In University Supercomputing, the Fastest May No Longer Be the Best - Technology

By Jeffrey R.
Young

Kristen Schmid Schurter for The Chronicle

For Blue Waters, a supercomputer to be housed here at the U. of Illinois at Urbana-Champaign, quickness is not of the essence.

Champaign, Ill.

The warehouse-sized supercomputer under construction here at the University of Illinois at Urbana-Champaign comes with a price tag of nearly half a billion dollars, making it one of the most expensive supercomputers ever devoted to academic research. And yet, when engineers turn on the machine this year, it very likely won't be the fastest computer in the world.

And its designers don't care.

"We're not looking to be on the Top 500 list," says Thom Dunning, who leads the computer's development as head of the university's National Center for Supercomputing Applications. Rather than hit a peak sprint speed measured by the Top 500, the most widely used supercomputer ranking, he wants to build a distance runner, capable, for example, of powering through intricate simulations of a tornado that can predict where a storm might strike.

Flat-out speed, for a long time the measure of a supercomputer's worth, may be going out of style. A recent report from an influential federal panel recommended more emphasis on software and alternative designs rather than computational Ferraris. Still, fast computers attract top faculty—and federal money. "Every congressman loves to sign his name to the latest, greatest machine," Mr. Dunning acknowledges. "That's the photo op. You don't get the same photo ops with software."

The 2012 federal budget request for high-end computing (including infrastructure, research, and development) is

about $1.6-billion. So the heightened debate over the need for speed could shape how this money gets divided up—and have a major influence on the universities that house supercomputers.

Like Mr. Dunning, some leaders now argue that a single test of top computing speed often doesn't reward clever software designs—and that software is increasingly the bottleneck that slows simulations that might otherwise lead to scientific breakthroughs, such as understanding that tornado or the intricacies of a biological cell, which are two of the tasks that the new machine here, called Blue Waters, will undertake.

That argument was made strongly in the federal report, issued in December by the President's Council of Advisors on Science and Technology. It calls for a more balanced portfolio of U.S. supercomputing development, and warns against an overemphasis on speed rankings like the Top 500 list. "Engaging in such an 'arms race' could be very costly, and could divert resources away from basic research aimed at developing the fundamentally new approaches" to supercomputing, the report states.

Think, say advocates, of the folly of a best-car list based only on top speed. So what if a Ferrari is faster than a Volvo station wagon when you have to take two kids to soccer practice?

But as any college leader knows, rankings are hard to ignore. Just last week, when Yale University unveiled a high-performance computer, Steven M. Girvin, deputy provost for science and technology, called it "among the top 100 academic machines in North America."

Even President Obama referenced the supercomputing speed race in this year's State of the Union address, noting that China, not the United States, now boasts the world's fastest computer. It's the first time a Chinese machine has been No. 1 on the Top 500 ranking, calling policy makers' attention even more to the race for speed.

So the pressure is on Blue Waters, which floats on federal money—the National Science Foundation has contributed $208-million over several years to build it. Mr. Dunning admits that the new computer has to lead to some clear scientific breakthrough. In fact, this center has made big news before. It's where the first popular Web browser, Mosaic, was invented as a tool to help scientists here manage their data. That advance ended up sparking an information and e-commerce revolution. If the new supercomputer can come up with something like that, nobody will care if it doesn't top a speed list.

Touring the Machine

Because Blue Waters is the kind of supercomputer you could use to design nuclear weapons, the building here in Champaign has plenty of spy-movie security measures not usually found on college campuses. Visitors must submit to a retina scan, for instance, then enter a holding chamber that checks their weight to make sure no unapproved freeloader is tagging along.

A football-field-sized computer room on the second floor will soon contain 300,000 processors operating in unison. The closely packed computer chips will generate so much heat that water will be continually streamed through cooling pipes (hence the machine's name). IBM produced the processors, which are so new that their specifications remain a trade secret. (The researchers all signed nondisclosure agreements.)

In the supercomputing world, getting a new computer is vastly more complicated than taking a machine out of a box and pressing the "on" switch. The entire design is a research project. Much attention is paid to the "interconnects" that link the chips: It's no use having superfast processors if the instructions that coordinate their activity are delivered inefficiently.

Blue Waters also takes a unique approach to data storage. Borrowing from the design of the human brain, some memory actually lies along the pathways between processors. Even traveling at the speed of light along fiber-optic cables, data moving from the large tape-storage library, elsewhere in the building, to the processors can cause split-

second delays that can degrade performance.

