


FPGAs BALANCE LOWER POWER, SMALLER NODES

# DRIP



# BY DRIP

THE FPGA INDUSTRY FACES THE SISYPHEAN TASK OF ADDRESSING DEMAND FOR LOW-POWER OPERATION, EVEN AS VENDORS FACE THE LURE OF PERFORMANCE, DENSITY, AND PRICE-PER-GATE ADVANTAGES OF THE 65-NM-PROCESS NODE.

ABOUT EIGHT YEARS AGO, just when FPGA vendors figured out how to increase the gate counts of their devices to rival those of ASICs, the market started demanding higher performance. It took the industry about four years to make these now-million-gate devices run at speeds comparable with those of ASICs. But it did so just as the market made low-power devices its top priorities. So, once again, the FPGA vendors are trying to address demand for low-power operation as they approach ever-smaller process nodes.

This time, however, the task of meeting market demand is more challenging because, in making FPGAs larger and faster over the years, FPGA-chip architects squeezed more power and capacity from silicon mainly at the expense of increasing power consumption. FPGAs got most of their speed increase over the years from using thin-oxide transistors that grow thinner with every process reduction. Thinner gate oxides come with a nasty side effect: They leak power, and leakage, or static power, produces heat. Starting at the 130-nm node, static power in transistors began to explode. It got worse at 90 nm, and, if manufacturers fail to address the issue, it would get exponentially worse at 65 nm (**Reference 1**).

In the race to have the fastest, highest

capacity parts in the 65-nm node, Xilinx and Altera have made power management a top priority. Neither has produced a low-power miracle, so it's unlikely that large FPGAs are going to give ASICs a run for their money as the primary chips in large-quantity consumer-handheld-electronics applications, such as cell phones (see **sidebar** "What drives FPGAs' demand for low power?"). FPGAs still consume 400 times more power than their ASIC equivalents. However, FPGA vendors have seemingly made admirable progress toward stopping the leakage at the 65-nm-process node and make devices at those nodes less power-hungry than their 90-nm devices.

Xilinx claims that it has stabilized leakage and reduced dynamic power from 10 to 50%, depending on configuration, so that its 65-nm Virtex-5, which the company released in May, has an overall lower power consumption than its 90-nm V4 device but with 65% greater density, 30% better performance, and 45% less die area. Meanwhile, Altera claims that users will be able to configure its upcoming 65-nm Stratix III device, due out next year, to consume on average half the power of its 90-nm Stratix. Further, it claims that the 65-nm family will be the highest per-

## AT A GLANCE

▶ FPGAs are 400 times more leakage-prone than their ASIC equivalents.

▶ Xilinx and Altera are stabilizing leakage in their 65-nm devices to the level of their 90-nm devices.

▶ Xilinx's and Altera's 65-nm devices will have lower dynamic power than their 90-nm devices.

▶ The companies achieve most of their dynamic-power reduction from shrinking the process to 65 nm.

formance, lowest power FPGA on the market, with a capacity double that of its 90-nm device.

To address power at the 65-nm node, both companies have attacked the low-power problem on multiple fronts: in circuitry and silicon at the architectural level and in power-savvy design tools to help users manage power in their FPGA designs.

## POWER AT 130 NM

Both Xilinx and Altera say that 70 to 90% of the power savings at the 65-nm-process node come from changes to the circuitry and overall FPGA-chip archi-

ture. FPGA vendors started tweaking their circuits and architectures for low power at the 130-nm node—the first node in which leakage became nasty. Derek Curd, senior staff applications engineer at Xilinx, says that starting at the 130-nm node, Xilinx started to become selective about the types of transistors it was using for each area of the device. In the 130-nm Virtex-2 family, the company used one transistor with higher threshold voltage and longer channels for I/O and used a second transistor with a thinner gate oxide for core logic, which operates at high speeds and lower voltages.

Starting with Virtex-4, the company added a third transistor, which had a middle oxide layer that addresses both gate leakage from the gate oxide of a transistor to the substrate and source-to-drain, or subthreshold, leakage (Figure 1). “We’ve traditionally been concerned with subthreshold leakage, but as we go down in process nodes, the gate leakage is becoming a bigger component of the leakage story,” says Curd. “At room temperature, it can be two-thirds of the total leakage. You can’t control that by making longer channels; you have to do something else. The midoxide gave us a dramatically lower gate-leakage component.”

Altera reacted to the need for low power at the 130-nm node primarily by

moving from a traditional, four-input look-up table to adaptive-logic modules, which users can customize to serve their speed-versus-power requirements. Each module contains look-up-table-based resources; two full adders; some carry-chain segments; and two flip-flops, which designers can mix and match to create logic functions with as many as seven inputs in an adaptive-logic module or a mix of two- to five-input logic functions. Altera also uses thicker oxide transistors in I/O, and its foundry, TSMC (www.tsmc.com), moved to a low-k dielectric. Each of these approaches adds another layer of protection against leakage.

