



WP246 (v1.2) February 1, 2007

Power Consumption in 65 nm FPGAs

By: Derek Curd

With the introduction of the Virtex™-5 family, Xilinx is once again leading the charge to deliver new technologies and capabilities to FPGA consumers. The move to 65 nm FPGAs promises to deliver many of the benefits traditionally associated with smaller process geometries: lower cost, higher performance, and greater logic capacity. However, along with these benefits, the 65 nm process node brings with it new challenges. This white paper addresses one of those challenges, power consumption in 65 nm FPGAs. As with the Virtex-4 family, Xilinx has implemented a number of process and architectural innovations in Virtex-5 devices to ensure that static power consumption is minimized and that the dynamic power benefits of moving to a new process node are fully realized.

© 2006–2007 Xilinx, Inc. All rights reserved. All Xilinx trademarks, registered trademarks, patents, and further disclaimers are as listed at <http://www.xilinx.com/legal.htm>. All other trademarks and registered trademarks are the property of their respective owners. All specifications are subject to change without notice.

NOTICE OF DISCLAIMER: Xilinx is providing this design, code, or information "as is." By providing the design, code, or information as one possible implementation of this feature, application, or standard, Xilinx makes no representation that this implementation is free from any claims of infringement. You are responsible for obtaining any rights you may require for your implementation. Xilinx expressly disclaims any warranty whatsoever with respect to the adequacy of the implementation, including but not limited to any warranties or representations that this implementation is free from claims of infringement and any implied warranties of merchantability or fitness for a particular purpose.

Introduction

Traditionally, performance has been the primary technical figure of merit for selecting an FPGA for a particular application. Power consumption, while given some attention in planning power supply design and thermal solutions, was rarely a determining factor in device selection. However, as FPGAs have increased in logic capacity and performance by migrating to smaller process geometries, power consumption is receiving much greater attention at the component selection phase of FPGA-based designs. Designers of next generation systems want to integrate more features and higher performance within the same (or often smaller) space and power budget. Furthermore, some applications have specific power requirements that must be adhered to in order to be compliant with system specifications and standards. The Virtex-5 family of 65 nm FPGAs was designed to offer the lowest possible power consumption while not compromising on device performance. This was accomplished by applying innovative techniques at the process, circuit design, and architecture levels of the device.

Benefits of Reducing Power

Achieving lower power operation in an FPGA design has many benefits beyond just meeting the operating conditions of the device. While operating within the specifications of the component is obviously critical to ensuring that performance and reliability expectations are met, how this goal is achieved can have a significant impact on the cost and complexity of the entire system.

First, lowering the power consumption of the FPGA has an immediate benefit to the power system design. Lower supply requirements enable less expensive power supplies that have fewer components and, therefore, consume less PCB area. The implementation cost for a high performance power system is typically between \$.50 and \$1.00 per Watt. Lower power FPGA operation, therefore, contributes directly to lower overall system cost.

Second, because power consumption is directly related to heat dissipation, lower power operation enables simpler, less expensive thermal management solutions. In many cases, designs might need no heat sinks or, at the very least, smaller heat sinks than would otherwise be necessary. In larger, higher performance designs, a passive heat sink could be selected in place of a more costly, less reliable active component. System airflow requirements could be reduced as well. Again, the benefits of lower power FPGA operation result in lower overall system cost and complexity.

Finally, because lower power operation can translate to fewer components and lower device temperatures, system reliability is improved. A decrease of 10°C in device operating temperature can translate to a 2X increase in component life. This fact clearly illustrates the importance of controlling temperature for systems with high reliability requirements. Hence, lower power FPGA operation provides direct benefit to the overall system quality.


65 nm Power Challenges

Total power in an FPGA (or any semiconductor device) is the sum of two components: *static power* and *dynamic power*. Static power results primarily from transistor leakage current in the device. Leakage current is the small current that "leaks," either from source-to-drain or through the gate oxide, even when the transistor is logically "off."

Dynamic power is the power consumed during switching events in the core or I/O of the device and it is, therefore, frequency dependent.

