

Reduce Power with Virtex-5 FPGAs

The world's first 65-nm FPGAs offer the lowest power without compromising performance.

by Derek Curd
Senior Staff Applications Engineer,
Advanced Products Division
Xilinx, Inc.
derek.curd@xilinx.com

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 benefits traditionally associated with smaller process geometries: lower cost, higher performance, and greater logic capacity. And although these benefits present exciting opportunities for advanced system designers, the 65-nm process node brings with it new challenges.

Power consumption, for instance, becomes increasingly important when selecting an FPGA for your application. More than likely, your next-generation design will require you to integrate more features and higher performance within a similar (or perhaps even smaller) power budget.

In this article, I'll explore the benefits of reduced power consumption. I'll also illustrate the many process and architectural innovations implemented in Virtex-5 devices to offer you the lowest possible power solution without compromising performance.

Benefits of Reducing Power

Implementing a lower power FPGA design offers advantages beyond simply adhering to the device's thermal operating requirements. Although meeting component specifications is obviously critical for performance and reliability, how you achieve this has a significant impact on system cost and complexity.

First, lowering FPGA power consumption allows you to use less expensive power supplies, which have fewer components and consume less PCB area. The imple-

menting a lower power FPGA design offers advantages beyond simply adhering to the device's thermal operating requirements. Although meeting component specifications is obviously critical for performance and reliability, how you achieve this has a significant impact on system cost and complexity.

Power: Challenges and Solutions

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, the small current that "leaks" from either source-to-drain or through the gate oxide of the

tremendous tool to fight leakage. In older FPGAs, two gate-oxide thicknesses were used: a thin one for the high-performance, lower operating voltage transistors in the FPGA core, and a thicker one for the larger, high-voltage-tolerant transistors in the I/O blocks. Simply put, "triple oxide" refers to the addition of a third, medium-thickness gate oxide (or "midox") transistor that has much lower leakage than the thin-oxide core transistor.

The "midox" transistor is used extensively in the core of the device for non-performance-critical circuits (like configuration memory) or circuits that do not require fast switching times in response to a changing gate voltage (like routing pass gates). The thin-oxide, highest leakage transistors are reserved only for the portions of the speed path that require very fast switching times. The net result is that total device leakage is dramatically reduced, while still offering a substantial performance improvement over previous-generation FPGAs.

The triple-oxide process allowed Virtex-4 devices to reduce static power consumption by an average of more than 70% relative to competing 90-nm FPGAs. The results were so successful that the Virtex-5 family again makes extensive use of this technology to reduce leakage at the 65-nm process node.

Figure 1 illustrates how the triple-oxide process enables 65-nm Virtex-5 devices to achieve comparable static power to similarly sized 90-nm Virtex-4 devices under worst-case (high-temperature) operating conditions, despite industry predictions that 65-nm devices would see a dramatic rise in static power consumption. Thus, the Virtex-5 family retains a substantial static power advantage over competing high-performance FPGAs.

Dynamic Power

Dynamic Power consumption presents other challenges for 65-nm FPGAs. The

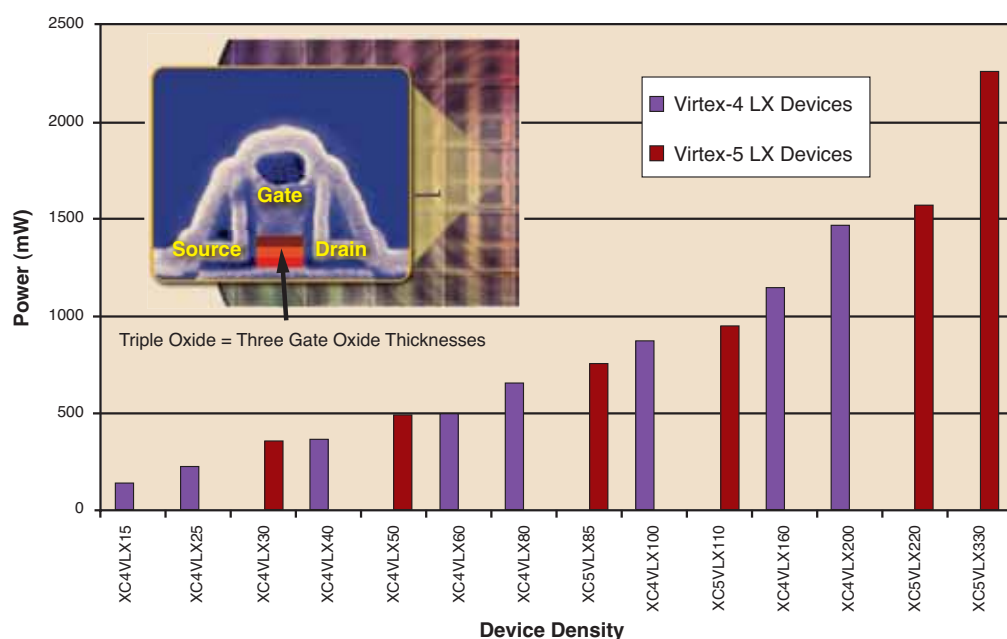


Figure 1 – Static power comparison at 85° C

mentation cost for a high-performance power system is typically between \$.50 and \$1 per Watt. Lower power FPGA operation, therefore, contributes directly to lowering overall system cost.

