



[Getting around multicore walls: The roads less traveled](#)

[Rich Pell](#) - February 23, 2012

For the better part of two decades, the processor industry has been running pell mell down the road of multicore design, packing more and more processor cores on a single chip. But a funny thing happened on the way to the personal super computer. It didn't work.

In 2007, a DARPA study on the potential of an [exascale computer](#) concluded that with the current architecture of processors, in other words x86 and PowerPC, we could not get there from here. As a result, in January 2012, DARPA announced the Power Efficiency Revolution for Embedded Computing Technologies, (PERFECT), to figure out what to do next.

Dr. David Patterson, a RISC pioneer and a leading voice in the development of multicore processors, suggested that the problem could be solved using FPGAs as an experimentation platform in a program he called Research Accelerator for Multiple Processors (RAMP) at UC Berkeley [in a series of RAMP presentations](#). In a related 2006 Berkeley white paper, 'The Landscape of Parallel Computing Research: The View from Berkeley,' Patterson said that the power consumption of the logic in the CPU, converted into heat, limits the performance.

Since any heat that cannot be removed by a heat sink reduces the performance of the transistors, the results are:

- If you increase the system clock to boost performance, heat rises, transistors slow down
- If you increase memory bus width and you increase the number of transistors, heat will increase and transistors slow down
- If you increase instruction-level parallelism (ILP) so more can get done at the same time, you increase the heat and...

The result of the RAMP effort? "The memory wall has gotten a little lower and we seem to be making headway on ILP, but the power wall is getting higher," Patterson says. One anonymous engineering wag put it more succinctly:

"We're screwed."

Throughout this process, however, there have been voices, crying by the roadside as it were, "Go back! you're going the wrong way!" And it may be time for those voices to be heard.



A significant wall facing multicore design is memory bandwidth. While each core can access a small amount of on-chip memory, the sharing of the bulk of stored data must pass through significant bottlenecks, which aggravates the problem of power. (Illustration by Doug Davis, courtesy of Venray Technologies.)

Going back to the turn of the century, companies like UK-based Celoxica were pointing out the weaknesses of the multicore approach in contrast to a heterogeneous approach incorporating FPGAs.

"The first problem is the architecture of the standard processor doesn't lend itself to parallelism," says Jeff Jussel, former VP of marketing for Celoxica and current senior director of technical marketing for Element14. "No matter how many processors you put on a chip, we are not seeing any one algorithm processing faster because it is too hard to program access to all the processors. What you end up with is a system that can do 12 things, but no actual system speed increase with an incredible power increase."

Celoxica's approach, according to Jussel, was to break up the algorithm over multiple processors inside the FPGA with dedicated memory. "You end up with millions of tiny processors optimized for the algorithm. When the algorithm changes, you just reprogram the FPGA."

At the time, the problem was not immediate and the market was entrenched. Celoxica ultimately spun off their tool business, eventually landing in the hands of Mentor Graphics and kept their board development business. That business was always focused on one-off applications, ranging from mining to financial services.

Patterson says their work in RAMP showed that an FPGA approach, especially for highly focused applications, was "more and more attractive" but there were two specific obstacles: power and design tools. "We found the CAD tools available were just not that easy to work with. We were actually surprised at how difficult it was and it formed a major stumbling block. And while FPGA providers have gotten better with power even as the number of transistors increase, they still need to get better with it before it can be a mainstream answer."