Static Power Challenges

As mentioned, static power is largely a result of transistor leakage current. Basic rules of semiconductor physics indicate that as you shrink the size of transistors (e.g., move from 90 nm to 65 nm devices), leakage current tends to increase. This predicted increase is directly related to the smaller physical dimensions of the CMOS transistors. The shorter channel lengths and thinner gate oxides that are generally used at the new process node make it easier for current to "leak," either across the channel region or through the gate oxide of the transistor. Figure 1 shows the two forms of transistor leakage, source-to-drain (also called sub-threshold) leakage and gate leakage.



WP246_01_050206

Figure 1: Forms of Transistor Leakage Current

Unless significant efforts are made to control leakage, FPGAs fabricated with these deep sub-micron process nodes can have serious static power problems. The first real indications of this potential issue occurred at the 90 nm process node. While Xilinx took significant steps to address static power in Virtex-4 devices and was able to actually decrease leakage relative to previous generation products, other 90 nm FPGAs on the market failed to do so.

For more information on a real-life example of what can happen when leakage concerns are ignored in 90 nm devices and a more detailed look at static power, please see Xilinx White Paper WP221, *Static Power and the Importance of Realistic Junction Temperature Analysis* (<http://www.xilinx.com/bvdocs/whitepapers/wp221.pdf>). At the 65 nm process node, the problem has the potential to become even worse, if proper measures are not taken to control static power.

Dynamic Power Challenges

On the dynamic power side, there are other challenges for 65 nm FPGAs. The simple equation governing dynamic power consumption is:

$$\text{Dynamic Power} = CV^2f$$

where C is the capacitance of the node switching, V is the supply voltage, and f is the switching frequency. The 65 nm process node enables FPGAs that have significantly

greater logic capacity and higher performance than previous generation devices. In other words, more nodes are switching at higher frequencies. All else being equal, this would tend to increase the dynamic power for the largest FPGA devices.

However, there is good news with respect to dynamic power at 65 nm. The core FPGA supply voltage (V) generally reduces with each new process node. And the node capacitance (C) also tends to decrease due to smaller parasitic capacitances (associated with the smaller transistors) and shorter, less capacitive interconnects between logic. These effects offer significant dynamic power benefits for 65 nm FPGAs.

Static Power Innovations

Triple Oxide Process Technology

As with most significant engineering challenges, innovation is the key to resolving them. Xilinx has a long history of innovation dating back to the invention of the FPGA. So it is no surprise that a Xilinx innovation was the key to solving the 90 nm static power problem. In the Virtex-4 family, Xilinx introduced *triple oxide* process technology, which gave Xilinx circuit designers a tremendous tool to fight leakage. The results were so successful that the Virtex-5 family again makes extensive use of this technology to dramatically reduce leakage in the 65 nm process.

In older generation FPGAs, there are typically two gate oxide thicknesses used; a very thin one for the very high performance, lower operating voltage transistors in the core of the FPGA, and a thicker one for the larger, higher voltage tolerant transistors in the I/O blocks. Quite simply, triple oxide refers to the addition of a third medium thickness gate oxide (or midox) transistor that is used extensively in the core of the device to control leakage. Gate oxide formation is one of the oldest and best understood processes in semiconductor fabrication. It is a well-controlled, highly manufacturable process and, thus, it is a relatively simple matter to include the third gate oxide thickness in the 65 nm process integration. Figure 2 shows the triple oxide concept overlaid on the cross-section of a 65 nm transistor.




Figure 2: Cross-Section of 65 nm Transistor Illustrating Triple Oxide Process

The slightly thicker gate oxide in the "midox" transistor dramatically reduces leakage current when compared to a thin oxide transistor of similar size. The medium oxide slightly increases the threshold voltage (V_t) of the transistor and reduces both source-to-drain leakage (also called sub-threshold leakage) and gate leakage. While the midox transistor has lower performance than the thin oxide transistor, this is inconsequential because there are many areas of the FPGA in which performance is not critical.

