

Flexible friends

New programmable solutions challenge the fpga in the communications sector. By **Graham Pitcher**.

A combination of ever changing standards, low volumes and the need for flexibility means there will always be a role for programmable devices in the communications world.

For some time, the sector has been the province of the 'usual suspects'. But new entrants are threatening that position, bringing different solutions to the design challenges.

The communications sector poses particular challenges. Speed is one of the more pressing issues and the solution comes, in general, through the use of asics. But asics bring with them the burdens of high development cost and high risk. Achronix is targeting this particular niche with its Speedster range of fpgas. Yousef Khalilollahi, Achronix' vp of sales and marketing, said: "The primary targets for these devices are those who would like the performance of an asic, but can't afford to specify one because of volume issues."

As an asic replacement technology, fpgas look to provide the speed required by the application, even if they can't match the optimised die size. But there's another requirement: memory.

Achronix has decided to target fabric

performance and its Speedster range supports data rates approaching 1.5GHz. "Many applications now require this level of performance," Khalilollahi continued. "Beyond that is a need for high speed serial interfaces and the Speedster range meets this need with a 10Gbit serdes."

The high fabric speed is supported by Achronix' picoPipe technology. This fabric uses handshake protocols to move data tokens along. The fabric is formed from an array of Reconfigurable Logic Blocks (RLBs), each



KHALILOLLAHI:
"THE PRIMARY
TARGETS ... ARE
THOSE WHO
WOULD LIKE THE
PERFORMANCE OF AN ASIC, BUT
CAN'T AFFORD TO SPECIFY ONE
BECAUSE OF VOLUME ISSUES."

containing eight Look Up Tables. In addition to RLBs, the picoPIPE fabric contains block RAMs and dedicated multipliers, connected through a programmable routing fabric.

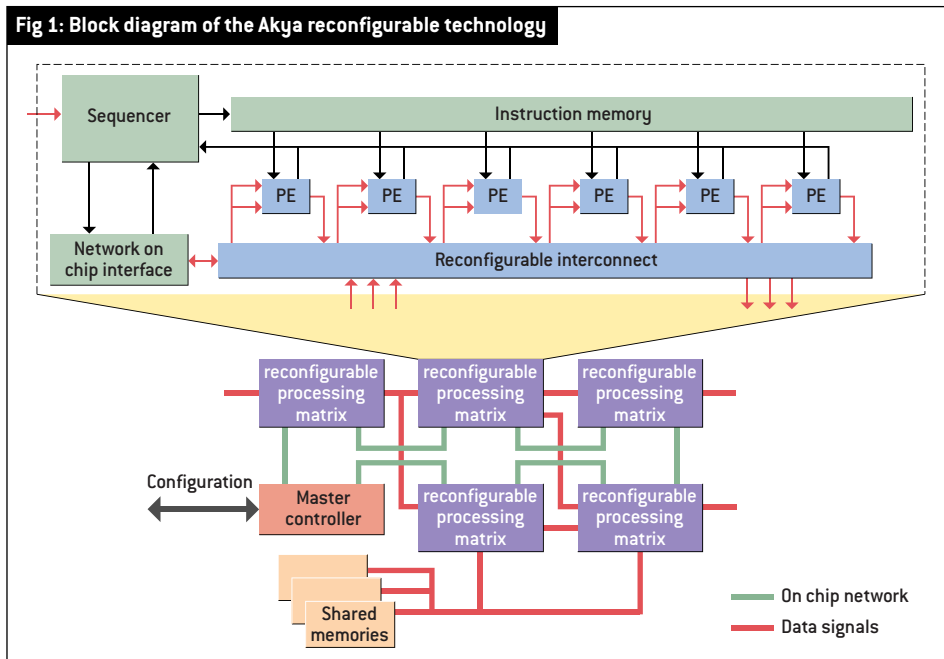
picoPIPE circuits use explicit data tokens instead of implicit ones. Instead of requiring an explicit piece of valid data to be transferred from the input to the output of the flop, the picoPIPE approach encodes the validation signal into the data, and therefore explicit data tokens are used.

As the clock information has been encoded into explicit data tokens, there is no need for a global clock. While data is still clocked between logic and pipeline stages, removing the need for a global clock allows data to be passed at a more optimal rate.