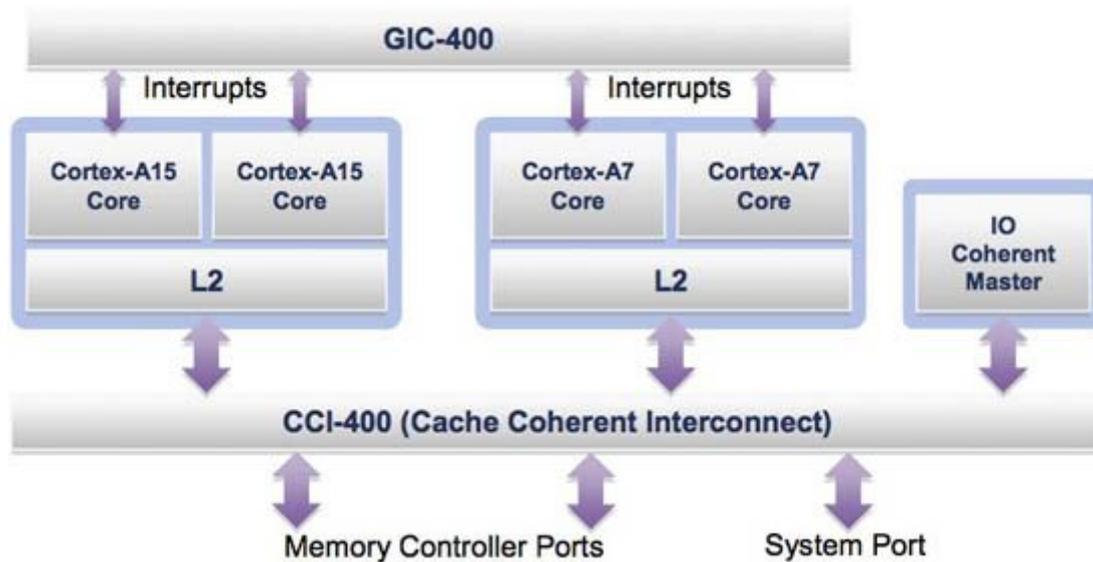
The reprogrammable nature of an FPGA has allowed several board- and system-level companies, ranging from Wall Street FPGA in financial analysis markets to Convey Computing in scientific analysis, to assign ARM cores or small hard-configured RISC blocks, like MicroBlaze, to a variety of tasks, with subroutines handed off to coprocessors on the same FPGA. But the dream of a fully retargetable FPGA system, touted in the mid-2000s by companies like Quicksilver, has been largely deferred because of the problem of developing parallel multithread software for such changing architectural types.

Think "many" not "multi"

ARM walks into the fray almost in a position of neutrality. While it still endorses the validity of Intel's homogeneous approach to multicore, as early as last fall it began discussing a "[many core](#)" as opposed to multicore approach. According to John Goodacre, program manager in the ARM Processor Division, the traditional approach of using full-performance cores still has a long road ahead of it, especially when you are considering dual- and quad-core designs, but it may not be necessary, especially in some consumer applications to use the large cores.

"Mobile applications are full of little processes," Goodacre explains. "If you put all those processes into four or eight big cores, you don't actually see a big performance improvement, but you see quite a big negative power impact. A many-/multi-processing approach duplicates the capability of a big homogeneous multicore design that is inherently more power efficient."

Goodacre points to ARM's [big.LITTLE](#) concept that marries an A15 - which he claims is capable of running more of today's dual-core-type software - with four small A7 cores, in a power-efficient formation.



ARM's big.LITTLE concept is a heterogeneous computing approach putting a number of dissimilar, specialized cores on the same slice of silicon - CPU and GPU cores, for example - and parcelling out tasks to each core as necessary.

"This approach is mostly targeting toward power, but it's also giving the next-generation programmers the concept that there's also a lot more power-efficient processes available for that next generation of software. The first next-generation software I anticipate will be in gaming, but as time progresses and [there is] more and more availability of more cores, there will be more software available."

From the software side

Architecture experts developing RISC instruction sets for a mix of server and embedded applications - dominated by ARM, but also including MIPS, Tensilica, and other companies - have offered their cores to standard IT developers, to FPGA and ASIC vendors, and to embedded specialists. Xilinx and Altera, among other FPGA vendors, say they see a mix of SMP and asynchronous RISC implementations.

Exascale expert says US efforts have failed

As multicore processing deals with physical and implementation walls, Exascale computing has hit both philosophical and fiscal walls. Thomas Sterling has been [delivering the depressing news that DARPA had given up](#) on the current path of computing architecture, which means funding goes away as well.