This computer is one of the first whose brain is arguably faster than ours, at least "by some metrics that are very rudimentary and very debatable," says Marc Snir, a computer-science professor at Urbana-Champaign who is a principal investigator on the Blue Waters project. (That claim may worry science-fiction fans. In the novel *2001: A Space Odyssey*, the murderous computer HAL was invented in Urbana. But Mr. Snir insists that fact won't follow fiction: This computer will have no possibility of developing a will of its own.)

It didn't used to take all this effort to make a supercomputer faster. For decades, improvements in individual chips brought speed gains without the cumbersome interconnects. But it has become impractical to make individual chips faster using today's technologies; hence the move to parallel systems with hundreds of thousands of processors pulsing together.

Adding so many chips brings expensive challenges, like providing enough power for the system, cooling the machine, and building the software that can manage dividing a simulation among so many processors, or cores.

"As we move into tens of millions and hundreds of millions of cores, we'll have to think differently about how we program the machines," Mr. Snir says. "The software may look very different."

The heyday of computers like Blue Waters—those as big as football fields—may soon be over, some scientists argue. Most machines to come may use a system that works more like a search engine, distributing computational problems among processors distributed across a broad physical network, often called a cloud.

That's the focus of Dan Reed, vice president for technology, policy, and strategy at Microsoft Research, who calls this research area the next generation of cloud computing.

"If you look at the insights we have gleaned from massive search engines and cloud infrastructure, those things are really bigger than any of the supercomputers that have been built before," he says.

Proponents of single-location machines like Blue Waters argue that the cloud will simply never be able to crunch numbers as fast as their systems, because of the time it takes data to move across a network.

And even fans of the cloud approach applaud building at least a few machines like Blue Waters, which many compare to the Hubble Space Telescope—an expensive instrument reserved for the most challenging problems, but not the tool that most scientists use.

The real question for research computing, then, is where to invest the most energy and money. If supercomputers are like cars, which ones belong in the national garage?

Measuring Speed

It's not clear how well Blue Waters will do on the Top 500 list. That will depend in part on whether or not other countries unveil even faster machines at about the same time—as Japan and possibly China may do.

But perhaps it shouldn't matter if Blue Waters tops the list, says Jack Dongarra, director of the Innovative Computing Laboratory at the University of Tennessee at Knoxville. He should know. He's one of the creators of the ranking.

"I criticize this list as well," he says, because it reflects only how fast a computer solves a series of algebraic equations. "These computers are complicated, and they have many facets, and we should evaluate the different components that go into the systems."

He recently helped create a new speed test, one that considers many variables, called the HPC challenge benchmark. It measures things like how fast data can be stored, retrieved, and moved within the computer. But the

test is not as popular as the method used by the Top 500 list, although the NSF and other U.S. agencies have used it, he says. It doesn't produce a ranking, since there is no way to fairly reduce the variables to one number for comparative purposes—in the same way *Consumer Reports* doesn't produce a single ranking for all cars, he says. "I wouldn't want to make policy based on one number."

In November, still another supercomputer ranking was unveiled, at a conference in New Orleans. This one, called Graph500, does produce a ranking, but it is based on how fast supercomputers solve complex problems related to randomly generated graphs, rather than on the simpler computation of the Top 500. Some computers that had ranked well on the Top 500 ran the Graph500, but their operators refused to announce the scores, most likely because they fared less well.

Does that mean China may not actually be ahead in the supercomputing race? Hard to tell, unless its computers participate in alternative challenges to Top 500. One scientist here speculated that the Chinese computer may have been designed simply to do well on that one test.

Mr. Dongarra says he saw the Chinese machine when he visited China's National Supercomputing Center, about two hours' drive from Beijing. "It's state-of-the-art in many ways," he says, noting that he was impressed by the unique interconnections among processors that researchers there had developed.

Mr. Reed, of Microsoft Research, says he, too, has been impressed by China's efforts. "I used to say that in the high-performance computing race, the U.S. was laps ahead," he says. "Now it's steps ahead."

No Guarantees

Why do researchers need such fast machines, anyway?

Klaus Schulten, a physics professor at Urbana-Champaign and a leader in exploring how biological cells function, is one who says he needs a machine that can perform quickly enough, and consistently enough, to run computer models so complex that they would take previous machines years to calculate. Now the same simulation should run in a couple of months. He began his work at top German universities and says he was drawn to Illinois because of the university's supercomputers.

His plan for Blue Waters is to create a high-resolution simulation of a key part of a living cell, called an organelle, on a scale that has never before been attempted. "A living cell has as many proteins as the United States has citizens," he says. "For the last 50 years, we were investigating the individual members of the society, the individual proteins. Now we want to go to the society."