Denny Scharf, Achronix' strategic marketing manager, noted: "Data bandwidth is something you hear a lot about, but memory bandwidth is just as important because there is the need to buffer, gather statistics and so on. The most obvious application is to take packets and store them while you work out what to do with them."

Supporting this latter aspect, Achronix claims it has gone beyond the DDR2/3 PHY by

Fig 1: Block diagram of the Akya reconfigurable technology



embedding a controller on chip which has the ability to run at up to 1066Mbit/s. “We went to a lot of customers,” Khalilollahi continued, “who said they had trouble getting DDR to work. We have supplied something that works out of the box.”

Yet Achronix doesn’t see the changing standards landscape of the communications sector as a major issue; unlike many other companies serving that market. “Memory and data bandwidth are potential bottlenecks,” Scharf noted. “Dealing with changing standards is possibly number two or number three on their list. Right at the top of their list is the need to avoid the NRE costs of an asic. The cost of making mistakes is too high.”

The company does, however, see increasing use of asics for lower volume designs. Khalilollahi added: “If companies expect to sell hundreds or thousands of products a year, then they will use fpgas. But, beyond 100,000, they have to go to asic.”

Colin Dente is ceo of Akya, another company developing an alternative to the fpga. “Mobile handsets feature a lot of places where our technology can fit – it would sit happily with multimedia and wireless tasks,” he said. “Multi standard handsets have a lot of potential. When you get the signal into an intermediate frequency range, there’s a lot of potential for our technology to process data directly, performing

data processing in a flexible way.”

He also sees potential for Akya in the basestation sector. “We can provide the flexibility of an fpga at much lower power consumption and installation cost. This is an area we are very keen to explore and it could have a major impact.”

Akya’s reconfigurable technology (ART) is a dynamically reconfigurable logic technology that can be used to implement general digital logic functions. Dente says it is particularly efficient at implementing dsp functions.

The basic processing block of an ART2 device is a Reconfigurable Processing Matrix (RPM), which consists of a dynamically reconfigurable datapath made up of design time selected Processing Elements (PEs), connected by a Reconfigurable Interconnect (RI).

A program running on the Interconnect Sequencer determines the connections between PEs and the operation of each PE on every clock cycle. The entire configuration of the datapath in an RPM can change each clock cycle under control of the sequencer. Two or more RPMs can be connected together to form an ART Core [see figure 1].

ART2 is delivered as IP, allowing anything from an asic to a part of an SoC or ASSP to benefit from its flexibility and reconfigurability.

The approach is said to bring three major advantage: shorter design times than RTL;

flexible designs which allow modifications without new silicon; and greener products.

In general, basestation developers have elected to use fpgas. The one exception is Ericsson, which has elected to pursue a configurable dsp solution. “We could see our technology being used in a similar way to Ericsson’s approach because it’s all about power; we can provide dsp power in an fpga for something like 10% of the power consumption.”

The reason for this latter statement is Akya’s technology, according to Dente, uses only 10% the number of transistors. “We don’t have the overhead associated with an fpga,” he continued.

Akya refers to its approach as an application area specific device. “It’s a chip which has been designed to be flexible, but which can address the opportunities available in the communications sector.”

Having said that, Dente admits his technology isn’t as flexible as an fpga. “That’s where the application specific area comes in. If you’re designing devices for mobile basestations, there’s a limit to the number of things you want to do. If you want an infinitely flexible design, then you’ll buy an fpga. But, by realising that a design doesn’t need to be general purpose, you can reduce overheads, power consumption and cost.”

Meanwhile, in the infrastructure sector, Dente sees two attractive applications. “One is deep packet inspection. To do this correctly, you need a lot of processing, which our technology can support. Today, the cost per gigabit is too high and this makes deep packet inspection hard to deploy.”

The other application addresses the burgeoning demand for high speed internet access. “Our combination of flexibility and low power is well suited to fibre to the cabinet,” he claimed. “As this rolls out – which it must – we’ll find that the boxes at the end of the street will be stuffed with more and more devices. The problem is the electrical power budget is restricted and performing the necessary processing with a programmable solution is a clear opportunity.”

ne For further information on any subject visit: www.newelectronics.co.uk