For example, the midox transistor is used extensively in the configuration memory of the device. The configuration memory contains millions of memory cells that are used to store the unique design image created by the FPGA user. These memory cells are static during device operation, and, therefore, do not need to be constructed with the highest performance transistors. There are literally tens of millions of transistors in the configuration memory that use the midox transistor to reduce leakage by orders of magnitude.

In addition, the "pass gate" transistors used in the routing architecture of the FPGA to connect signals between logic functions can take advantage of the midox transistor. While these transistors need to be able to pass a high-speed signal from source to drain, the gate of these devices is always fixed, or static, during operation. In other words, these transistors do not need to have fast switching times in response to a changing gate voltage. Therefore, the midox transistor is a perfect fit for the millions of pass gate transistors in the FPGA. Again, leakage current is dramatically reduced.

The triple oxide process opens up tremendous opportunities for Xilinx circuit designers to use the right transistor for the right job. In combination with selecting the appropriate oxide thickness, circuit designers can adjust the transistor channel length and other parameters to make sure that each component of the FPGA is optimized for performance and power. The thin oxide, highest leakage transistors are reserved only for the portions of the speed path that require very fast switching times. The result is that leakage current, and hence static power, for the overall FPGA device is significantly reduced.

Architecture Innovations

In addition to the process innovations in Virtex-5 devices, there are new architectural features that further reduce leakage current. Most notable is the new 6-input Look Up Table (LUT6) architecture, the first true LUT6 architecture in the FPGA industry. The new LUT6 design allows roughly a 50% increase in logic capacity per LUT. The net effect is that more logic happens locally within the LUT where smaller transistors are used.

Since transistor leakage is measured in current per unit width (for source-to-drain leakage) or current per unit area (for gate leakage), clearly smaller transistors have less leakage. Fewer large transistors are required because more logic is implemented within the LUT, reducing the need for large drivers to buffer signals between logic functions. The end result is that the new LUT6 architecture has a lower ratio of transistor size to logic capacity, and, therefore, has less leakage than older architectures would have if implemented on a 65 nm process.

Static Power Conclusions

Despite predictions that the move to the 65 nm process node would result in a dramatic increase in static power, the process and architectural innovations implemented to control leakage current allows Virtex-5 devices to exhibit comparable static power consumption to prior generation Virtex-4 devices. Virtex-4 devices

already offer a 40% reduction in static power relative to the prior generation Virtex-II Pro devices and a more than 70% reduction compared to competing 90 nm FPGA products. Thus, the Virtex-5 family continues to hold a commanding leadership position in static power consumption.

Dynamic Power Innovations


As mentioned previously, migrating an FPGA to a 65 nm process brings with it both challenges and benefits with respect to dynamic power. The desire for ever-greater logic capacity and performance certainly tends to push dynamic power consumption higher for users who want to extract the most speed out of the largest FPGAs available in a given product generation. However, the previously discussed benefits of voltage scaling and node capacitance reduction for 65 nm devices can offer dramatic dynamic power reductions for the vast majority of applications. It is also important to recognize that there are two components of dynamic power in FPGAs: *core dynamic power* and *I/O dynamic power*.

Core Dynamic Power

In Virtex-5 devices, the core supply voltage (V_{CCINT}) decreases from the 1.2V used in Virtex-4 devices to 1.0V. Plugging this value into the dynamic power equation (CV^2f) shows an immediate reduction of approximately 31% because of the quadratic relationship between voltage and dynamic power.

In addition, the process “shrink” to 65 nm also tends to reduce the average node capacitance within the FPGA. As shown in Figure 3, this capacitance reduction is the result of two primary factors. First, the smaller 65 nm transistors tend to have lower parasitic capacitance associated with them. This would include capacitance from the gate to the substrate, source, or drain of the transistor.

Second, the routing capacitance between elements tends to be lower in a 65 nm device. This is a result of the physically shorter distances that a piece of interconnect must travel in a 65 nm product. And in addition, Virtex-5 devices use a reduced-K dielectric material between metal interconnect layers to further minimize routing capacitance.



