

Select Font Size: [A](#) [A](#) [A](#)

Sponsored By



Multicore Is Bad News For Supercomputers


By **Samuel K. Moore**

IMAGE: SANDIA

TROUBLE AHEAD: More cores per chip will slow some programs [red] unless there's a big boost in memory bandwidth [yellow]

With no other way to improve the performance of processors further, chip makers have staked their future on putting more and more processor cores on the same chip. Engineers at Sandia National Laboratories, in New Mexico, have simulated future high-performance computers containing the 8-core, 16-core, and 32-core microprocessors that chip makers say are the future of the industry. The results are distressing. Because of limited memory bandwidth and memory-management schemes that are poorly suited to supercomputers, the performance of these machines would level off or even decline with more cores. The performance is especially bad for informatics applications—data-intensive programs that are increasingly crucial to the labs' national security function.

High-performance computing has historically focused on solving differential equations describing physical systems, such as Earth's atmosphere or a hydrogen bomb's fission trigger. These

systems lend themselves to being divided up into grids, so the physical system can, to a degree, be mapped to the physical location of processors or processor cores, thus minimizing delays in moving data.

But an increasing number of important science and engineering problems—not to mention national security problems—are of a different sort. These fall under the general category of informatics and include calculating what happens to a transportation network during a natural disaster and searching for patterns that predict terrorist attacks or nuclear proliferation failures. These operations often require sifting through enormous databases of information.

For informatics, more cores doesn't mean better performance [see red line in "Trouble Ahead"], according to Sandia's simulation. "After about 8 cores, there's no improvement," says James Peery, director of computation, computers, information, and mathematics at Sandia. "At 16 cores, it looks like 2." Over the past year, the Sandia team has discussed the results widely with chip makers, supercomputer designers, and users of high-performance computers. Unless computer architects find a solution, Peery and others expect that supercomputer programmers will either turn off the extra cores or use them for something ancillary to the main problem.

At the heart of the trouble is the so-called memory wall—the growing disparity between how fast a CPU can operate on data and how fast it can get the data it needs. Although the number of cores per processor is increasing, the number of connections from the chip to the rest of the computer is not. So keeping all the cores fed with data is a problem. In informatics applications, the problem is worse, explains Richard C. Murphy, a senior member of the technical staff at Sandia, because there is no physical relationship between what a processor may be working on and where the next set of data it needs may reside. Instead of being in the cache of the core next door, the data may be on a DRAM chip in a rack 20 meters away and need to leave the chip, pass through one or more routers and optical fibers, and find its way onto the processor.

In an effort to get things back on track, this year the U.S. Department of Energy formed the Institute for Advanced Architectures and Algorithms. Located at Sandia and at Oak Ridge National Laboratory, in Tennessee, the institute's work will be to figure out what high-performance computer architectures will be needed five to 10 years from now and help steer the industry in that direction.

“The key to solving this bottleneck is tighter, and maybe smarter, integration of memory and processors,” says Peery. For its part, Sandia is exploring the impact of stacking memory chips atop processors to improve memory bandwidth.

The results, in simulation at least, are promising [see yellow line in “Trouble Ahead




PHOTO: INTEL

THE FUTURE: Intel's experimental chip has 80 cores.

To Probe Further

The complete Sandia simulation is still in progress and will be published later when additional computing benchmarks are finished. (Each benchmark takes about one month to produce.) However, Richard C. Murphy [explained the main results we reported on at the IEEE International Symposium on Workload Characterization 2007](#), in Boston, Mass. [abstract view in IEEE Xplore; subscription required for full text

And he and University of Notre Dame's [Peter Kogge](#) analyzed the [memory access patterns of informatics applications](#) in an article in the July 2007 issue of *IEEE Transactions on Computers* [abstract view in IEEE Xplore; subscription required for full text