Second, because power consumption is directly related to heat dissipation, lower power operation allows you to use simpler, less expensive thermal management solutions. In many cases, designs will not need heat sinks, or they will need smaller, less expensive heat sinks.

Finally, because lower power operation means fewer components and lower device temperatures, overall system reliability improves. A decrease of 10° C in device operating temperature can translate to a 2x increase in component life, which clearly

transistor even when it is logically "off." Dynamic power is the power consumed during switching events in the core or I/O of the device and is therefore frequency-dependent.

Static Power

As you shrink transistor size (for example, move from 90-nm to 65-nm devices), leakage current tends to increase. The shorter channel lengths and thinner gate oxides 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.

In the 90-nm Virtex-4 family, Xilinx introduced "triple-oxide" process technology, which gave Xilinx® circuit designers a

equation governing dynamic power 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 older devices. In other words, more nodes are switching at higher frequencies. All else being equal, this tends to increase dynamic power.

However, there is good news with respect to dynamic power at 65 nm. The core FPGA supply voltage (V) and node capacitance (C) generally reduce with each new process node, providing substantial dynamic power savings over previous-generation FPGAs.

In Virtex-5 devices, the core supply voltage (V_{CCINT}) decreases from the 1.2V used in Virtex-4 devices to 1.0V. Node capacitance tends to decrease because of smaller parasitic capacitances (associated with the smaller transistors) and shorter, less capacitive interconnects between logic. Additionally, Virtex-5 devices use a reduced-K dielectric material between metal interconnect layers to minimize routing capacitance.

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 a 35-40% reduction – at least – in core dynamic power for Virtex-5 devices.

Although the “process shrink” to 65 nm provides an inherent 35-40% dynamic power reduction, architectural innovations in Virtex-5 devices offer additional power savings for every design. Most of the node capacitance that contributes to dynamic power is attributed to the routing or interconnect between logic functions. The new Virtex-5 architecture fundamentally reduces routing capacitance in two ways:

- Virtex-5 configurable logic blocks (CLBs) are based on a six-input look-up table (6-LUT) logic architecture, as opposed to the 4-LUT architecture used in older devices. This means that more logic is implemented within each LUT, translating to fewer levels of logic and thus a reduced need for higher capacitance routing between logic functions.

- The Virtex-5 routing architecture now includes diagonally symmetric routes, meaning that every CLB now has a direct “one hop” connection to all of its neighbors, including diagonal neighbors. When a connection is required between logic functions, it is now more likely that this connection is a less-capacitive “one hop” connection, whereas previous routing architectures may have required two or more hops for the same connectivity.

comparison to implementing these functions in general-purpose FPGA logic.

Unlike the FPGA fabric, these dedicated blocks contain only the transistors necessary to implement the required function. And there are no programmable interconnects, so routing capacitance is as small as possible. Fewer transistors and lower node capacitance benefit both static and dynamic power consumption. The net result is that these dedicated blocks can perform the same function in as little




Figure 2 – Dynamic power comparison of counter benchmark design

Together, the 6-LUT architecture and improved routing pattern reduce core dynamic power by lowering average node capacitance beyond the level achieved purely from 65-nm process scaling. Figure 2 shows the core dynamic power measurements from a benchmark design in which a Virtex-5 device and a Virtex-4 device are each filled with 1,024 8-bit counters. These actual silicon measurements illustrate that the combined process and architectural benefits to dynamic power reduction can exceed 50%.

Hard IP Blocks

Virtex-5 devices contain more hard IP blocks (circuitry dedicated to commonly used functions) than any other FPGA in the industry. FPGA designs that utilize these blocks see additional power savings in

as one-tenth the power of an equivalent implementation using the general-purpose FPGA fabric.

In addition to adding new types of dedicated blocks, many blocks that existed in Virtex-4 devices have been redesigned in Virtex-5 devices to add features, improve performance, and reduce power. For example, the 18-Kb block RAM memories in the Virtex-4 family have been sized up to 36-Kb block RAMs in Virtex-5 devices; each of these block RAMs can be broken into two independent 18-Kb memories for backward compatibility to Virtex-4 designs.

Interestingly, from a power perspective, each of the 18-Kb sub-blocks is constructed from two 9-Kb physical memory arrays. For the majority of block RAM configurations, any given read or write request to the block RAM only needs to access one of the 9-Kb

physical memories at a time. The other 9-Kb memory can therefore be effectively “powered down” while it is not being accessed. This reduces power consumption by nearly an additional 50% beyond those reductions resulting from the 65-nm process migration. This “ping-pong” accessing of the 9-Kb blocks is inherent to the new block RAM architecture, meaning that no user or software control is required to take advantage of this capability. It occurs dynamically and automatically, pro-

additional capabilities. In many cases, you can achieve dynamic power reductions as high as 75% when utilizing the full capability of the new DSP slice. If you are not designing a DSP application, keep in mind that you can use the DSP slices for many standard logic functions (counter, adder, barrel shifter) at a substantial power savings compared to implementing the same function in standard FPGA logic.