WP246_03_050206

Figure 3: **Reduced Node Capacitance in Virtex-5 Devices**

The estimated reduction in average node capacitance for Virtex-5 devices is 15% compared to Virtex-4 devices. Taken together with the voltage reduction benefit, this translates to an expected 35%–40% total reduction in core dynamic power for Virtex-5 devices, assuming frequency is held constant. Even at 50% higher performance,

Virtex-5 devices demonstrate a 12% power improvement over Virtex-4 devices. [Table 1](#) summarizes the benefits of voltage and node capacitance reduction to core dynamic power.

Table 1: Core Dynamic Power Reduction from Voltage and Node Capacitance Reduction

	Virtex-4 Family 90 nm	Virtex-5 Family 65 nm	% Change	Power Ratio
V_{CCINT}	1.2	1.0	-16.6%	0.69
C_{TOTAL}	1.0	0.85	-15%	0.85
Power	1.44	0.85	-40%	0.59

I/O Dynamic Power

The power consumed from switching of I/O components in an FPGA obeys the same fundamental CV^2f dynamic power equation. However, since the I/O voltage supply (V_{CCO}), drive strengths, and termination requirements are largely driven by the I/O standard being implemented, the I/O dynamic power component is expected to remain about the same in Virtex-5 devices, assuming the same frequency of operation.


Architecture Innovations

As with static power, Virtex-5 family architectural innovations provide additional dynamic power reduction, particularly in the core of the device. FPGA designs that are able to take advantage of these new architectural features can exceed the 35%–40% minimum core dynamic power reduction mentioned above.

First, the LUT6 architecture, that was already discussed with respect to static power, has a similar benefit in reducing core dynamic power. Because more logic is implemented within a given LUT, there is a greater percentage of local, low capacitance nodes. Larger capacitance nodes that would result from the programmable interconnect between logic functions are consequently reduced.

In addition, Virtex-5 devices have a new, more uniform routing architecture that introduces "diagonal" interconnect elements. The result is that there is a higher percentage of 1-hop or "direct connect" routes between logic elements of an FPGA design. This greater usage of 1-hop connections not only tends to increase performance but also has the added benefit of reducing the average node capacitance between logic elements. [Figure 4](#) shows the new Virtex-5 routing architecture with direct connections to diagonal neighbors.

Together, the LUT6 architecture and improved routing pattern reduces core dynamic power by lowering the average node capacitance beyond the level achieved purely from the 65 nm process scaling. The degree to which these features contribute to additional dynamic power reduction is design and implementation dependent. FPGA designs that are optimized to take full advantage of these new architectural innovations receive the most benefit.



WP246_04_050206

Figure 4: Virtex-5 Routing Architecture with “Diagonal” Interconnects

Embedded Blocks

Virtex-5 devices contain more embedded (or hard IP) blocks than any prior generation FPGA in the industry. FPGA designs that utilize these blocks properly can see additional dramatic dynamic power reductions in comparison to implementing these functions in general purpose FPGA logic.

Unlike the FPGA fabric, these hard IP blocks contain only the necessary transistors to implement the required function. There are no programmable interconnects, so routing capacitance is as small as possible. The result is that these hard IP blocks can perform the same function in as little as one-tenth the power of the equivalent implementation in general purpose fabric.

In many cases, embedded blocks that existed in Virtex-4 devices have received significant design overhauls in the Virtex-5 family to improve features, performance, and power consumption. For example, the Virtex-4 family’s 18 Kb block RAM has been redesigned. Virtex-5 devices now contain 36 Kb block RAM modules that, logically, can be used as a single 36 Kb memory or two individual 18 Kb memories.

