

# Putting pieces together

Moore's Law is still driving the fpga market, but compromises are needed. By **Graham Pitcher**.

Since Gordon Moore postulated his famous Law in the 1960s, the semiconductor industry has managed to adhere to its spirit: double the number of transistors on a given area of silicon every 18 months or so.

Because the fabric of a programmable logic device is much like that of a memory, both have been at the leading edge of process technology, taking advantages of the benefits of density increases and cost reduction. Both Altera and Xilinx have announced their intention to manufacture at 28nm and both expect to have samples available towards the end of 2010. But technological problems associated with leading edge process technology have required radical design solutions.

Bob Blake, product marketing manager for Altera Europe, said silicon area – and hence product cost – and power consumption were the two major factors pushing fpgas down the technology curve. “Customers always talk about the lowest possible cost and power consumption,” he noted.

The cost issue is dealt with relatively easily: move as quickly as possible to the next process node. But that's becoming increasingly problematic, because while you solve the cost issues, you create a range of other challenges – and power is becoming one of the hardest to crack.

The problem is that as process features get smaller and smaller, the oxide layer in transistors becomes so thin that some current flows even when the transistor is switched off. For leading edge processes, static power consumption is now a major element of the overall power figure.


There are two ways you can deal with the issue. One is through process innovation, the other is through clever design. A number of leading semiconductor companies have decided to follow the high K metal gate (HKMG) approach to reducing static power consumption. Both Altera and Xilinx will be using TSMC's HKMG process for

their 28nm parts. Both companies will be using 'clever design'. TSMC offers a triple oxide process, which means three oxide thicknesses are available and that allows the designers to select the right one for the task – the thinner the oxide layer, the higher the performance of the transistor.


Altera has paid particular attention to the power consumed by transceivers in its 28nm products. One of the major customers for the parts will be those companies designing products for high bandwidth communications. David Greenfield, senior director, Hardcopy, with Altera used an 400Gbit/s data stream as an example. He said this would require 80 of the current 10Gbit/s transceivers used in Altera's 40nm products – 40 in and 40 out. Each transceiver consumes 220mW, giving an overall consumption of 17.6W. But the company has developed transceivers capable of running at 28Gbit/s. The same 400Gbit/s throughput can be accommodated using 32 of these transceivers running at 25Gbit/s. In a 28nm device, each transceiver consumes 200mW/channel, giving an overall consumption of


6.4W. “Not only is power consumption reduced,” he noted, “but there is less board design complexity. A further benefit is the transceivers will work with 25Gbit/s optical systems.”

Perhaps more radical is Altera's decision to allocate far more of the fpga die to Hardcopy blocks – effectively asics. According to Greenfield, the benefits of this approach are obvious. “ASIC



**Fig 1: Altera's 28nm fpga concept**





blocks in the first Stratix parts is one example. What's changed is what customers want to do with our parts. Their requirements have become far more complex and this needs hundreds of thousands of logic elements (LEs) behind the scenes. The ability to hard code these blocks not only saves fpga die area, but also power." But customers may have to wait for a while if they are hoping to include their IP in the Hardcopy blocks in the new fpgas. Blake said it would be a three phase roll out. "In the first step, Altera has decided on some hard IP blocks to include; PCI-Express is an example. In the next stage, we will be looking at emerging standards and may include IP from a partner if it's beneficial. In the final stage, we could take designs from customers with the right volumes, 'harden' them and include them in an fpga."

technologies can give an x20 benefit in terms of die area. While this brings advantages, these blocks are not as fast as they would be in standard cell technology, but do bring enough performance.

Blake expanded on the point. "Hardcopy makes more sense at 28nm because it brings more flexibility. But you have to remember that Altera has always had hard IP in its fpgas; memory and dsp

Blake gave an example of where he believes hardening will bring benefits. "If you look at wireline interfaces, there are applications which could have a 100Gbit/s input to the fpga, which requires about 150,000 LEs, and an Interlaken output, which requires another 150k LEs. At that point, you have already consumed more than half the fabric in the largest Altera device. If you can harden those blocks, you get a lot of LEs back."

Altera has also decided to support partial reconfiguration in the 28nm product family. While not new to the fpga world, it is a first for Altera. Why has it decided to support the approach now? "We're seeing applications which need the ability to reconfigure while the fpga is running," said Blake. "Wireline applications is one area, where different channels need to be supported. Software defined radio is another, where algorithms need to change while the radio is running."

Blake admitted the ability could have been provided in earlier devices. "It's the business case which has driven us to offer it now. It's now the right time for us to get in this market."

"There are customers who want to see reprogrammability in an asic fabric," Greenfield added. "It the market's big enough, we'll talk to anyone."

Greenfield said the approach drew on Altera's incremental configuration technology, which has evolved from block based design. "Designers can choose which blocks can be reconfigured and these can be as small as two or three logic array blocks (LAB)."

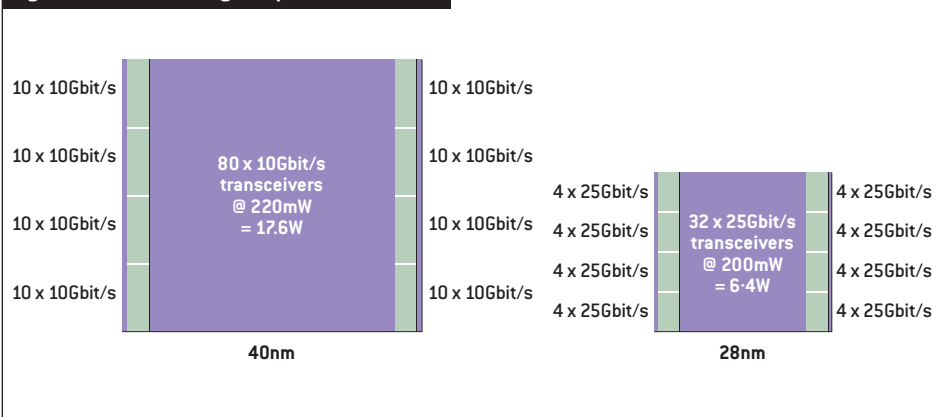
Partial reconfiguration allows fpga I/Os to stay in a known state while the device is repurposed. Although each LAB can be reprogrammed in less than 100ms, the whole device will take about 5s to be ready for its new use.

It's somewhat ironic that it has become increasingly more difficult to differentiate fpgas from the asics and assps which they once challenged. "We have been in conversation with assp companies for some years," Greenfield noted, "but our technologies have different cost structures. However, fpgas at the 28nm node will offer great opportunities for alignment. Using our technologies, it may well be possible for third party IP companies to become silicon vendors."

And with the right IP blocks – Blake's wireline communication example is a case in point – it's going to be harder in future to determine what is an fpga with communications functionality and what is an assp with some configurable logic.

Blake believes Moore's Law is still working for fpga companies. "It's still happening," he concluded, "but we're having to make compromises along the way."

**Fig 2: The benefits of higher speed transceivers**



**ne** For further information on any subject visit: [www.newelectronics.co.uk](http://www.newelectronics.co.uk)