

INNOVATIVE FPGA DESIGNS CAN MODERATE THE ESCALATING POWER PROBLEM AT THE EXPENSE OF PERFORMANCE, BUT A TECHNOLOGY SHIFT MAY PROVE TO BE YOUR BEST BET.

Heat wave

FPGAs confront increasing, evolving power consumption

At a glance.....62
CPLDs also suspect64
Web-site launching pads.....67
For more information67

HISTORICALLY, those of you who include programmable-logic devices in your system designs probably focused your power-consumption prediction and reduction efforts on dynamic, or “active,” power draw. This emphasis existed for good reason: Until recently, dynamic power consumption dominated the total power

profile of a chip, and it was a factor you could influence through your design decisions. Static power consumption was primarily an issue with CPLDs’ sense amplifiers, whose power-versus-performance behavior you often could also control. *EDN*’s coverage of programmable-logic power consumption, whose recommended design techniques remain valid, appropriately focused on these topics (**Reference 1**).

Fast-forward to today, though, and static, or “standby,” power draw is be-

coming an increasingly significant factor in overall FPGA power consumption, too (**Figure 1**). If you employ FPGAs that are based on SRAM configuration elements, power-up current surges also demand attention. Although the chip vendor’s design decisions, rather than yours, define a device’s static power consumption, an understanding of static power consumption’s root causes can guide your FPGA-technology, vendor, architecture, device, and packaging selections. Accurate power-consumption prediction, before pro-

prototype implementation, also enables you to appropriately design your system's power-generation and -distribution subsystems, along with its heat-removal apparatus, thereby maximizing the probability that your system correctly functions the first time you turn it on.

WHAT'S CHANGED?

Escalating power consumption isn't an issue that's restricted to FPGAs, of course; it's a phenomenon that's generic to any CMOS-based device that vendors manufacture on today's advanced near- and less-than-100-nm lithographies. Any of you who've, for example, followed Intel's travails associated with its 130-nm-Northwood-to-90-nm-Prescott-CPU transition are familiar with the topic. Intel to date has been unable to translate the 130-to-90-nm transistor reduction to appreciable speed gains, as it has accomplished in past lithography conversions, because to do so would have further bloated the chips' already-hefty 100W and larger power budgets.

Semiconductors are now at a point, with conventional silicon-dioxide process technology, at which ever-shrinking transistor dimensions have in effect run into a power-consumption brick wall. Two transistor specifications are of particular interest: oxide thickness and minimum channel length. The well-known, oft-quoted Moore's Law neatly explains *why* these dimensions are decreasing; chip designers can squeeze more of the smaller transistors onto a given-sized

AT A GLANCE

- ▷ FPGAs aren't immune from the power crisis now facing modern CMOS-based chips.
- ▷ Leakage-current effects are now hindering shrinking transistors' performance improvements.
- ▷ SRAM-based FPGA leaders' application-tuned transistors reduce the leakage-current problem.
- ▷ Switching to an antifuse- or flash-based alternative enables additional improvements.
- ▷ EDA tools are making incremental progress in enabling you to accurately predict chips' power consumption on your design.

cost-effective piece of silicon (**Reference 2**).

Today's transistors' narrow oxides mean that they cannot tolerate the high voltages of times past, but, from a dynamic power consumption standpoint, this increased voltage sensitivity is a plus. Look at the equation $P = 1/2CV^2f$ for dynamic power, where C is the capacitance being charged or discharged, V is the voltage swing the transistor experiences during its transition, and f is the transistor switching frequency. Decreasing voltage counterbalances the increasing tran-

sistor operating frequencies and number of per-chip transistors.

Low-voltage-driven transistors, however, switch more slowly than high-voltage-driven ones unless process- and chip-design engineers also scale down the switching threshold voltage, V_{TH} , or V_T . After all, ever-faster switching with incremental process generations is also an objective of Moore's Law's various corollaries. Unfortunately, at 130 nm and smaller process nodes, the decreasing threshold voltage combines with the earlier mentioned decreasing transistor dimensions to create transistors that are never fully off or on. They constantly leak current; to use a bipolar-transistor analogy, they operate in the active, or linear, region, and they exponentially leak more current with each lithography reduction. This static current comes in subthreshold, gate, and reverse-biased leakage types.

WHAT CAN YOU DO?

The larger the number of conventional CMOS transistors on a chip, predictably, the bigger the leakage-current problem. SRAM-based FPGAs, which use these transistors not only for logic- and embedded-memory-array functions, but also for chip-configuration elements, garner the bulk of today's FPGA business but are most susceptible to leakage-current effects. To combat the problem, Altera employs two types of transistors on its ultradense Stratix II FPGAs. A high-performance transistor with low voltage threshold and small minimal channel length finds use in speed-critical areas of the chip such as DSP blocks and the circuitry within a logic element. A low-power transistor with higher threshold voltage and larger minimum channel length implements less performance-demanding areas of the chip, such as configuration RAM and memory blocks.

