

washingtonpost.com

## Science: Supercomputers

Christopher Lee and David Turek  
Washington Post Staff Writer and Vice President of  
Supercomputing, IBM  
Monday, December 3, 2007; 12:00 PM

Washington Post staff writer **Christopher Lee** and **David Turek**, vice president of supercomputing for IBM, were online to discuss technological advancements in computer science, including the creation of the first "petascale" supercomputer.

Read more about this in today's story: [Faster Computers Accelerate Pace of Discovery](#), by Christopher Lee, and watch a [video about the Road Runner supercomputer](#) courtesy of IBM.

**The transcript follows.**

---

**Christopher Lee:** Welcome to today's chat everybody. Thanks for joining us. Supercomputers are becoming so powerful it's scary. Next year the computing industry will unveil a machine that can perform 1,000 trillion calculations per second. That's almost twice as powerful as today's machines. I don't know about you, but numbers like that make my head hurt. Fortunately, we've got an expert with us today to help us understand the significance of the advances science is making in supercomputing. He's Dave Turek, vice president of supercomputing for IBM, and we're lucky to be able to pick his brain for the next hour. These monster machines are vital tools for electronically modeling things like climate change, the effects of new drugs on the human body, and the likely implications of earthquakes. They've become so much a part of science that some experts argue they have changed the scientific method. Now, in addition to theory-hypothesis-experimentation we've got another step: Computation. What do you think about that? Or anything else involving these amazing machines? Let's get started.

---

**Reston, Va.:** What are the best steps for learning about supercomputing technology?

**David Turek:** One of the best ways to learn about supercomputing is to follow the news coming out of companies who follow the industry. Two very good examples are IDC (whom many consider to be the bible of the industry) and HPCWire (which has up to the minute news blurbs and lots of other information). You can also check out the web site for the US Council of Competitiveness which has a section on HPC (high performance computing) in layman's terms. Look at companies web sites (for IBM the keyword is DeepComputing) but check out Cray and SGI, HP, and Sun as well. The Office of Science for the Department of Energy has some interesting material as part of their INCITE program and look at the HPCS program that is managed by DARPA. A high level technical and historical view can actually be obtained from Wikipedia (caveat emptor on veracity of everything one finds on the internet). Also look up the PCAST report (Presidential Council of Advisors on Science and Technology). They came out with a report earlier this year.

Advertisement



**East Brunswick, N.J.:** What are potential applications of super computers in the field of medical science?

**David Turek:** Medical science is a rich area for supercomputer application. One of the key things happening in the medical field is the relentless adoption of digital technologies in place of the old familiar analog approaches. An example in my life is the fact that my dental Xrays are now all digital. Well, once you get into digital imaging clever people can start writing algorithms that can actually assist the physician in making better diagnoses of what the image shows. This doesn't replace the judgement of the physician but provides an enhanced tool to better assess the medical situation. Examples abound and include mammography, brain imaging, digital colonoscopies and so on. There is also opportunity in drug design and personalized medicine where a supercomputer is used to help decode the genetic profile of a patient so that a customized therapy could be provided yielding the highest possible probability of a successful outcome. This is a very vibrant area of research and I would expect an extraordinary amount of innovation in the next few years coming out of this area.

---

**Alexandria, Va.:** Why use a supercomputer instead of a cluster?

**David Turek:** One of the interesting expenditures of energy in our field is the debate among different factions over the difference between supercomputing and clusters. Rather than engender more smoke on that topic I would say it's about horses for courses, meaning one picks the architecture appropriate to the problem at hand. For years extraordinary work has been done in the petroleum industry with very large clusters using relatively low speed interconnects like ethernet. Then there are applications like the Blue Brain project at EPFL in Switzerland that have a need for something quite different. So for me the question starts with what you want to accomplish and then match the attributes of the application to the nature of the platforms available. I have no better dogma than that to provide

---

**Boston:** What is the biggest worldwide issue/problem that has yet to experience the benefit of supercomputing?

**David Turek:** One of the interesting aspects of answering your question is to reflect on all the ways supercomputing is used today. Supercomputing is very pervasive and we all encounter the benefits of this type of technology in our everyday lives. They are used to design products we use (efficient autos and airplanes), used to help discover petroleum deposits, used to make animated films, help design and discover new pharmaceuticals, route traffic through busy cities, as well as numerous applications in areas like homeland security. One of the things I would expect to see in the future is the progressive inclusion of more of the human element in the science simulations that take place today. In fact, there are some researchers who are beginning to use supercomputers to study issues in sociology. So inclusion of human behavior in some of the modeling activity might produce bigger and better returns than we have currently experienced

---

**San Antonio, Tex.:** Questions for Mr. Turek:

Are there any fundamental roadblocks to making ever super-er supercomputers that you can see? If so, when might we bump up against them?

And, what problems do we know about now that need multi-peta, exa-, zeta-, yotta-flop, etc. computation to address? Doubtless completely unexpected ones will pop up in the future, but which can we name with some plausibility now? String theory? Complete simulation of a bacterial cell? Or what?