"When I came to America and I saw my first football game, I didn't know what was going on. I just found it beautiful—the people and the colors were nice and the girls dancing, that was nice. But that was about the most I got out of it," Mr. Schulten says. That's a good metaphor for scientists's minimal understanding of molecular behavior when they watch cells under the microscope, he says. "We have no clue how the cells do things. Now we have to find out the rules of the game."

He hopes the computer simulation will do that—but there's no guarantee that it will yield a breakthrough. In case it doesn't, he hopes to make an even more high-definition model for the next generation of supercomputers, whatever form they take.

For much of his career, Mr. Schulten has been able to stay in front of peers doing similar research, he says, because of his focus on getting time on advanced supercomputers. He showed a visiting reporter a parallel supercomputer he and his colleagues built themselves years ago, which is now in a Plexiglas case outside his office, underneath a colorful printout of the cell model the machine devised.

The professor has also had to build his own software, since there wasn't any available when he started. Now the

software his team built has become a standard in biophysics, with more than 200,000 registered users, he says.
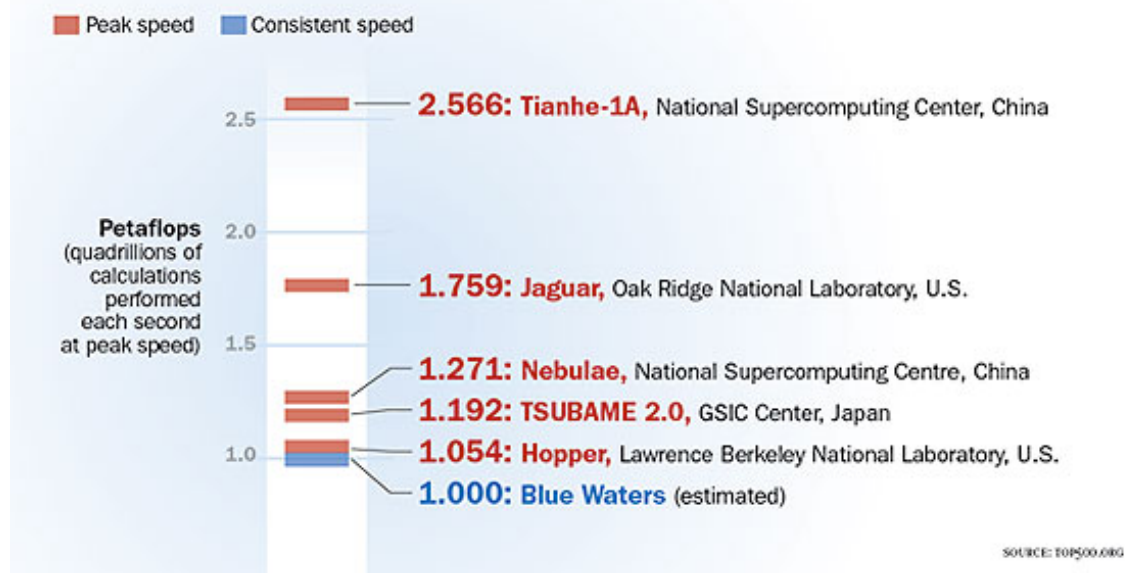
But getting attention—and, most important, grant support—for software can be tough, many computer scientists say.

The biggest help in that PR battle may have come from an unlikely quarter—a TV game show. In February *Jeopardy!* ran two episodes in which a supercomputer named Watson faced off against two human champions. The computing problem boiled down to a clever algorithm as much as to speed, since Watson had to understand puns and other language oddities in the questions so it could hunt down the answers.

And, at least for now, it inspired confidence in a new approach to supercomputing research. Watson won handily.

---

6 Supercomputing Speedsters

The Top 500 Supercomputing Sites ranks the world's fastest machines at their peak speeds, measured in petaflops. Here are the top five, plus a new entry from the University of Illinois at Urbana-Champaign called Blue Waters. It has yet to be measured at peak speed, and designers think consistent speed during scientific simulations, shown below as an estimate, is more important.

■ Peak speed   ■ Consistent speed

Petaflops (quadrillions of calculations performed each second at peak speed)

- 2.5 — **2.566: Tianhe-1A,** National Supercomputing Center, China
- 2.0
- 1.759: **Jaguar,** Oak Ridge National Laboratory, U.S.
- 1.5
- 1.271: **Nebulae,** National Supercomputing Centre, China
- 1.192: **TSUBAME 2.0,** GSIC Center, Japan
- 1.0 — 1.054: **Hopper,** Lawrence Berkeley National Laboratory, U.S.
- 1.000: **Blue Waters** (estimated)

SOURCE: TOP500.ORG