


# PREDICTING POWER

## What's Ahead For High-Performance Computing




**S**ometimes it seems that personal computer breakthroughs are an almost daily occurrence. Consumers have grown accustomed to hearing of this performance enhancement or that critical improvement, responding to needs they didn't know they had. Once bought, PCs can appear obsolete, quickly superseded by next week's model.

Such is not the case for supercomputers. Despite hardware innovations and creation of specialized software, the performance breakthroughs can't come rapidly enough. Supercomputers cannot answer any number of scientific and technological questions.

Nature's complex interactions require enormous calculating capability, which defies even the latest generation of machine. Until, possibly, now.

**QUEST** editor James Schultz spoke with David Keyes, director of Old Dominion's Center for Computational Science and acting director for the Lawrence Livermore Laboratory's Institute for Scientific Computing Research, to get an idea of what the next decade may hold for what insiders call "high-performance" computation.

Keyes' comments have been edited for length and clarity.



## Where do we stand with large-scale computing?

Computational simulation, as a way of forecasting and policy support, is making its way into the federal arena, in terms of procurement and policy. We're on the verge of having computational simulation taken seriously for the first time. In the scientific computing arena, important decisions are being made regarding carbon-dioxide emissions or whether we can forgo nuclear weapons testing, using simulations.

We're continuing to discover we don't have enough range of scales on the computer to match the range of scales that exist in the real world. At the moment, for instance, a really good global climate model might have a [resolution] of one data point per square

kilometer of the Earth's surface, which is not enough to resolve clouds and other surface phenomena known to have effects on local weather. There have been some exciting successes in forecasting thanks to modeling to fill in the missing scales, though.

Wildland fires — like the ones in Los Alamos and the recent ones in Australia — can cause widespread danger and damage to property and health. These fires can pile up huge costs, from the property they damage and the costs incurred in fighting them. Several groups are now addressing, with supercomputers, how these fires evolve. The more detailed models may take weeks to figure out what would happen in the next hour, so that's not too useful to fire chiefs. The cruder models are on fire chiefs' laptops, which is something, at least.

September 11 has had an effect. A lot of new simulation pro-

jects have been jump-started, accelerated to see if they can be of help in disaster preparedness. The expectations tend to run ahead of the scientific capability, at least for now. It's still an exciting area to be in. I definitely feel I was born for just this moment. The motivation and the technology are coming together, even though the number of people with the interdisciplinary background to take advantage of [the increasing capabilities] is relatively small.

It's like the Manhattan Project or Apollo. There's an urgency to acquire a certain technological capability as soon as possible. Take the three "dash oh" prefixes — bio, nano, info — driving federal procurements. Just in the Department of Energy, nano[technology] and bio[technology] are being kick-started in a vigorous new way. Some decisions are being made at the very edge of what we know or can forecast.

## Are computers becoming more powerful?

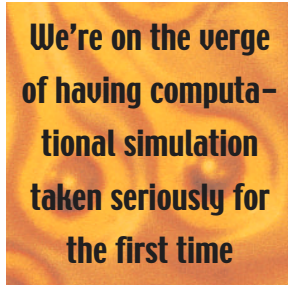
They are. [But] it takes big jumps to keep up with three dimensions and time. Fab[rication] lines are \$3 billion a pop. It's a very risky thing to develop a new generation of microprocessor. That technological capability is very rapidly soaked up by three-dimensional simulations, if you don't improve the models to put the memory storage exactly where the problem is.

If you want to do something as modest as double the mesh resolution in one dimension, then you need a factor of eight to double it in all three dimensions. If you want to improve resolution by a factor of 10 in each dimension, then you need much more [microprocessor capability], which can take the better part of a decade to achieve. And there are many phenomena for which the next range of factor-of-10 spatial resolution is not enough. Take turbulence, for example: the range of spatial scales is huge, from the size of an eddy shed off an airplane wing or, even worse, some climatological feature down to the scale at which [molecular] motion is dissipated into heat. We can't pretend the models we're using now are capable of resolutions down to the smallest scales. We need bridging models to go between the scales.

One of the trends in the leading edge of simulation is getting material properties from fundamental molecular and atomic interactions. You plug those properties into the models you have that average over the continuum. There's a real need for scientists who can look over the wall and understand the computational technology, to see what realistically can be computed. You have the hammer and so you go looking for the right nail.

When I talk to colleagues who span a couple of generations, there's a sense that this is the time that we've all

been waiting for. We can see the chance to move beyond the limitations of physical experimentation to computationally designed and computationally investigated fundamental science ... There are things you can do that you could never do as a physical experiment.



