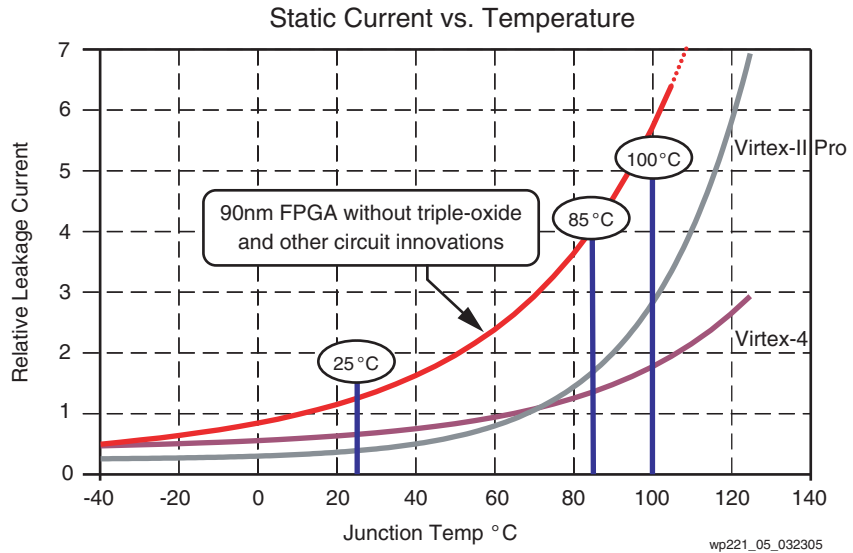


Figure 5: Relative Leakage Current in the Virtex-4 Device versus the Competitor for Similar Logic Density

Xilinx has measured thousands of devices; the overall leakage, and hence static power, is low. Figure 6 shows measurements of actual parts from Xilinx and Altera, specifically the Virtex-4 XCV4LX100 and the Stratix II EP2S90.

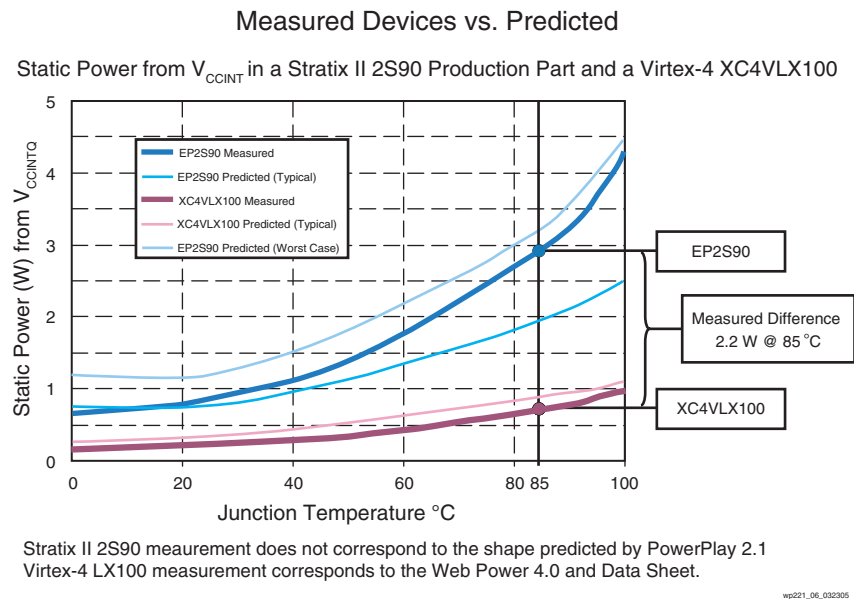


Figure 6: Measured versus Predicted Static Power from V_{CCINT} in Stratix II EP2S90 and Virtex-4 XC4VLX100

The measured power savings of 2.2 W for the XC4VLX100 versus the EP2S90 at 85°C is significant, especially considering that the XC4VLX100 is a larger FPGA than the EP2S90. Also note that the measured data for the Stratix II EP2S90 compared to the

worst case and typical prediction data from Altera's PowerPlay Early Power Estimator for Stratix II, version 2.1 is much closer to the worst case prediction at 85°C junction temperatures and beyond. On the other hand, the measured data for the XC4VLX100 is actually lower than the typical predicted static power versus junction temperature from the Xilinx Virtex-4 device (see the Virtex-4 data sheet <http://www.xilinx.com/bvdocs/publications/ds302.pdf> and the Web Power Tools versions 4.0 and 4.1).

FPGA Operating Environment and Static Power


Operating environment for the FPGA must be considered carefully when looking at Static Power consumption in the high performance 90 nm FPGAs. Designers look at the normal ambient environment for various equipment that uses FPGAs, but it is the T_J of the FPGA die that is most important when looking at static power. As in [Figure 2](#), there is 2.5 times more leakage at a junction temperature of 85°C than there is for a junction temperature near 25°C, and there is greater than 3 times more leakage at 100°C than at 25 °C.

In [Table 1](#), some of the typical junction temperatures might be from greater than 75°C to beyond 100°C, depending on the type of product.

Table 1: Environmental Conditions for Various Types of Equipment

	Operating Temperature Environment		
	Ambient Outside Equipment T_{AO} (°C)	Ambient Inside Equipment T_{AI} (°C)	Device Junction Temperature T_J (°C)
Type of Equipment	High	High	High
Computing, Storage, Server Racks	45	55	75 - 85
Metro and Access Networking	85	100	> 100
Automotive, Aerospace and Defense	100	> 100	> 100

Typically, designers using FPGAs have power budgets ranging from < 1 W for a small device like a Virtex-4 LX15, to 4-6 W in an LX60, to 10 W in an XC4VLX100 and so forth. Figure 7 shows the static power consumed by the Virtex-4 devices and Stratix II devices at 85°C based on Altera's PowerPlay Early Power Estimator for Stratix II version 2.1 and the Xilinx Web Power Tools versions 4.0 and 4.1. As mentioned earlier, the Xilinx tools and data sheets seem to be more conservative than the competitor's tools. Nonetheless, a large static power savings is gained in the Virtex-4 FPGAs compared to the Stratix II FPGAs across comparable part densities from the two companies.



Note: Based on Altera PowerPlay Predictor v2.1 and Xilinx Web Power 4.0 and Data Sheets.

wp221_07_032305

Figure 7: Static Power at 85°C in Virtex-4 and Stratix II Equivalent Density FPGAs

Conclusion

When considering thermal operating environments for various industries, many consider junction temperatures from 85°C to 100°C. Even at these junction temperatures, the Xilinx Virtex-4 FPGAs present an acceptable amount of static power, which leaves plenty of design margin for the dynamic components of power consumption in the FPGAs. The closest 90 nm competitive FPGA has 2-3 times the worst case static power consumption and up to 5 times the typical static power consumption. Virtex-4 devices give you the design edge to meet your performance without breaking your power budgets.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
3/23/05	1.0	Initial Xilinx release.