But what is more interesting from a power perspective is that each of the logical 18 Kb memory blocks is actually composed of two 9 Kb physical memory arrays. To minimize dynamic power consumption, most block RAM configurations require only one of the 9 Kb physical memories within each 18 Kb block to be active (powered up) during any given Read or Write operation. Control logic on the address, input, and output ports of the block RAM ensure that the proper 9 Kb physical array is selected for each transaction. In this manner, dynamic power consumption occurs in only one half of the 9 Kb physical arrays at a time. To the user, however, the block RAM appears as one continuous memory. Figure 5 shows the 36 Kb block RAM in Virtex-5 devices.

The embedded DSP elements in Virtex-5 devices have also been redesigned to incorporate more functionality at higher performance and lower power consumption. On a slice versus slice comparison, the new Virtex-5 DSP slice has roughly 40% lower dynamic power consumption relative to the Virtex-4 DSP slice. This is mostly attributable to the voltage and capacitance scaling factors of the 65 nm process that were discussed earlier.




Figure 5: 36-Kb Block RAM with 9-Kb Physical Memory Arrays

However, because the new Virtex-5 DSP slice has greater functionality and wider interfaces, many DSP operations experience greater dynamic power reduction by taking advantage of these additional capabilities. For example, a 25 x 25 multiply can be implemented in only two DSP slices in Virtex-5 devices, whereas this same function requires four DSP slices plus some additional fabric logic to implement in Virtex-4 devices. In this example, the Virtex-5 implementation results in a 75% dynamic power reduction when compared to the Virtex-4 family. Figure 6 shows the results of this test case based on actual silicon characterization data.

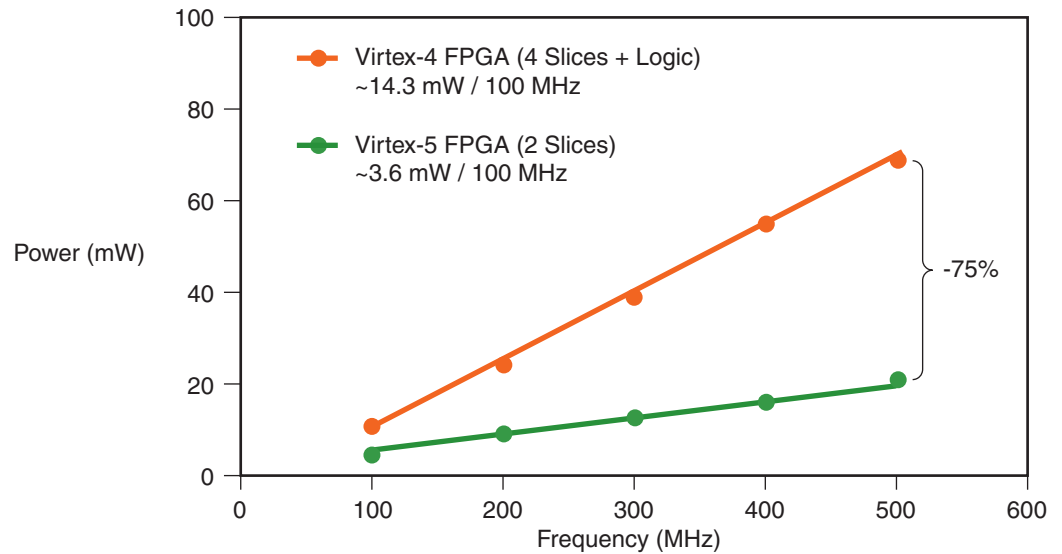


Figure 6: Dynamic Power Comparison of 25 x 25 Multiply

As a final example of power-efficient embedded blocks, the LXT and SXT platforms of the Virtex-5 family include integrated serial transceivers running at rates up to

3.2 Gb/s. These serializer/deserializer blocks were implemented with an emphasis on reducing power consumption. Each full-duplex transceiver in a Virtex-5 LXT or SXT device consumes less than 100 milliwatts of total power at 3.125 Gb/s, representing roughly a 75% reduction relative to Virtex-4 serial transceivers (see Figure 7).