Altera's lower-density Cyclone II chips are intended for more cost-sensitive system designs that can't afford exotic power supplies and expensive thermal-management schemes, such as thermal-slug-inclusive ceramic-chip packages, heat sinks, and multiple high-output system fans (**Reference 3**). Therefore, Cyclone II employs a much higher percentage of Stratix II's power-thrifty transistor option. Conversely, Xilinx's Virtex-4 chips sell into the same high-end system designs that Stratix II targets. Xilinx and

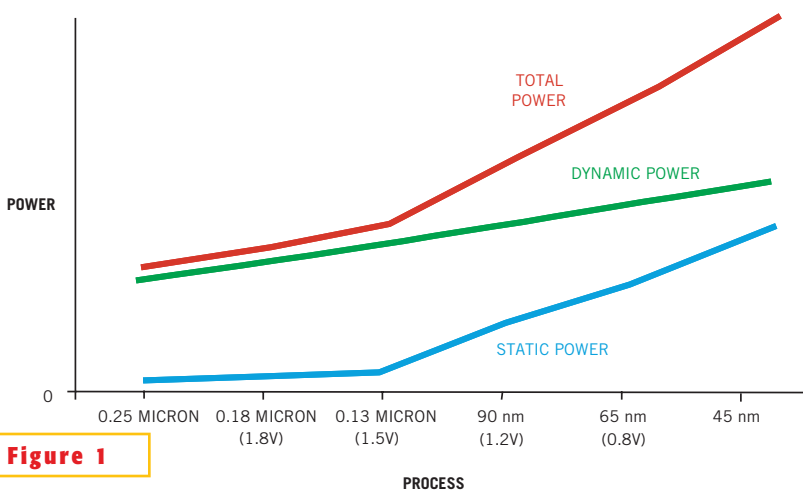


Figure 1

Without architecture and transistor optimizations, whose results this graphic omits, newer FPGAs would exhibit much higher dynamic and—especially—static power draw, even with the same number of I/O buffers and at lower operating voltages than predecessors, due to increased transistor integration, faster clock frequencies, and higher leakage current (courtesy Altera).

foundry partner UMC have reportedly figured out how to fabricate, within a single device, three 90-nm transistor structures with varying oxide thicknesses to hit differing performance and leakage-current targets.

Transmeta's upcoming 90-nm variants of the Efficeon CPU family will take transistor customization to an even higher level by using the LongRun2 technique, dynamically altering transistors' threshold voltages as system-performance needs change. No FPGA vendor has yet announced a similarly complex response to the static-power problem. But perhaps Lattice Semiconductor, which, like Transmeta, uses Fujitsu as its foundry partner, will employ the technique in some future 90-nm-based FPGA family. And, longer term, a number of semiconductor companies, including Intel, as it announced last November, are investigating and implementing high-k dielectric materials that can construct thick, leakage-resistant transistor gates that don't suffer from silicon-dioxide's degraded performance.

Actel and competitor QuickLogic have another idea: Dispense with the transistor-gobbling SRAM-based configuration elements (see sidebar "CPLDs also suspect"). Actel offers two technologies for your consideration: its multiple antifuse-based FPGA families, along the flash memory-based ProASIC and Pro-ASIC Plus chips, whose standby current, according to the company, is as low as 15 mA for commercial-temperature devices and 20 mA for industrial-temperature-screened parts. QuickLogic's latest Eclipse II antifuse family runs at 1.8V and specifies standby current as low as 17 μ A for the smallest device in the family and with internal charge pumps disabled. More generally, both Actel and QuickLogic point out that, because their products run on less aggressive and less leaky 0.13-micron and older processes, their devices' transistors aren't to the same degree subject to static-power effects as the latest generation Altera and Xilinx products are.

Because antifuse and flash FPGAs power up in predefined logic and routing states, you don't need to compensate in your system design for the significant inrush current that results from SRAM-based FPGAs' unpredictable start-up configurations (Figure 2). Your system also needs to endure neither the power draw of the SRAM-based FPGAs'

CPLDs ALSO SUSPECT

Although the primary focus of this article is FPGAs, any transistor-rich device manufactured on advanced-process lithographies is also susceptible to high static-power consumption. Specifically, SRAM-based CPLDs that are or will soon be candidates for the leakage-induced phenomenon include Altera's Max II, which, despite its market positioning, is an FPGA; Cypress Semiconductor's Delta39K; and Lattice Semiconductor's ispXPLD. See "EDN's fourth annual programmable-logic directory," *EDN*, June 10, 2004, pg 49, for more information.

nonvolatile-to-volatile-memory-configuration-transfer process, nor the redundant power draw of the separate nonvolatile-configuration-storage device. Alternatively, you could employ a mask-programmable HardCopy FPGA variant from Altera in your system design or a device with user-programmable logic but hard-wired routing from eASIC or Leopard Logic.

POWER-COGNIZANT EDA

Once you select an FPGA technology, vendor, product family, and device, you can do little in the design process to influence its static-power draw. Compile your design to be as compact as possible, therefore fitting into the smallest device with the fewest transistors, but don't unduly constrain its performance in the

process. Drive the device's inputs to full CMOS levels and connect unused inputs to V_{CC} , too, but accept that you can't ultimately control the leakage current of internal device nodes. However, the chip vendors are addressing your reasonable request for an early estimate of power consumption before you fire up your first system prototype.