Xilinx also saved power on its 90-nm-node devices by placing more standard-cell hard IP (intellectual property) in its FPGA fabric. Xilinx offers three platform FPGAs at the 90-nm node, each containing hard IP for specific applications. It offers the SX ultrahigh-performance, signal-processing platform and the FX embedded-processing and serial-connectivity platforms. Meanwhile, Altera takes a one-size-fits-all approach with the 90-nm-node Stratix II, gaining most of its power savings from its adaptive-logic-module-based architecture. The company last year somewhat followed the Xilinx model by offering the Stratix GX specialized-platform FPGA, which adds high-performance transceiver IP to the

## WHAT DRIVES FPGAs' DEMAND FOR LOW POWER?

With power consumption and leakage greater than those of comparably sized ASICs, it is unlikely that FPGAs will soon displace ASICs as the main SOCs (systems on chips) in the next generation of cell phones. According to Tim Saxe, vice president of engineering at QuickLogic, “green” requirements drive much of the demand for low-power FPGAs. “Chances are that you spend more money powering up the clock on your microwave

than you do cooking food with it,” he says. “When you use it for cooking, it runs at 1000W only for a few minutes, but that little clock is drawing 9 or 10W 24 hours a day. If you can decrease those 9 or 10W to 4 or 5W, you can make a huge difference.”

The other factor driving FPGAs to lower power is the fear of overheating. Heat increases leakage, and leakage increases heat. More and more, FPGAs

are finding use in applications such as base stations, which can be within units that withstand the elements. This exposure raises ambient temperature. FPGAs may also find use in large, high-speed network equipment. The lack of ventilation, the exposure to sun, or both can increase heat and cause transistors to leak and yield more heat, which leads to thermal runaway and ultimately results in system failure.

Nevertheless, users expect FPGAs to be power hogs in some applications and thus don't lower the power budgets for systems incorporating FPGAs. Vendors such as Altera and Xilinx are stabilizing the power levels of their high-performance FPGAs, doubling capacity, halving die size, and improving performance. All these improvements ultimately allow them to decrease the number of devices in a system.

Stratix II fabric. The company's trump card in low power is Hard-Copy, which allows customers to mass-produce their devices at lower power in a structured ASIC (Reference 2).

To attack power in 65-nm FPGA fabric, both Xilinx and Altera have again significantly changed circuitry and chip architectures. Xilinx has released its V5, and Altera will next year release its 65-nm device.

### INNOVATION AT 65 NM

With its 65-nm Virtex-5 FPGA, Xilinx is using a "smarter mix" of its three transistors, but the biggest change is that it steps beyond the traditional four-look-up-table architecture to a new six-look-up-table architecture (Figure 2). This approach allows the company to use fewer large transistors because more logic processing occurs inside a look-up table, says Curd. Xilinx has also changed the clustering of these six-input look-up tables. In Virtex-4, each configurable-logic block has four slices, and each slice has two look-up tables and two flip-flops. To reduce power consumption, the V5 has four six-input look-up tables and four flip-flops. The total remains the same at the configurable-logic-block level, allowing the company to employ multiple look-up tables, build larger memories and multiplexers, and build wider functions, according to Anil Telikepalli, senior marketing manager for Virtex products. Xilinx is also adding V5 diagonal routing, similar to Cadence's




Figure 1 Xilinx's continued use of triple oxide and hard IP helps the Virtex-5 achieve 35% lower dynamic power and maintain the same leakage rate of its 90-nm Virtex-4.

X-Architecture, as well as traditional, north-to-south, east-to-west routing. "You can now get to the diagonal neighbor directly," says Curd. "One hop gives you lower capacitance than two hops."

The end result is that the V5 has approximately the same leakage as the V4. "If we had done nothing, we would expect a big increase in leakage," says Curd. Xilinx's goal with 65-nm devices is to keep pace with leakage and not follow the predicted upward curve in process and architecture, he says. The V5 has 12 to 40% lower dynamic power than do V4 devices. Most of that dynamic-power sav-

ings results from the process reduction, but some of it comes from the architectural changes. Whereas 90-nm devices have 1.2V core power, the 65-nm Xilinx devices have 1V core power. The 65-nm V5 devices also offer about 15% improvement in internal-node capacitance over V4.

