

**EETIMES**
ONLINE

EE Times:

High and dry at high end?

Rick Merritt and Nicolas Mokhoff

(11/15/2004 10:00 AM EST)

URL: <http://www.eetimes.com/showArticle.jhtml?articleID=52601292>

Pittsburgh — There was plenty to cheer about at Supercomputing 2004 here last week, as both custom and off-the-shelf designs took major steps forward in performance and market reach.

But a team of top computer experts warned that the pipeline for fresh ideas in computer science is nearly dry, and urged the U.S. government to make big changes in its procurement and research activities to prime the pump. Other industry watchers voiced concern about nagging problems in supercomputer software and benchmarking.

"Our top message is that when we look to the future, we cannot figure out where our supply of [high-end supercomputers] is going to come from," said Susan L. Graham, computer science professor at the University of California, Berkeley, and co-chair of the National Research Council's "Getting up to Speed with the Future of Supercomputing" report.

Today the government is relying too heavily on relatively low-cost clusters of commodity microprocessors, and each agency buys systems independent of any overarching plan. But slowing gains in single-thread microprocessors and growing memory and network latencies will undercut the benefits of these off-the-shelf systems in the future, the report said.

That slowdown will lead to a generation of subpar supers, with nothing in the pipeline to take their place, threatening the needs of top government and scientific researchers, according to the report.

"We are extrapolating that by 2020, a computer node can execute a million instructions in the time it takes to communicate with another node," said Marc Snir, head of the computer science department at the University of Illinois at Urbana-Champaign and co-chair of the report. "That's not tenable," he said. "We don't know how to write algorithms for such machines."

The two-year report commissioned by the Department of Energy urged the government to draft — with broad support from industry and academia — a road map for supercomputer technologies and needs for the next five to 10 years, then procure systems that follow that road map. It should also spend an additional \$140 million a year on hardware and software research into custom supercomputer designs, the study said.

Under current government procurement practices, Cray is the only U.S. company developing custom processors specifically for supercomputers, and no companies exist in some key supercomputing-software sectors, said Snir. "Supercomputer companies not focused on commodity microprocessors have found the market is too small to build a company," he said.

The committee that prepared the report is now presenting its findings to leaders in Congress and the White House. However, the war in Iraq and record budget deficits cast a cloud over their message. "The chance that budgets for science will see major increases are not as good as we would like, but we can only make the best case we can for them," Snir said.

"I'm optimistic that we're laying the foundation now for good things in this area in the future," said Alan Laub, a Department of Energy program director and a professor at the University of California, Davis. Laub co-chaired the High-End Computing Revitalization Task Force that released a similar report in May.

Glacial speed

Problems in software and benchmarks also dog the supercomputing community, said Charles J. Holland, deputy undersecretary of defense for science and technology.

Progress in software has been moving at "glacial speed" compared with hardware, said Holland in an invited talk here. End-to-end system latencies have been poor and "much of [today's] code cannot utilize an entire [high-performance system] effectively," he said.

Meanwhile, code size is growing. By 2010, it will take nearly 20 million lines of code to control avionics in fighter aircraft.

"We now have a multiagency effort to use common analytical methods and tools to map different applications to appropriate hardware and software architectures," Holland said.

As for benchmarks, the Department of Defense uses High Performance Linpack to rate supercomputers because it measures not only speed but also accuracy of calculations within limited performance parameters. That's superior to the straight Linpack method currently used for the popular Top 500 list that ranks the world's largest installed supers, Holland said.

Beyond the DOD's metrics, the government is spurring collaborative development of a so-called HPC challenge (see <http://icl.cs.utk.edu/hpcc/index.html>), a suite of seven benchmarks that will attempt to provide a measure of real-world applications performance.

The Top 500 list got an update at last week's conference, with IBM Corp.'s BlueGene/L making its debut on the list as the most powerful installed supercomputer on the planet. Its performance of 70.72 teraflops showed at least one successful custom system emerging from the labs. Columbia, a cluster of Intel Itanium-based systems built by SGI for NASA, took second place at 51.87 Tflops (see www.top500.org).

