



# Law breakers



Electronics researchers continue their fight against the Laws of Physics.

By **Graham Pitcher**.

**R**esearch and development is fundamental to any industry. But electronics can be singled out for one good reason: much of the R&D work is taking the Laws of Physics on squarely.

Moore's Law set the industry on this course some 40 years ago, when Gordon Moore predicted the number of transistors on a chip would double every 12 months. Ever since, the next process technology node has been approximately 0.7 ( $1/\mu 2$ ) of the existing one, because Moore's Law is area based.

So 90nm production is being superseded by 65nm and so on. All very well, but these dimensions are substantially smaller than the wavelength of the light used to create them. Currently, manufacturers are coping with the problem, but not for much longer.

This means there has been a lot of work in developing lithographic technologies. So it's no surprise that breakthroughs have been made – and strategies changed.

In February, IBM announced its belief that 193nm lithographic techniques may last for another seven years. Dr Robert Allen, manager of lithography materials at the company's Almaden research centre,



noted: "Our goal is to push optical lithography as far as we can so the industry does not have to move to any expensive alternatives until absolutely necessary."

Collaborating with JSR Micro, IBM managed to print 29.9nm features using its NEMO technology in conjunction with JSR's SOLOnX approach. NEMO uses two intersecting 193nm laser beams to create an interference pattern, whilst SOLOnX brings the benefits of immersive lithography.

ARM's asynchronous ARM996HS isn't the speediest thing it has built, but that isn't seen to be an issue when power consumption and noise are considered.

A change in strategy came from Intel – previously an enthusiastic supporter of extreme ultraviolet (euv) lithography. Apparently dogged by technical issues – including power levels and light sources – Intel has now widened its focus to embrace immersion lithography. However, euv



**"Older asynchronous designs would have been two or three times the size of a clocked design."**

Richard York, **ARM**

research leader ASML is said to be close to resolving at least some of the issues.

From the manufacturers' point of view, both TSMC and TI are 'happy' with immersion lithography. According to TSMC, tests have been promising enough that it can now 'focus on throughput improvement for high volume'. TI, meanwhile, announced recently that it will use 193nm immersion lithography for its 45nm process.

Whilst equipment manufacturers were looking to 45nm processes and beyond, semiconductor companies were beginning to address the reality of 65nm production.

Illustration: Ely Walton



Xilinx was amongst the first to trail its move to 65nm when, last November, Richard Terrill, senior manager for Spartan and cpld solutions, said: "We're looking past 90nm and have extended our relationship with UMC and Toshiba to 65nm and beyond." And it duly unveiled its 65nm product in May, with the launch of Virtex-5 (covered elsewhere in this issue).

Altera also discussed its use of 65nm technology. Taking advantage of manufacturing partner TSMC's process, Altera believes a 30% drop in clock rate from 90nm to 65nm will bring a 70% cut in power consumption. When Stratix III appears later this year, it will also offer programmable power – the ability to put performance where it's needed, said vp of European business operations Tim Colleran.

Meanwhile, Xilinx underlined the growing importance of Europe by opening a research centre in Dublin – the first time its Research Labs organisation has set up outside of the US. Ivo Boelens, Xilinx' cto, noted: "We're honoured to join forces with the best and the brightest researchers in Ireland and throughout Europe in our quest to invent the future of scalable, reconfigurable computing architectures."

Power consumption is one reason why microprocessors are going multicore. High end microprocessors have high power consumption and around one third of that power is consumed by distributing clocks and another third by leakage. Simon Moore, senior lecturer at the University of Cambridge's Computer Laboratory, said: "That means only one third of the power consumed is available for computation and data movement. The logical extreme is to remove the clock."

And that's what ARM announced in February. Richard York, ARM's cpu product manager, admitted it had taken longer than anticipated to produce, but claimed the ARM996HS 'looks like a typical embedded ARM9 core'. The device consumes half the power of its clocked

equivalent and has an equivalent gate count. York pointed to the latter achievement: "Older asynchronous designs would have been two or three times the size of a clocked design."

Simon Moore believes asynchronous design is now entering the mainstream. "There are text books appearing and the consensus on how to do things is a healthy sign." But he admits there's still work to be done.