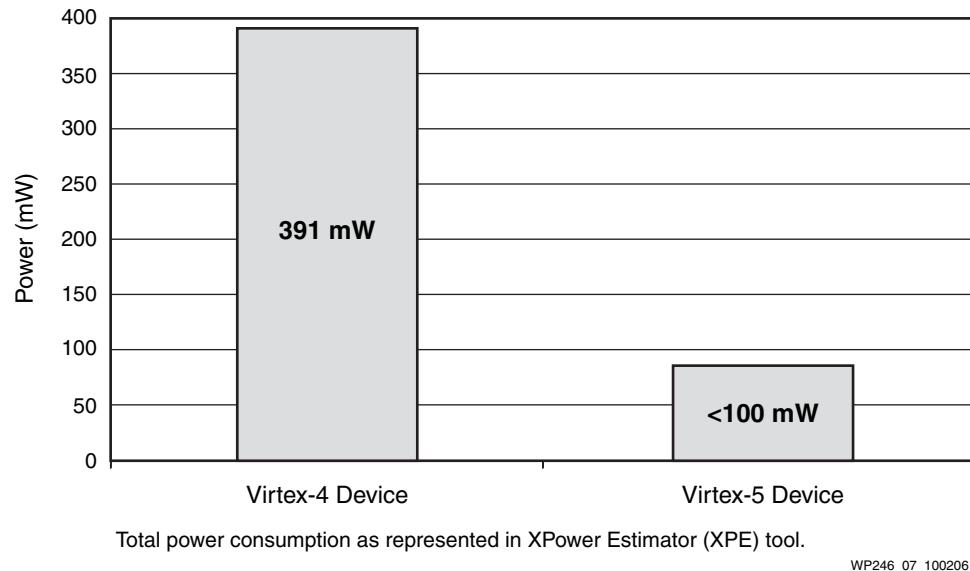


Figure 7: Power Comparison of Serial Transceivers Operating at 3.125 Gb/s

Dynamic Power Conclusions

Virtex-5 devices offer FPGA users a minimum of 35%–40% core dynamic power reduction over 90 nm generation FPGAs. Designs that are able to leverage the new architectural features and various hard IP blocks see even greater reductions. Figure 8 shows the measured results from a single characterization design used to compare early Virtex-5 silicon against Virtex-4 devices on core dynamic power consumption. The design instantiates 1024 8-bit counters into the fabric of the device to create an average toggle rate of 25%. Identical HDL source code was used to target both devices using the same version of ISE tools. While this design does not represent a “real world” application, it does provide one example of the level of dynamic power reduction that can be achieved in Virtex-5 devices.




Figure 8: Example Design Comparing Core Dynamic Power

Total Power Consumption

As outlined in the discussions above, Virtex-5 devices are targeted to have comparable static power to Virtex-4 devices under worst-case temperature conditions. Since Virtex-4 devices already had a commanding 70% lower static power than competing 90 nm FPGAs, the Virtex-5 family continues to offer a dramatic static power advantage in the high performance FPGA market.

While I/O dynamic power is expected to remain relatively constant for any given I/O standard, users should expect to see a minimum of 35%–40% reduction in core dynamic power. Even greater power reductions can be achieved by taking full advantage of the new architectural features and improved embedded blocks found in Virtex-5 devices.

Total power reduction in a Virtex-5 device depends on the ratio of static power, I/O dynamic power, and core dynamic power for any given design. However, core dynamic power is often the largest of these three components, and, therefore, the major reductions that Virtex-5 devices offer in this area should provide a significant total power reduction for the vast majority of designs.

Conclusion

The Virtex-5 FPGA family was designed from the earliest stages of development to offer the lowest possible power consumption without compromising on performance. By using innovative techniques at the process, circuit design, and architecture levels of the device, Virtex-5 devices offer the performance and features required for next-generation systems without breaking the power budget. As the first 65 nm FPGA on the market, the Virtex-5 family continues the Xilinx tradition of setting the standard for the industry.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
05/11/06	1.0	Initial Xilinx release.
10/12/06	1.1	Updated “ Embedded Blocks ” section and added Figure 7 .
02/01/07	1.2	Updated with SXT platform information.