

# Bringing giant FPGAs to a new node

*Altera Corp's design team for its 40-nm Stratix IV FPGA provides an example of the huge amount of work it takes to be the first to use a new process.* **By Ron Wilson, Executive Editor**

*“This could work only because we both had the same goal: a successful FPGA launch.”*

**MOLY CHIAN**  
Altera

Engaging with TSMC's (Taiwan Semiconductor Manufacturing Co's) 40-nm process-development effort toward the end of 2006, Altera had to simultaneously work with TSMC on process development, develop the FPGAs that the company would build using the process, create the IP (intellectual property), and design tools for the new chips.

This need for near-simultaneity led to a multitrack project organization. One substantial team at Altera worked with TSMC just on process development. Another team worked on the silicon design of the Stratix IV. Two additional teams worked on tool and IP development for the new FPGAs. “We chose to schedule our first tapeout of the design shortly after process qualification,” says Moly Chian, Altera's vice president of technology development.


“To achieve that [goal], all through the schedule the chip design and the process had to advance together.”

#### **Technological development**

In 2006, when Altera first became involved with the 40-nm node, TSMC's process was barely more than a set of equipment choices and design goals. At that stage, the process models suggested only how a 40-nm transistor might look. From this early stage, however, Altera's and TSMC's work would have to converge. To do so, Chian explains, Altera and TSMC had to agree on a two-phase cooperative relationship: “target-driven process development.” During the development, Altera would turn process data into device models, using these models to check against their FPGA design requirements and suggesting changes to TSMC.

“In the first phase,” Chian says, “the models tracked the process development.” During this time, the process

The San Jose, CA, portion of Altera's design team included (left to right) Neville Carvalho, Chong H. Lee, Jeff Watt, Srinivas Reddy, Shawn Wang, Bergen Hung, Peter McElheny, Andrew Leaver, CK Sung, Zunhang Yu Kasnavi, Chris Finan, Elvis Fu, Sergey Shumarayev, John Xie, Bill Liu, and Renxin Xia. Not pictured: Qi Xiang





was evolving rapidly, and the two teams would periodically take a snapshot of the process models, translate these into transistor models. This phase ended nearly as soon as TSMC and Altera could project trend lines for the process parameters, says Chian.

The two companies then moved into the second phase. "At this point, we knew enough about what the process could do and what we needed of it that we could define targets for the process," Chian continues. "Once we agreed on a set of targets, the work changed. Now, the process had to converge on the targets, rather than the models tracking the process."

This technique was a major excursion beyond the comfort zone for many at TSMC. The process engineers in Taiwan have a tradition of being conservative about their commitments until they have silicon in hand. With Altera's team, however, they agreed to commit to a set of targets that looked achievable but that they had not actually measured in silicon. "This could work only because we both had the same goal: a successful FPGA launch," Chian says.

## Chip design

The chip design began with architectural exploration long before Altera and TSMC froze the transistor model. The basis of the early work was the internally developed FMT (FPGA-modeling tool). This system-level simulation environment allows architects to breadboard the characteristics of an FPGA architecture—the circuit and logic cell designs, interconnect-fabric details, power, and performance figures—and then examine the behavior of a complete FPGA with those characteristics. The result is a simulation of the FPGA product that the customer would see: the capacity, performance, power, compilation time, and so on, for real customer-IP blocks.

Thus, the FMT became not only the architects' primary tool for exploring chip-design alternatives, but also the primary way of understanding the impact of process, circuit-design, and design-tool choices on the end product. Reflecting the growing importance that design tools are assuming in the customer's view of an FPGA product, the architectural modeling started in the tools group, not in the chip-design area.

Early on, according to Altera's vice

president for IC design, Richard Cliff, the architects made some rather conservative decisions based on marketing's requirements and results from the FMT. "We chose not to change either the logic cell design or the routing fabric from the 65-nm Stratix III," Cliff says. "This was primarily a time-to-market decision."

Once Altera and TSMC froze the process targets, the chip design could progress from architecture to design of specialized cells, synthesis of the major FPGA components, and test-chip builds. This workflow required teaming not just with TSMC, but with Altera's EDA partners, such as Synopsys. "We had to rely on internal tools early on," Cliff says. "The design exceeded the capacity of the commercial tools. We were seeing runtimes an order of magnitude greater than we had experienced for the previous generation of chips. So, we worked with TSMC and Synopsys to abstract the models and evaluate them faster and to use the tools in novel ways."

A key part of the chip-design effort was the creation of eight test chips during the course of the project. These chips ranged from small structures to help understand the process and define the targets to a complete transceiver channel to thoroughly verify the critical analog-transceiver design. Interestingly, one of the early test chips was a complete Stratix-III-class FPGA, which the designers automatically ported from its 65-nm database. This approach not only gave the team confidence with a full chip in the new process, but also highlighted the areas that would require more than just resynthesis with the new libraries.

The complete FPGA also gave TSMC a test vehicle for its process. TSMC thought of the chip not as a logic device but as a programmable fabric for detecting and diagnosing defects. TSMC then ran the full-FPGA mask set on each of the monthly process runs during development.

As the process neared completion and the Stratix IV neared initial tapeout, white-knuckles time arrived. "You had to be good with uncertainty," Cliff says.

## Software and IP

The first team to jump on the FMT was the tool group—and with good reason. "We had software available to key customers a year before we had silicon," Chian says. "That was not long after we froze the

process targets. So, we really had to trust those targets." It was particularly important, according to Altera's vice president of software and IP engineering, Udi Landen, for the tools team to understand the impact of the chip architecture on compilation times. "For big designs, compiles can run 24 hours," Landen says. "We are always working to reduce that time."

What might be less obvious is that IP development had to start early, as well, partially because some of the critical IP, such as the PCIe (peripheral-component-interconnect-express) Generation 2 core, had both hard-wired and programmable components. So Altera had to complete these blocks by the time that chip tapeout took place.

It was also vital to have the IP design process running in parallel because there was constant interplay between IP, chip, and process design. Early on, for instance, the team determined that the specification for the transceiver blocks was too limiting and that Altera would have to relax it. This change would influence how the highest-speed serial-interface standards would be implemented. Conversely, the performance requirements for PCIe Gen 2, the IP team found, implied that the embedded RAM in the controller block would have to run at 500 MHz. This parameter defined another requirement for the chip-design team.

All of this interaction required central supervision in the form of a program-management office, according to Landen. The office included representatives from each design group within the overall team and had a single program manager for each product family in the 40-nm process. There were more than 100 engineers at six sites worldwide, not including the people at TSMC, says Landen, so there had to be central coordination.

The Altera design illustrates an ambitious design collaboration. But any first adopter of a process node in this range will face similar complexity and risk. Perhaps the most important lessons here involve how Altera attacked that complexity—with centralized management of separate design tracks; designed, rather than accidental, communications links and decision points between the tracks; and an intentional change of the atmosphere between groups from adversarial to collaborative. These are no small achievements in themselves. ■