Pushing to 32nm

ST develops advanced platform for networking asics. By Graham Pitcher.

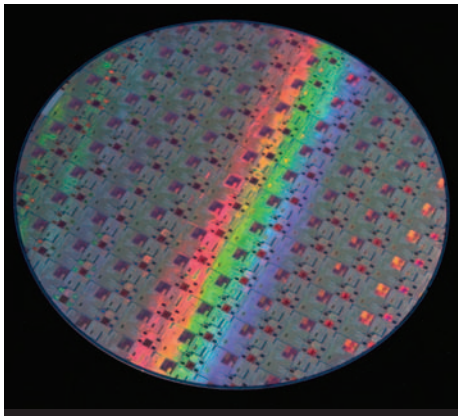

Devices built using STMicroelectronics' 32nm network design platform are likely to include 100million transistors in 200mm² of silicon

Despite many people claiming otherwise, asics are alive and well. And, again in contrast to recent opinion, asics are also alive and well at the leading edge of technology. Although the development costs at the leading edge are escalating rapidly, these are outweighed by the benefits which an asic brings to the application.

For many companies, the leading edge for asic design is 65nm; some are working on newer technologies. But many are using what can be seen as legacy processes.

Firmly on the 'bleeding edge', however, is STMicroelectronics, which has just announced the availability of a technology platform that supports the design and development of networking asics on a 32nm process. Not only does the platform support the leading edge, it also enables devices with a silicon area in excess of 200mm^2.

Developing a 32nm technology platform is not a trivial matter; being first to market with such a solution is equally challenging. What has driven ST in this direction? Riccardo Ferrari is general manager of ST's networking and storage division. "The big problem today with networking is power dissipation. Many data centres not only consume 1MW for processing, they also consume another megawatt to keep devices cool. Another problem is complexity; because data rates are increasing, companies are trying to put as many functions as possible into silicon. With the 32nm platform, we are enabling the next generation of networking equipment and meeting challenging power consumption and silicon integration goals."

Integrating more functions

Counterintuitively, putting more functions on one chip, makes handling connections between those functions easier than it would be if data had be routed between discrete devices, hence the move to 32nm.

ST made the strategic decision to jump the 45nm node in order to be first to market at 32nm. "A couple of years ago, when considering the move from 65nm to 45nm, we looked at the amount of work that needed to be done," said Ferrari. "We decided we could either be late at 45nm or first at 32nm."

The decision was made easier due to ST's membership of the IBM Joint Development Alliance. "That allowed us to start work on 32nm straight away," Ferrari continued, "and we have been able to work with silicon well in advance."

Although describing the technology at the time as 'primitive', Ferrari said that, through cooperation with Philippe Magarshack's central R&D department, ST had been able to tape out the first 32nm device in July 2009.

Magarshack, group vp for technology R&D, noted: "Our role as central R&D is to support all designs and there is more focus on the first design of a new platform, where there is more uncertainty. What we did in this case was to develop technology created by the IBM Alliance partners. We have more than 50 R&D engineers in Fishkill and that allows us to drive ST's requirements and to bring developments back to ST's fabs."

While all partners in the Alliance have access to the 'vanilla' baseline process, each can customise this for their particular needs. "In the case of this design platform," Magarshack observed, "we defined a 10 layer metal stack and decided to develop an embedded dram

STMicroelectronics' Crolles facility houses the 32LPH process

option. We will now be able to make the complete design in Crolles." Magarshack also noted the central R&D group had developed a high speed clocking methodology that allowed clocks for 12Gbit data transfer to be placed and routed.

Reinforcing Ferrari's opinion, Magarshack said the technology was still in its alpha stage when ST taped out the test chip. "Because we wanted to be first to market, we had to go with it."

More than a physical design kit

ST's 32LPH (low power, high performance) design platform for networking applications supports up to 10 metal layers to increase routing efficiency. While based on the 32nm high K metal gate process developed by the IBM Alliance, it also incorporates specific IP and devices developed by ST. These include embedded dram at a density of 10Mbit/mm^2.

How does the design platform differ from, say, a physical design kit (PDK)? "The design platform adds many layers on top of a PDK," said Magarshack. "For example, it will include cell libraries, a cad flow and IP which is integral to the solution – in this instance, embedded dram. The design can then move to the rtl level and on to place and route. In this instance, there's an emphasis on floorplanning and a hierarchy within the flow supports top down timing and bottom up closure."

The design platform is designed to accelerate the development of next generation networking asics for high performance applications such as enterprise switches, routers and servers, as well as optical cross connect and wireless infrastructure.

The 32LPH process can also support large asics. According to Ferrari, die area can be more than 200mm^2, with designs featuring around 100million gates. Having said that, Ferrari pointed out that the largest networking asic designed by ST at 65nm had a die area of 420mm^2. "And we have done three devices with areas of more than 300mm^2."

"Moving to 32nm has advantages compared to 65nm," Ferrari asserted. "Customers like it because higher transistor density allows more complex functionality to be included, but a limiting factor is power consumption." Existing 65nm networking asics can consume up to 50A. "Moving to 32nm means we can reduce current consumption and a further benefit is a reduction in leakage current."

Magarshack: "There are many more transistors available in a given area and that means we can include hundreds of megabits of embedded dram. In fact, in one design, we are looking to include 1Gbit of embedded dram."

However, Ferrari noted the designs being developed are not being targeted at smaller areas; rather, customers are looking to integrate more functions into the same area. At 65nm, asics generally contained around 25 functions. "By going to 32nm," Ferrari pointed out, "that number can be increased – possibly as many as 50." A consequence of this is that power consumption may remain high – often as much as 65W.

An important piece of IP in the 32LPH platform is a serdes, known within ST as S12. This is said to be 'vital' for the development of networking asics, supporting chip to chip, chip to module and backplane communications.

S12, which can be scaled to support eight 12.5Gbit/s transmit/receive channels per macro, has been designed with an optimal footprint for flip chip bga packages. "When you have so many components in a design," said Ferrari, "the problem becomes how to connect to the outside world."

Packaging challenges

Existing complex ST designs use a substrate with 12 layers of interconnect. "We're dealing with something like 6000 bumps and 2500 balls," Ferrari noted. "But we may have to go beyond that at 32nm because customers will want to include as many connections as possible into their designs."

Embedded dram is also important and 32nm enables more of this to be included. "There are many more transistors available in a given area," Magarshack said, "and that means we can include hundreds of megabits in designs. In fact, in one design, we are looking to include a gigabit of embedded dram."

But communications devices don't always need large blocks of memory; what they need is lots of blocks distributed around the chip – often just a few hundred kbits. "Networking applications don't need so much memory," Ferrari concluded. "What is needed is a lot of connections."