As a final example of redesigned dedicated blocks, the LXT platform of the Virtex-5

architectural innovations aimed at offering the lowest possible power consumption, while still enabling performance increases of 30% or more.

As Figure 3 illustrates, with static power levels comparable to Virtex-4 devices, the Virtex-5 family provides a clear advantage relative to competing FPGAs. As the only available 65-nm FPGA, Virtex-5 devices also offer a minimum of 35-40% core dynamic power reduction over other high-performance FPGAs on the market. Architectural innovations such as the new 6-LUT and diagonally symmetric routing are likely to enable actual core dynamic power savings up to 50% or more. And taking advantage of the unprecedented level of dedicated blocks lowers power consumption even further.

To find out more about how you can harness the low power of Virtex-5 devices, visit www.xilinx.com/power.

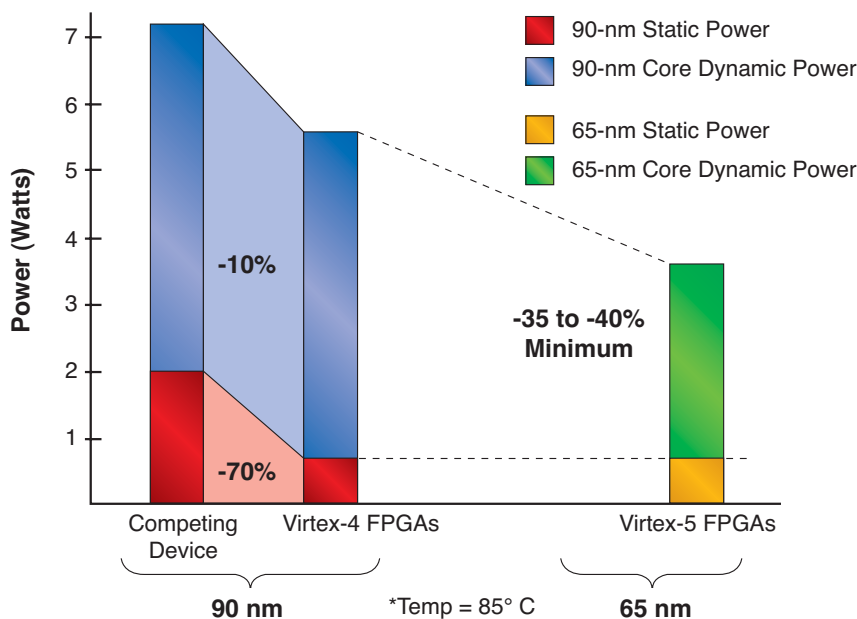


Figure 3 – Power comparison between available FPGAs for a typical design

viding dramatic power reductions for all designs that use block RAM without compromising block performance.

The dedicated DSP elements in Virtex-5 devices have also received a significant design overhaul to incorporate more functionality at higher performance and lower power consumption. In 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 discussed earlier.

However, because the new Virtex-5 DSP slice has greater functionality and wider interfaces, many DSP operations experience even greater dynamic power reduction by taking advantage of these

family includes integrated multi-gigabit serial transceivers, running at rates as fast as 3.125 Gbps. These “SERDES” blocks are implemented with an emphasis on reducing power consumption. Each full-duplex transceiver in a Virtex-5 LXT device consumes less than 100 milliwatts of total power at 3.125 Gbps, representing roughly a 75% reduction relative to Virtex-4 serial transceivers.

Conclusion

Xilinx has a long history of innovation dating back to the invention of the first FPGA more than 20 years ago. So it is no surprise that Xilinx was the first FPGA company to make reducing power a top priority in deep sub-micron technologies. As with the Virtex-4 family, Virtex-5 devices employ a number of process and

Xilinx Power Estimator (XPE)

Introduced in January 2006, the Xilinx® Power Estimator (XPE) spreadsheet-based power tool supports the Virtex™-4 and now Virtex-5 and Spartan™-3E FPGA families. XPE was designed to replace the Web Power tool as the premier pre-design power estimation tool for all new Xilinx FPGA families. The key advantages of XPE over previous power estimation tools are an improved user interface, improved accuracy, and better presentation of important data.

XPE’s summary page displays a complete summary of power usage, first by resource type and then by voltage supply. You can use the navigation buttons on the summary page to access more detailed information. XPE automatically displays several graphs to complete the power usage picture.

Since the initial release, Xilinx has introduced newer versions of XPE that include many additional features and improvements in accuracy. These versions, plus those supporting Virtex-5 and Spartan-3E devices, are available at www.xilinx.com/power.

– Kevin Bixler
Power Tools Product Marketing Engineer
Xilinx, Inc.