Sterling, an Indiana University professor and a former distinguished visiting fellow at Sandia Labs, says the efforts of the past year have been "a disappointment" and have created the need for radical experimentation outside the current realm of development.

"The disappointment has come from DARPA; the agency historically noted for leading HPC system research and development, Sterling said. "OHPC was an explicitly Exascale technologies-related program that would have augmented the principal UHPC projects. OHPC was cancelled shortly after the awards were made and there appears little likelihood that UHPC will extend beyond its first phase although it was originally planned for four such phases to produce final proof-of-concept experimental platforms."

Sterling blames the "relatively narrow view exerted as to how the problem of Exascale software should be addressed with strong resistance to considering revolutionary approaches." He said those advocating radical change have been given short leashes while conventional practices are given preference. "This intransigence has spilled over to the US debate as a whole, slowing down progress in needed advances."

Sterling points out that more radical approaches in computing architecture are getting better play in arenas outside the US.

Some ARM licensees, including Freescale Semiconductor, Texas Instruments Inc., Qualcomm Inc., and Broadcom Corp., utilize ARM as part of non-SMP designs that use a central control-plane processing environment, in conjunction with on-chip coprocessors for functions such as encryption,

deep packet inspection, and fast list searches for tasks such as routing. Performance numbers for the asynchronous world, realized in benchmarks such as SPECmarks and EEMBC, are harder to define than simple SMP benchmarks, but the throughput as realized in packet-processing, etc., seems to be better than the simple four-way or eight-way SMP sought in simpler server architectures.

Jim Ready, CTO of MontaVista Software, has seen multiple cores develop for single-chip processing from a perspective vastly wider than the SMP cores that characterize high-performance server computing. Ready was founder of Hunter & Ready, Ready Systems, and MontaVista - all companies that addressed both high-performance computing and embedded computing.

Last year, MontaVista was acquired by Cavium, a security-processor company that implemented some MIPS cores in traditional parallel domains, but also uses a variety of on-chip coprocessors that do not operate from the same clock as the primary control processor on chip. Ready has studied Patterson's paper, as well as the seminal paper of Phil Colella that describes the "Seven Dwarves" of computing algorithms (cited by Patterson), and he thinks he's taking a balanced approach for the Linux community.

"If you're a commercial Linux vendor like a RedHat, you really don't have much interest other than parallelizing integer throughput to the maximum extent possible. The SMP model is sort of automatically preferred as the baseline. So the server sector of the market sees nothing wrong with dumping as many thread processes as possible into the kernel," Ready says.

"Where we at MontaVista, and to some extent at Wind River, diverge from the server community, is that we see more value in accelerating functions like security and list search. So we try to keep the kernel as simple as possible, and optimize the software for the task at hand. More software advantages can be realized for optimizing for many cores performing many types of activities on one chip, than for multiple cores on a chip all operating in lockstep."

In Patterson's 'View from Berkeley' paper, Colella's original list of 'Seven Dwarves' was expanded to at least a dozen radically different computing models, and Ready says there may be as many as "15 distinct dwarves by now." Some linear algebra computing models can take advantage of SMP multicores, he says, but newer models like finite state machines, identified by Patterson, simply are too different to be tackled by SMP architectures in an optimal way.

"This is why Patterson made such a strong distinction between the high-performance computing community and the embedded community," Ready says. "There is no single right answer as to how best to implement parallel cores, and there may not be for quite some time. In fact, the methods seem to be diverging."

At Ready's parent company, Cavium, new generations of complex processors like Octeon may implement up to 48 MIPS cores on one chip. But of equal importance to the central RISC cores are the co-processing Deep Packet Inspection engines, up to 64 on a chip, and Neuron search processors, similar to ternary CAMs. The coprocessors and MIPS cores all are connected through an Octeon interconnect fabric. Since different tasks are passed to different processor types at all times, memory becomes less an absolute bottleneck than it is for traditional SMP integer platforms. The Cavium model is used to some extent by controller specialists like Freescale in PowerQUICC, i.MX, and Q-or-IQ families, and by FPGA vendors, particularly in Xilinx's ARM-heavy Zynq-7000 FPGA.