Most vendors' Web sites provide for download, at a minimum, various power calculators, which they often base on Excel spreadsheets (see sidebar "Web-site launching pads"). Some of them are frankly crude, although, as a rule, they're more robust than the last time I examined them in 1997. They typically allow you to specify only a device name, thereby determining its transistor count; an internal-clock frequency; and the percentage of the device's internal registers that your design employs, thereby predominantly focusing their attention on nominal active-power consumption. Others are more sophisticated; comprehending static-power estimates and the full spectrum of the device's operating temperature and voltage ranges, enabling you to account for the presence of multiple internal clocks and including the incremental power consumption of I/O buffers, DLLs and PLLs, DSP and memory block, and other dedicated-function circuits.

After you synthesize your design, you have a more detailed understanding of the device logic and memory resources it

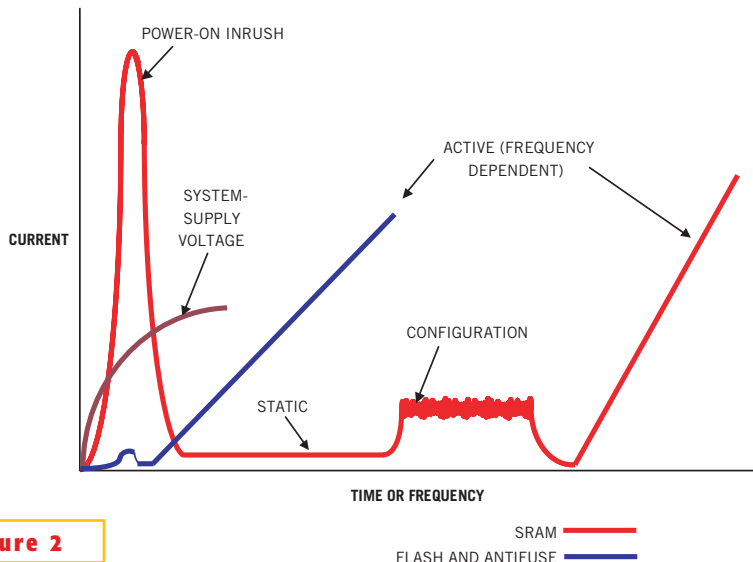


Figure 2

SRAM-based FPGAs are subject to high start-up currents and lengthy configurations that don't plague nonvolatile antifuse- and flash-based alternatives (courtesy Actel).

WEB-SITE LAUNCHING PADS

In researching this article, I found the following power-centric Web pages. Other companies' Web-site resources, although, in some cases, as equally comprehensive as the following list, aren't centralized and therefore require a bit more search-engine assistance to uncover. On the other hand, I may not have stumbled across the relevant summary pages:

- Actel: www.actel.com/products/rescenter/power/
- Altera: www.altera.com/support/devices/dvs-timing_power.html,
- QuickLogic: www.quicklogic.com/lowpower, and
- Xilinx: www.xilinx.com/products/design_resources/design_tool/grouping/power_tools.htm.

employs and how fast they're running than you did when it existed only as HDL code. Some of the tools import the compiled design files for a more accurate estimate of the power consumption you'll ultimately see in real life. One important




Figure 3 Nothing beats lab-bench testing for the most accurate estimates of your design's power-consumption profile (courtesy QuickLogic).

device resource is still missing at the postsynthesis point, though; you can factor in routing effects on the device's operating speed and power only after the design successfully completes the subsequent placement-and-routing process. Actel's SmartPower, Altera's Quartus II Simulator, and Xilinx's Xpower, for example, can all operate at the postlayout stage. SmartPower and Quartus II even estimate power consumption based on a design-specific simulation vector test suite you create, for the—theoretically—most accurate result.

Nothing beats real life for ultimate precision, though. For that, you need to toss your design at one of the vendor's or its partners' evaluation boards (**Reference 4**). QuickLogic, which, with its FPGAs, hopes to claim power as a leadership plank, just as Xilinx does with CoolRunner in CPLDs, has even crafted a power-centric daughtercard for its Eclipse II evaluation platform, containing analog circuitry, a microcontroller and an LCD (**Figure 3**). □

AUTHOR'S BIOGRAPHY



Technical editor Brian Dipert sends good-luck vibes to the engineers and solid-state physicists hard at work solving the static-power problem and thereby extending

Moore's Law's incredible run. Reach him at 1-916-454-5242, fax 1-617-558-4470, bdipert@edn.com, and www.bdipert.com.

REFERENCES

You can find the references to this article on the Web version at www.edn.com.

TALK TO US

Post comments via TalkBack at the online version of this article at www.edn.com.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers directly, and please let them know you read about their products in *EDN*.

Actel

www.actel.com

eASIC

www.easic.com

Leopard Logic

www.leopardlogic.com

Xilinx

www.xilinx.com

Altera

www.altera.com

Fujitsu

www.fujitsu.com

QuickLogic

www.quicklogic.com

Cypress

Semiconductor
www.cypress.com

Lattice

Semiconductor
www.latticesemi.com

Transmeta

www.transmeta.com