### Bringing structure

With devices being made on ever smaller process technologies, asics become increasingly the province of those with the deepest pockets. This has not only seen fpga fortunes rise, but also the return of the gate array in the form of the structured asic. And this technology turned out to be something of a minefield in the past 12 months.

We asked in September 2005 whether structured asics had reached maturity. With a number of companies entering the market, it looked as though the technology was 'hitting the spot'.

Our comment at the time was that structured asics seemed to be defining themselves more at the physical layer, for example, featuring prediffused PHYs. Proving the point, ChipX launched its CX6200 family in October, with the 12 devices in the family each boasting an integrated PCI Express PHY.

Yet, whilst companies were joining the structured asic market in September last year, a founder of the technology – LSI Logic – decided to leave it. Did this leave the structured asic market between a rock and a hard place: the rock being the traditional asic; the hard place being the fpga? Structured asics may well survive, but they may not have the impact that we might have expected.

All these technologies require good design tools. One three letter acronym that has suddenly burst into the forefront is

EDA tools – and electronic system level tools in particular – are helping designers overcome the complexity of working with leading edge processes.



"If you don't have the latest process technology, you won't get the results."

Andrew Benn, **Atmel**

DFM – design for manufacture. It's a particularly sticky problem at 65nm; so much so that IBM has released a 'fully statistical' toolset to help selected companies produce better chips. This kind of data has been closely guarded in the past, so the move emphasises the importance with which DFM is being regarded.

The other big problem at 65nm is verification. The two big eda names – Cadence and Mentor Graphics – both introduced verification packages recently.

Mentor introduced Calibre nm in response to 'orders of magnitude' changes in verification complexity. Included are such elements as lithographic friendly design, design rule checks and resolution enhancement. Joe Sawicki, general manager of Mentor's design to silicon division, claimed Calibre nm would remain 'the fastest, most complete verification tool ... for the next five process nodes'.

Cadence, meanwhile, has launched a range of verification tools. In particular, it unveiled a verification kit for ARM processor based designs said to offer a 'low risk path to verification closure'.

Design complexity was also reflected in the development of more electronic system level – or ESL – tools. Two major European projects are pursuing ESL design: Sprint is an ECSI project focused on IP reuse and integration; whilst a MEDEA+ project is developing a design flow focused on hardware/software codesign.



Xilinx also put its weight behind the technique with the launch of the ESL Initiative, aimed at making the tools and methodologies more widely available.

### Conversion

For all the work done on developing digital techniques and processes, there is still work to be done in the analogue domain. Promoting a different approach to a/d converter design was start up Signal.

Its technology is continuous time delta-sigma (CTDS) conversion. The advantage of CTDS is there is no acquisition phase, which means a high performance sample and hold stage is eliminated. Neither does the approach need high gain bandwidth stages to force settling, so power consumption is reduced.

Signal's first product is a low power CTDS alternative to pipeline converters, said to be a complete data conversion system designed to operate seamlessly over a range of sample rates.

Engineers at Atmel and National Semiconductor were pushing conversion limits earlier this year. Atmel launched what it claimed was the fastest monolithic 12 bit a/d converter. The 500Msamples/s AT84AS001TP will allow 250MHz signals to be digitised. The secret, said Andrew Benn, the company's data converter marketing manager, is process technology. "If you don't have the latest process, you won't get the results."

National Semiconductor addressed gigasampling requirements with a dual

Sitting between fpgas and asics, structured asics – or platform solutions – help fill the design gap.

channel 8bit device. Jumping to the gigasample rate required more advanced process technology – 0.18µm cmos – and a change to a folding architecture, rather than the previous two step approach.

### Memories ...

Alternative approaches to embedded memory have been researched for some time, but Freescale claimed a major breakthrough at the end of 2005 with the announcement that it had 'overcome substantial process challenges' to produce a 24Mbit nanocrystal array.

"We have a nanocrystal process that's reproducible," said Barry White, manager of the company's advanced materials, memory and interconnect operations. "We're targeting products at the 65nm node in 2008," he added.

But just a couple of weeks ago, Freescale brought products based on its magnetic ram (mram) technology to market. Aimed at combining sram speed with flash's non volatility, mram allows data to be changed more quickly.

The biggest challenge for Freescale was depositing a 1nm barrier layer across the whole of a wafer to a very tight tolerance. Now, a 4Mbit mram is available and Freescale says the process is scalable to 65nm and beyond.