In more traditional server applications, developers are trying to re-define the role of central memory. Intel Corp. and Micron Technology Inc. have formed the Hybrid Memory Cube Consortium,

in which stacked 3D DRAMs with vastly accelerated I/O channels will be used as a means of leaping beyond DDR3, to address what Patterson called the memory brick wall. In March, the [ISQED Symposium](#) in Santa Clara will kick off the event with a day-long series of tutorials focused on making HMC technology work.

This is the "next best thing" on the horizon, according to Dean Klein, CTO at Micron. "This solution opens the door to lower power, higher bandwidth, lower latency. The DRAM die can be optimized for the DRAM process and for the logic die to be optimized for the desired logic process. There is a trend towards 3D integration with through-silicon-vias that will allow a closer coupling of CPU to memory in a wide variety of applications. For instance, the HMC is a great solution for servers and HPC, but at the other end of the spectrum, Wide-IO is gaining steam for mobile applications."

But Russell Fish, founder of startup Venray Technology, believes Intel "is only trying to push the I/O pins faster, when what is needed is a radical architectural shift." Fish called the Intel/Micron HMC "a Rambus redux, and little more."

Enter Venray's TOMI Borealis, a four-inch-square circuit board that combines 128 cores (unique to Venray) and 2 Gbytes of DRAM. The global optimization model used by Venray is ideal for solving "big data" problems in medical analysis, financial analysis, and global search through vast data sets.

"I actually have to go on record as saying that, at some time, this (TOMI) would be the way to go," Patterson concurs. "I wrote an article for Scientific American in 1995 about what microprocessors would look like in 2020 and I said there would be a merging of processors and memory called 'intelligent RAM.' At that time we actually started looking into the idea and [did several talks on the idea](#). Eleven years later it's sounding even better."

ARM's Goodacre weighs in as well. "ARM's partners have been integrating ARM processors close to different forms of memory for a long time. Whether this is on the same silicon, or on some package in package/3D arrangement, the benefits of putting memory closer to the compute is good for both power consumption and performance."

Micron's Klein has a different opinion.

"While I applaud every attempt to integrate processors and memory, I am not holding my breath any longer waiting for it to happen. Dr. Patterson is correct in saying that the idea's time may have come, from a computer architecture standpoint, however, there is nothing about the commercial DRAM process technology that makes this one inch closer to reality than it was ten or 20 years ago. At some point will the CPU-memory wall become so insurmountable that some brave soul will finally say: 'Damn the economics, full speed ahead!' I won't hold my breath waiting for that to happen, either."

In truth, the TOMI problem may be one of getting to market. Fish, a veteran of Motorola, Fairchild Semiconductor, and other companies, has designed award-winning architectures, such as the Sh-Boom, co-designed with Chuck Moore, but he and Moore have designed many radical instruction set architectures that have not gained traction with major semiconductor players. Nevertheless, the 22,000-transistor CPU/DRAM core defined by Venray may be simple enough to allow TOMI to gain ground.

In 2007, Patterson did a speculative piece for IEEE Spectrum, in which he said that in an ideal world, better parallel software and more efficient multithreading middleware would combine with lower-power CMOS processes to allow SMP multicores to move from 16 or 32 cores to dozens or

even hundreds of cores operating in lockstep. But the more likely scenario, Patterson added, was one in which a variety of approaches - new memory architectures, asynchronous cores, and software utilities outside the kernel - all helped in solving little chunks of the multicore problem, albeit in a haphazard and ad hoc way.

In the new realms of multicore outside traditional SMP server models, Patterson may have gotten it right in his latter prediction. There are many good ideas for specific vertical markets, along with a good deal of making things up as we go along.

Related articles:

[The future of computers - Part 1: Multicore and the Memory Wall](#)

[Future of computers - Part 2: The Power Wall](#)

[Future of computing - Part 3: The ILP Wall and pipelines](#)