**We're on the verge  
of having computa-  
tional simulation  
taken seriously for  
the first time**

## Are we then moving into an era of super supercomputers?

Around 1995 or so, Thomas Sterling of Cal Tech (who happens to be an Old Dominion University alumnus) and JPL [NASA's Jet Propulsion Laboratory, in Pasadena, Calif.] laid out a plan to get to petaflops computing [thousands of trillions of mathematical calculations per second] which depended on a specific architecture. We're at teraflops now [trillions of calculations per second], but it's still pretty expensive. There are maybe 15 teraflop-capable systems in the world that are tightly coupled, i.e. confined within the same basketball court-sized room, as opposed to loosely spread out across the Web.

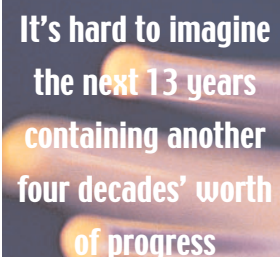
At the other extreme, the National Science Foundation has a teraflop facility that spans between San Diego and Illinois, also linking in two supercomputer centers at Cal Tech and Argonne: actually 13.6 teraflops of potentially poolable processing power, and the memory and the disk space that go with it. The majority of interesting problems, though, can't tolerate even the small amount of telecommunications latency and synchronization problems you get when you try to

compute across a large geography.

It's still very expensive to do a single, sustained teraflops-scale computation, because the interconnections aren't commoditized [so that they're cheaper]. The processors are commodities, but the interconnects are custom. The software is custom. A lot of people have to write their own distributed operating system, graphical interface and debugging tools. The vendors are producing these systems "on a shoestring;" they're not going to sell many copies.

Terascale is still a struggle [and] most scientists cannot afford it. If you're at a teraflop and want to get to a petaflop, it's a factor of 1,000. You can get a factor-of-10 improvement with straightforward hardware improvements, but you still need that factor of 100, and the only way you can do that is with concurrency [interconnections].

Right now, the biggest teraflop system we have is less than 10,000 processors: 8,192, to be exact. If you're talking about going to another factor of 1,000 in processor capability, you're talking about systems with a million processors. This produces several challenges all at once. We're already having a hard time developing networks that are reliable enough to do computations distributed over that many processors. The simple statistics of mean time to failure suggest that, with 10,000 processors, you lose one node a couple of times a week. With a million processors, you might lose one every hour. You might even lose a second one while you're still doing your backup to guard against losing the first. One of the major challenges is in system software ... In the future, I think scientific users will have to construct their own codes to be more tolerant of losing processors and (with them) pieces of the overall data.



**It's hard to imagine  
the next 13 years  
containing another  
four decades' worth  
of progress**

## When will these problems be overcome such that we'll have petaflops computing capability?

By 2007 is a reasonable date. The Department of Energy was planning to get to 100 teraflops by 2004, which I now doubt they're going to be able to do. I'm fairly confident the technology to get to petaflops by 2007 will be there. I'm not sure it will be cheap enough for any one agency to make the plunge. The one agency most likely to do it is one that doesn't talk about what it does: the National

Security Agency. They've actually been funding [petaflops research] more than any of the other science agencies.

... One technology that is being developed to get to petaflops is called "processors in memory," PIM for short. IBM already has one such processor available and there are Japanese vendors that have another couple. The basic idea is to attack the prime weak link in today's architectures, which is the memory bandwidth: getting data into the processor proves to be the slowest part of the computation for these simulations.

When you bring the processor down into the memory system, it's not on a separate chip. It doesn't take a lot of steps to move data in wide parallel

slices ... The down side is that chip densities you can get doing that are a factor of six smaller. The single chip packs much less memory when you fab it in the same line as a processor. As a result, we lose memory per node, even as we gain much higher bandwidth and improve interprocessor communications.

... Fusion power is one area — and I'm not talking about the discredited desktop variety — that, if we're going to have it, should get its next scientific and engineering shot in the arm from high-speed computing. I'm very excited about that. Fusion is being seriously entertained, and we may see it because of these [modeling] simulations.

## So what's beyond petaflops computing?

We've experienced a factor of 10,000 in performance improvements in about 13 years. It's hard to imagine the next 13 years containing another factor of 10,000, at least based on the current architecture.

The frontier limits to silicon may be pushed out of the way before we hit them. Biocomputers and quantum computers are the two that are the most often named as the next generation ... One possible jump is to continue to use current [architectures], but instead of using silicon and copper wires, you use carbon nanotubes, which leads to much smaller devices, with power requirements per calculation that are very much smaller as well.

Beyond petaflops isn't automatic; it's high-risk and there's a dimly lit path



**Creative people do seem to find solutions**  
—David Keyes

in that direction. It's one of those situations into which you boot-strap yourself. You use computers to understand nanotechnology, and then you use nanotechnology to build the next generation of computer, and so on.

In my own brief career window, watching wild projections, I often find myself on the skeptical side waiting for things to materialize ... But creative people do seem to find solutions.

[Simulations] have been jump-started, accelerated to see if they can be of help in disaster preparedness