"The transistors are getting smaller, so you have fewer parasitics from the transistor itself and shorter distances between logic," Curd says. "Fundamentally, you get a 15% capacitance reduction. When you multiply that figure with the voltage reduction, you get in the neighborhood of a 40% dynamic-power reduction from the process reductions." Curd says that figure can rise to perhaps 50% power reduction if your design maps well into the V5's six-input-look-up-table architecture, which contributes to the dynamic-power savings, too. He says that, if you tune a V5 LX to run at its highest frequency, 550 MHz, it still has 12% less dynamic power than the V4. Part of the device's dynamic- and leakage-power savings results from Xilinx's weaving in hard-IP blocks. Xilinx plans to offer the Virtex-5 LX platform for high-performance logic, the Virtex-5 LXT for high-performance logic with serial connectivity, the Virtex-5 SXT for high-performance digital-signal processing with serial connectivity, and the Virtex-5 FXT for embedded processing with serial connectivity. Xilinx V5 devices require one 1V power supply for core logic, one 1.8 or 2.5V supply for I/O, and a third for auxiliary power.

### NOT STANDING STILL

Paul Ekas, senior product-marketing manager for high-end FPGA products at Altera, says that, in creating the architecture for its 90-nm Stratix II FPGAs, the company evenly distributed a mix of power-resistant and thin-oxide transistors throughout the device's fabric. Altera also cranked down the transistors' clocks to save power. Ekas says that, in approaching the 65-nm node, Altera created an architecture that reflects real-world applications that require the




Figure 2 Xilinx's Virtex-5 has a six-look-up-table architecture that reduces the amount of necessary routing and logic levels, which also reduces power.

fastest transistors for the critical path. The rest of the design doesn't require the fastest, most leakage-prone transistors. With Stratix III, Altera complements its high-performance logic elements with new, low-power logic and power-down elements for critical paths (Figure 3). "We can change anything that is not the critical path to be low-power logic in the silicon via programming," says Ekas. "During programming, we tell each logic element to be either fast or low-power. For unused logic, you go into power-down mode, making it as little prone to leakage as you can, and you don't route clocks to it, so you isolate it from all signals."

The Stratix III will have a core voltage of 1.1V and higher standard-I/O voltages, such as 1.8 and 2.5V. "For the baseline 65-nm device, you can use a Stratix II power supply, and, if you add a second power supply, you can add a second core voltage," says Ekas. "If you port a design you implemented in Stratix II to a new Stratix III device, you will see a 50% power reduction. If you raise the clock rate of that design 20%, you see a 30% power reduction, and, if you decrease the clock rate of the design by 30%, you get a 70% power savings."

## NEW POWER TOOLS

Both vendors claim that they are adjusting their EDA suites to reduce the number of steps users need to take for managing power. As with their 90-nm offerings, both companies will offer power-estimation, -analysis, and -optimization tools for users concerned about power but whose tool sets automatically manage most of the power. Xilinx's power-optimization tool plugs into the Virtex-5 tool set, and the company is also moving into power-optimized synthesis and physical synthesis. "You get 80 to 90% of the benefit from the architecture itself, but, if you need to scrape off some milliwatts to get into an application, you can use the tool flows," says Curd. Xilinx made its place-and-route algorithms more cognizant of low power. Rather than cluster similar functions in tighter

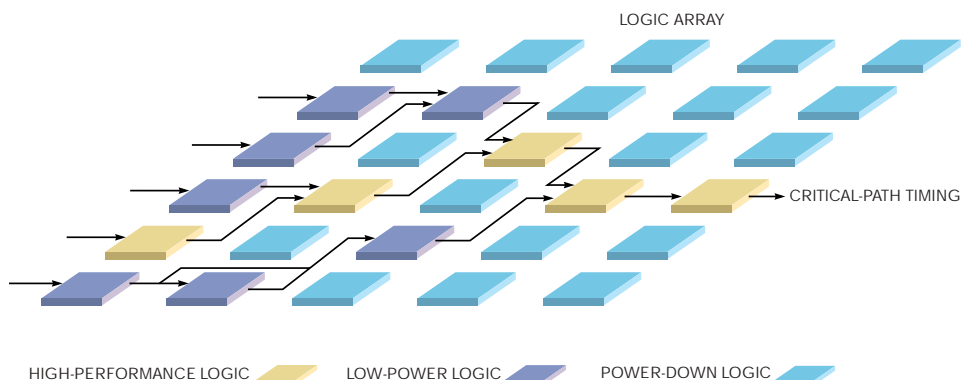


Figure 3 Altera's Stratix III assigns high-performance logic elements to critical paths and low-power logic elements to noncritical paths, and it powers down unused logic elements to reduce leakage.

spaces, the router identifies and optimizes those nodes having the highest switching activity to reduce power. "A popular generic approach to saving power is to pack things as tightly as possible to minimize distances and thus minimize capacitance and therefore power," says Curd. "To bring it to the next level, you have to bring in activity rates. What critical nodes have the highest activity rates? Optimizing those will give you the most benefit." The company plans to add more power-enabled tools this summer when it launches the ISE (integrated software environment) Version 8.2 software.

