

Published in issue of Chip Design Magazine

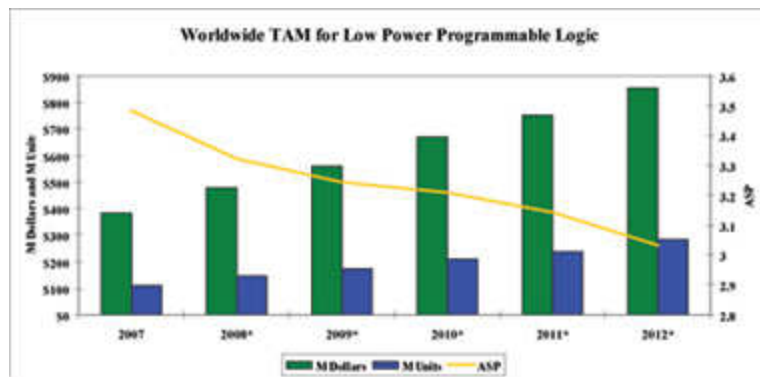
FPGA Vendors Throw Kitchen Sink at Power-Consumption Issues

By Brian Fuller

In the storied history of semiconductors, each era finds vendors generally attaching their strategy to a trendy application segment to differentiate themselves. For years, IC vendors were “computer companies.” Then they were in the “communications” business and more recently they were all about “consumer.”

But the evolution of technology has forced a re-assessment not only of technology but of positioning – to the point where most vendors are attaching themselves to low-power, which spans application segments. Nowhere is this truer than in the programmable logic industry, which for years was seen as a power-hogging alternative to ASICs. In recent years, however, market pressures and product churn have prompted FPGA vendors not only to deliver more features and functionality but lower power, as power-consumption issues bedevil not only consumer applications but take on more global significance as the world worries about dwindling energy resources.

For FPGA vendors, the battle is now joined all the way from their R&D departments and foundry partners to their marketing armies. The available market for low-power programmable logic devices (targeting, for example, battery-based or power-sensitive designs) is estimated to be \$670 million next year, rising to \$855 million in 2012 or roughly 18% of the total programmable logic market, according to Rich Wawrzyniak, market analyst at Semico Research.



“There is a tremendous amount of interest in low-power FPGAs,” Wawrzyniak says. “The low-power phenomenon is something people are going to embrace, assuming they have the right price point and the right performance.”

Different Approaches, Same Message

While the technology positioning may differ in vast and subtle ways, the approach among vendors really doesn't: Just throw the kitchen sink at power consumption and find any way to sand down power, cork leakage and optimize architectures.

Three of the smaller FPGA vendors make an overt show of their low-power strategy: Actel, Silicon Blue, and Abound Technology (formerly M2000).

Silicon Blue, which unveiled its low-power iCE technology a year ago at Computex in Taiwan, focuses on the consumer market where low power and small form factor are paramount design considerations. Its non-volatile CMOS-based SRAM FPGA families claim to draw as little as 3 micro amps in sleep current and 3 milliamps of dynamic current. The iCE Mobile FPGA family, built in 65nm CMOS, also claims not to scrimp on size to achieve power goals with versions ranging from 100,000 to 800,000 FPGA system gates.

Silicon Blue compares its family to Altera's Cyclone or Xilinx's Spartan devices “except we're optimized for power,” said Denny Steele, director of marketing and applications at Silicon Blue.

Because Silicon Blue is a startup, “Every time we have a design choice to make, we could optimize for low power. We didn't have any legacy to consider” as older FPGA vendors do, he added.

At Actel, the differentiation comes at the technology level. Its flash-based FPGAs differ from SRAM-based alternatives from Silicon Blue, Altera and Xilinx. SRAM-based devices experience power spikes at boot-up, which can drain batteries. Flash-based devices do not require an external configuration device to support device programming, Actel notes. That can cut system power by as much as 70%, according to Christian Plante, director of product marketing at Actel.

“You use one chip, not two chips,” Plante says. An architectural feature Actel calls Flash Freeze can drop standby draw to as little as 2 microwatts, he says.

For an increasing number of battery-powered electronics applications, this can be a key differentiator as the flash technology helps reduce leakage at standby.

The low-power challenge is a bit more complicated at the larger, older FPGA vendors, Altera and Xilinx, which spent their early years bulking up on performance and functionality while attacking each other’s market share and nibbling away at the low-end of the ASIC market. The relatively new low-power mandate has forced them to attack the problem in myriad ways.

Xilinx began really attacking the power problem five years ago this month when it introduced its triple oxide process for the Virtex 4 family, says Matt Klein, a member of Xilinx’s technical marketing team who has focused on low-power issues for several years. Thin oxides are key to higher performance, but as they shrink to a dozen angstroms or so in thickness, electrons tunnel away and leakage soars. Xilinx and UMC offered three different thicknesses of the insulating gate oxide layers to lower both static and dynamic power by as much as 50%. Designers could trade off thickness in places where they had to have high performance transistors and places where they needed to optimize and lower leakage.

Since then, the company has taken a holistic view toward static, dynamic and i/o power issues. These includes offering four different transistors at 40nm (low leakage to high performance) for designer tradeoffs; more direct metal connections to cut capacitance and draw less dynamic power; adding mid-level clock gating to allow users to toggle clock trees on an off to manage power draw; making termination for memory interfaces dynamically switchable to select high performance or low power.

These and other steps—including process shrinks— have helped reduce power consumption significantly depending on product, sometimes up to 70%.

Altera first began to attack the problem through software, specifically its Quartus design environment. It implemented programmable power technology that allowed Quartus to bunch up high-speed paths together into logic array blocks (LAB).

“Then what Quartus is free to do is put some LABs into the normal high speed mode, some in low power mode and some can be turned off, so there’s no static power consumption,” notes Umar Mughal, manager for low-cost products at Altera.

In addition, the company has adopted similar tactics to rival Xilinx by focusing on power-hungry memory interfaces, which are getting wider and faster. The company has implemented resistors required for termination onto the chip itself, eliminating the need for extra circuitry. It can save power by as much as 20%, Mughal says.

Power’s Future

The advances FPGA vendors have made in recent years have helped them expand into markets their former power-hogging ways once would have found closed. John Birkner, an FPGA industry veteran who is now vice president of strategic marketing for Silicon Blue, says net PCs are one hot application for low-power FPGAs because the electronics architecture and some standards are still evolving.

Few vendors are willing to tip their power-optimization strategies at this point. Leakage issues will only increase in ultra-deep submicron nodes where process variability will soar. Some surprisingly long-life, low-cost, small form factor battery technology could emerge to save the day, but don’t bet on it. FPGA vendors aren’t.

In preparing for the 32nm manufacturing node, Xilinx has assembled a cross-functional team from IC design to technical marketing to manufacturing which has identified no fewer than 30 concepts for reducing power at that future node.

“We’re pushing these ideas into the IC design group, and a number of them have already been accepted,” Klein says.

For Actel, "the keyword is integration," says Plante. "Another way to put it is system partitioning will be one knob we'll give to the end user to decide how to solve their power problems."



Brian Fuller brings more than 25 years of journalism and new-media experience to his work as writer, editor and communications consultant. Editorial positions he had held include six years as editor-in-chief of the EE Times. He is a well-known speaker and panel moderator in the electronics industry. Brian received his bachelor's degree in English from the University of California, Los Angeles (UCLA).