



Why academics and companies alike are pursuing asynchronous circuit designs.

By **Graham Pitcher**.

Time is the fundamental measure of digital electronics. The generation and distribution of clocks is something of an exact science and, as clock rates increase, the science becomes even more precise as designers deal with nanosecond time intervals.

So there is more than a degree of culture shock when the idea of clockless processors – or even clockless systems – is raised. Yet asynchronous design dates back to the early days of digital electronics. Simon Moore, senior lecturer at the University of Cambridge's Computer Laboratory, includes self timed circuits amongst his research activities. "The idea was tried initially in the 1950s, when people were dabbling with digital circuit design. There was a revival of interest in the 1990s, when a happy band of researchers started looking at the fundamentals."

But why should there be any interest in asynchronous designs? One answer is power consumption – a problem that

Beat the clock

starts to get worse as process technology continues to shrink. Moore noted: "High end processors burn around a third of their power consumption distributing clocks and around a third is accounted for by leakage current. That means only a third of the power consumed is available for computation and data movement. So the logical extreme is to remove the clock."

However, this logical extreme is not easy to achieve and doesn't always bring the benefits that may be imagined. Moore noted: "It's not clear what the advantages of pure asynchronous design are, but there are two immediate benefits. One is lower emi, because comput-

ing activity is not correlated to the clock. The other is composition. People talk about the pain of getting timing closure. If you don't meet timing requirements, then your design won't be correct. In the asynchronous world, the design will always be correct, but will you meet your performance goals?"

This apparent contradiction has resulted in people looking at how to mix synchronous and asynchronous circuits to the best advantage. And Moore is interested in this area, developing ideas that subscribe to the GALS – Globally Asynchronous, but Locally Synchronous – concept. "Clocked systems continue to



ARM goes clockless

Whilst there has been quite a lot of academic activity, commercial exploration of asynchronous devices has been less evident. One recent instance was in November 2004, when ARM announced a joint venture with Philips' incubator company Handshake Solutions to design a clockless ARM processor. The project was expected to produce results within a few months.

Richard York, ARM's cpu product manager, said the process has taken longer than expected. "We wanted to have something for mid 05, but the last year or so hasn't been wasted. We've been building a good proof of concept, which is as well tested and validated as any other ARM processor."

Why would ARM want to explore this area? There are two obvious reasons: lower power consumption; and lower emi. "It's an interesting area to get into," York noted, "and we're willing to chance our arms."

Now – later than expected – ARM has unveiled the fruits of its work. "It's a point solution for low cost microprocessors in such sectors as automotive and medical, where self timed circuits work," said York. "It's ideal for applications where you don't want to generate noise and for applications where there's virtually no power supply." Originally, the design goal was an



York: "It's not the speediest processor we've built, but that's not an issue."

ARMv5TE equivalent, but the team has created something that resembles an ARM9. "It's not the speediest processor we've built," York admitted, "but that's not an issue compared to noise, for example." The part has been designed on a 0.13µm process and early results are said to show that power consumption is less than half that of an equivalent synchronous ARM9.

The five stage pipeline device looks like a typical embedded ARM9 core and has a similar gate count. "That's an important achievement," York continued. "Earlier asynchronous designs have been twice or three times the size of a synchronous equivalent."

work, so why not keep doing it?" he asked. The GALS concept creates systems with 'synchronous islands' joined by asynchronous interconnects. "This allows you to put asynchronous technology where it's needed. There's quite a lot going on here," he added.

Richard York, ARM's cpu product manager, agreed: "Clock distribution across big chips is getting horrendous, so local self timed islands could be a good way of dealing with this."

In Moore's opinion, a lot of people concentrate on distributing clocks when they should be looking to distribute frequency. "The key thing about clock distribution is

getting the edges to appear simultaneously – low clock skew – and that's hard. Distributing frequency is easier. The issue then is how to move data between clock domains knowing these domains produce and consume data at the same rate."

Another challenge which Moore is addressing is how to 'mix and match' technologies to do useful things. "What are the challenges when we get to 10nm processes?" he asked. "There'll be a lot more device variability, which will make static analysis more difficult. It can be done, but with big error margins. At that point, conventional clock design doesn't look so attractive."

Asynchronous design involves, at times, some suspension of belief and Moore gives an example. "We've been working with Cambridge Consultants to develop an asynchronous implementation of its XAP processor. In order to test the asynchronous version, we included a clocked XAP. But we had to clock it." Moore said the solution was a self calibrating delay line, driven by a watch crystal. "And we used an asynchronous approach to build the delay line."

Belief needs to be suspended in another of Moore's research areas. "We've also done some work on the generation and distribution of clocks. One of the problems is that they are normally H trees and you get dislocations in time because clocks come from different branches. It's much better if you have the clock skew distributed uniformly and we've done this by building an oscillator grid that distributes time normally across the chip at high frequency."

Moore's solution is based on an asynchronous fifo. "If you imagine this wrapped in a circle, then empty or full data items spin around the ring," he explained. "There are two effects to be observed. One is a 'draughting' effect, in which clock edges catch up with each other, producing 'bunching'. But you also have the effect that, when inputs change simultaneously, the gates work more slowly. We use this to push the events apart. If you then look at the ring at different points, you observe a multiphase clock."

This can also be used to distribute as well as generate time, though a ring is rather one dimensional. To generate and distribute a clock over a 2d area, Moore's team has designed an oscillating grid.

He sees asynchronous design as becoming more of a mainstream technology. "There are text books appearing and the consensus on how to do things is a healthy sign. As people start to get on top of a subject, you get clarity. Whilst this has started to happen with asynchronous circuits, there's still work to be done."

As pure clocked design becomes intractable, Moore concluded, it will become increasingly important for engineers to understand how asynchronous techniques can ease the pain. 