Both systems topped the Earth Simulator, which is now ranked third at 35.86 Tflops. Built by NEC Corp. and installed in 2002 in Yokohama, Japan, the Earth Simulator had held the No. 1 position for five consecutive editions of the list, which comes out twice a year.

Clusters of off-the-shelf CPUs like those used in SGI's Columbia have increasingly dominated the Top 500 list in recent years. While that trend continues in the latest list, the BlueGene/L shows custom architectures are making a comeback among the most muscular systems.

"In the last few years there's been a growing interest in new computer architectures at the very high end and BlueGene is one of them," said Erich Strohmaier, a computer scientist with Lawrence Berkeley Labs who helps maintain the Top 500 list. "I'd expect to see more custom architectures on the list in the next few years."

Custom computing

Indeed, Cray, IBM and Sun Microsystems are competing to deliver by 2009 a custom supercomputer under a program of the Defense Advanced Research Projects Agency (see www.eet.com/showArticle.jhtml?articleID=18308889). In addition, NEC has announced a follow-on to the system used in the Earth Simulator (see www.eet.com/showArticle.jhtml?articleID=50900052) and Cray is developing a hybrid system using custom interconnects and commodity Opteron CPUs from Advanced Micro Devices Inc. (see www.eet.com/showArticle.jhtml?articleID=19502121).

IBM's BlueGene/L has plenty of headroom. Only 32,000 of its custom PowerPC 400 processors have been installed so far. Ultimately the system will sport 128,000 CPUs when it is delivered to Lawrence Livermore National Laboratory in Livermore, Calif.

IBM announced it will make a commercial version of BlueGene available using off-the-shelf Power5 processors. It will deliver peak performance of 5.7 Tflops for a rack of 1,024 dual-processor nodes that fit into a 3 x 3 x 6-foot floor space. Up to 64 racks can be linked together, IBM said.

Meanwhile, off-the-shelf clusters are making great strides on the Top 500 list. There are 296 systems so labeled on the list, making clusters the most common architecture in the Top 500. That's up from 208 clusters at the end of last year and 149 on the June 2003 list.

A total of 320 systems on the list are now using Intel processors, compared with 287 six months ago and 189 one year ago. The second most commonly used CPUs are the IBM Power processors (54 systems), Hewlett-Packard's PA RISC processors (48) and AMD processors (31).

While not at the No. 1 spot, clusters are contributing to great performance advances among the Top 500 systems. Total combined performance of all 500 computers has for the first time exceeded the 1-petaflops mark with the list released a week ago. The exact figure is 1.127 Pflops, against a total 813 Tflops six months ago.

Only one system in the top 10 now ranks below 10 Tflops. Strohmaier of Lawrence Berkeley Labs said he expects the next version of the list, to be released in six months, will only include systems exceeding 1 Tflops.

While off-the shelf architectures are moving upward in performance, they are also stretching outward in the number of users and applications they embrace. Indeed, one researcher here spoke of a democratization of supercomputers.

In an invited talk, Stan Ahalt, director of the Ohio Supercomputer Center, pointed to two new high-performance systems shown at the conference: the \$5,500 Rocketcalc and a \$10,000 Orion Multisystems computer. Both plug into regular ac sockets without special wiring.

Companies such as Procter & Gamble are making investments in high-performance systems. At the other end of the scale, it is possible now for a dentist using a desktop high-performance computer in his office to analyze the form of a crown for a tooth and sculpt that crown within 20 minutes, in a regular office visit, Ahalt said.

Ahalt cited a recent survey report from the Council of Competitiveness (Washington) in which more than 70 percent of respondents indicated their companies could not function without a high-performance computer, and 25 percent could quantify its return on investment.

Seeing this high-end market broaden, Microsoft Corp. has reportedly laid plans to beta test in 2005 a high-performance version of its Windows Server operating system for 64-bit Intel Xeon and AMD Opteron CPUs (see www.informationweek.com/story/showArticle.jhtml?articleID=51201325). It will support a version of the Message Passing Interface software that has become one of the key software components of off-the-shelf clusters.

All materials on this site [Copyright © 2008 TechInsights, a Division of United Business Media LLC](#). All rights reserved.

[Privacy Statement](#) | [Your California Privacy Rights](#) | [Terms of Service](#) | [About](#)