Altera's power-management functions will be automatic, pushbutton features. Altera also offers PowerPlay software in its Quartus II suite for users who need to design for low power. The suite includes a power estimator for use before synthesis and a postroute-power analyzer. A third power tool performs toggle analysis and helps users interconnect and select logic. However, power management isn't the biggest concern for users of Stratix III, says Ekas. "The big challenge for designers is going to be what you can do with another doubling of gates," he says. Dou-

bling the number of gates means that more designers must work on one FPGA project, so Altera is ramping up team-based design software to go with its 65-nm devices.

Timing closure will also re-emerge as a primary concern for these large devices, so Altera provides the TimeQuest timing analyzer, which features incremental synthesis and a design-space explorer to automatically meet timing constraints. The analyzer runs SDC (Synopsys-design-constraint) format in native mode. Both vendors are also working with commercial EDA vendors to develop power-saving FPGA tools.

## TACKLING LOW POWER

Although Xilinx and Altera are taming leakage in their high-end FPGA devices, many vendors offer smaller, slower devices that suit low-power use. Some devices even specialize in low power. For example, Lattice Semiconductor this year introduced its high-performance, high-gate-count, SRAM-based SC (system-chip) family (Reference 3). Whereas Xilinx's and Altera's 90-nm parts operate from 1.2V core supplies, designers can tune down the 90-nm Lattice SC family to 1V if customers require power savings. "You get a 50% power reduction if you run it at 1V, and it impacts the performance by only 15%," says Stan Kopeck, vice president of corporate marketing at Lattice. "By designing the devices to work over this expanded voltage range, we pro-

MORE AT EDN.COM ▶

+ We encourage your comments!  
Go to [www.edn.com/060608cs](http://www.edn.com/060608cs) and click on Feedback Loop to post a comment on this article.

vide a useful tool to help the system designer dial in performance and power consumption," he says. Both Lattice and Actel also have lineups of nonvolatile FPGAs. The devices have inherently lower power than SRAM-based devices but lack the top performance and capacity of the Virtex and Stratix devices.

Martin Mason, director of silicon-product marketing at Actel, believes that moving to a 65-nm-process node may be

## FOR MORE INFORMATION

Actel Corp  
www.actel.com  
Altera Corp  
www.altera.com  
Cadence  
www.cadence.com

Lattice Semiconductor  
www.latticesemi.com  
QuickLogic  
www.quicklogic.com  
Xilinx Inc  
www.xilinx.com

a bad move for SRAM vendors. "What are they going to give customers at 65 nm? Is it speed, price, power, or are they going

to try to compromise on all three and not do any of them well? Maybe the 65-nm node doesn't bring an awful lot to the party in any of those areas," he says. He asserts that the 65-nm node brings power headaches and that customers, especially those in the "value market," aren't looking for higher performance FPGAs. "From a price perspective, they are pushing the burden onto the board and out of the device," says Mason. He believes that these vendors will increase the total system-cost requirements with additional high-tolerance power supplies, power sequencing, and power management, all of which are driving the analog business to double-digit growth. Actel prefers instead to integrate more of the board and the system by using unique process technology. The company's latest device, Fusion, has a deep-sleep mode, which lets you power it down to 10  $\mu$ A of standby current (**Reference 4**).

Low power has also become QuickLogic's theme. The company's one-time-programmable, antifuse PolarPro and Eclipse II devices require little current and act as gatekeepers to power down power-hungry devices when not in use (**Reference 5**).<sup>EDN</sup>

## REFERENCES

- 1 Dipert, Brian, "Heat wave: FPGAs confront increasing, evolving power consumption," *EDN*, Aug 5, 2004, pg 61, [www.edn.com/article/CA438310](http://www.edn.com/article/CA438310).
- 2 Santarini, Michael, "Structured ASICs deserve serious attention at 90 nm," *EDN*, July 7, 2005, pg 69, [www.edn.com/article/CA621659](http://www.edn.com/article/CA621659).
- 3 Santarini, Michael, "Lattice announces 90-nm high-end and economy-class FPGAs," *EDN*, Feb 8, 2006, [www.edn.com/article/CA6305011](http://www.edn.com/article/CA6305011).
- 4 Santarini, Michael, "Device incorporates mixed-signal circuitry," *EDN*, Dec 16, 2005, pg 19, [www.edn.com/article/CA6290793](http://www.edn.com/article/CA6290793).
- 5 Santarini, Michael, "Low-power FPGAs target portable market," *EDN*, Nov 7, 2005, [www.edn.com/article/CA6281905](http://www.edn.com/article/CA6281905).

You can reach  
Senior Editor  
**Michael Santarini**  
at 1-408-345-4424  
and [michael.santarini@reedbusiness.com](mailto:michael.santarini@reedbusiness.com).