**David Turek:** It's a never-ending journey for which technology problems continually arise. Obviously one of the key design issues today is the amount of energy the very high end systems require for both power and cooling. That has given rise to innovative designs like Blue Gene from IBM and the new machine that Sicortex is producing. But we also have problems with respect to I/O (how fast can you feed these things data) and emerging problems with respect to the ways in which memory needs to be packaged to keep these systems in computational balance. As an aside, the cost of memory is becoming a very dominant feature of future system designs. We also have issues with innovation (or the lack thereof) on new algorithms to utilize these new hardware designs.

In terms of new problems that require this kind of computing (and the orders of magnitude increases in compute power that we expect) I would expect that some will appear to be extensions to applications familiar today (materials science, weather forecasting, industrial design, logistics, financial modeling, digital media, petroleum exploration, etc) but pursued to much greater fidelity and radically more inclusive of elaborate context. For example, linkages of earth, atmospheric, and ocean models on a world wide basis to get a better feel of the interaction of industrial activity on weather patterns. But some of the applications will have a peculiarly pedestrian feel. For example, one of the most complex computational efforts being contemplated today has to do with modeling the attributes of plumbing in nuclear power plants. This is unusual plumbing subjected to unusual phenomenon but plumbing nevertheless. These kinds of applications need multi-petaflops today. And there is a wide ranging array of so-called multiscale or multiphysics problems that require huge amounts of compute power. These would certainly include modeling an organism from individual cell right up through organs, biological systems, and the fully integrated biological entity. This is one of the ways people are thinking about to get better medical therapies in place.

---

**Hampton, Va.:** I have noticed that the major advances in computing over the last few years has resulted from increasing the number of processors working together. However, the speed of the individual processors themselves has not increased very much recently. This has resulted in the strategy of increasing the number of "cores" in today's commodity processors to achieve faster speeds. What are the prospects for and difficulties faced in creating new chips with significantly faster core speeds than today's crop?

**David Turek:** Well the prospects for creating faster cores is kind of the crux of the problem. The rough physical analog of Moore's law is perhaps best understood by observing how many times you can fold a piece of paper. The first couple of folds are easy but pretty soon it becomes impossible. That is kind of what has happen with the doubling of densities in microprocessor design; you get to a point where it becomes impossible to make anymore folds. So one answer is multicore but that puts a premium right back on software and begs the question of balanced design: how do you feed all those cores. I think you'll start to see more work going on in 3D integration and of course there will continued to be more work on new materials. But this is just another phase in the journey and the games we played over the last few decades simply have new rules. This should spark some serious innovation

---

**Washington, D.C.:** Hi, two questions: How will the new supercomputer that comes out next year compare to the most powerful computers of years past? For instance, I would imagine that NASA had pretty powerful computers when they sent men to the moon. How many times more powerful will this one be?

Also, I have trouble believing that supercomputers will change the scientific method. Even microscopes and telescopes did not do that -- they simply gave scientists a better tool for investigation, right?

**David Turek:** The Roadrunner system we expect to deliver next year will be the fastest machine by roughly a factor of 2-3 over the fastest machine today. But this is a giant game of leapfrog and we

would expect to deliver even more powerful machines after that. IBM participates in a program sponsored by DARPA called HPCS for which we are designing a massively scalable petaflop system with prototype and production system targeted for 2010 and 2011. And you will observe a continuing growth in the Blue Gene install base as well. The irony to this is that when NASA sent men to the moon they used computers that would be dwarfed in compute power by the system on which I am writing this answer. NASA had and still has excellent scientists, engineers, programmers and mathematicians who made it all work well. So never underestimate the power of human intelligence to solve a complex problem. I think the point on scientific method is not the repudiation of the tenets set out by Sir Francis Bacon 400 years ago but rather the evolution of the method to incorporate as legitimate the notion of simulation as an arbiter of science.

---

**Alexandria, Va.:** So how close is a "quantum computer" to reality?

**David Turek:** Well, the waggish answer would be that every year for the last 15 years it has been 15 years away. That said, IBM has a very active program focused on quantum computing and there are a variety of academic and other endeavors in this field as well. But this is "hurt your head" hard stuff so I think it is a ways away.

---

**Newburgh, NY:** Will we ever have a home PC capable of a petaflop?

**David Turek:** It we do it is more likely to be a game device than a PC. The reason has to do with purpose (or applications). Most PCs are used for what appears today to be fairly mundane purposes (surf the internet, do a spreadsheet, photo processing, etc). Not much demand for FLOps for that stuff. But the gamers have a different paradigm and that is what led to the genesis of the Cell Broadband Engine we co-developed with Sony and Toshiba. Five of those gets you to a teraflop today (single precision) so maybe in the future we will get there. But please contact your local utility before you put yours together.

---

**Munich, Germany:** I once had the privilege of working in the computer room of a CDC Cyber Supercomputer. As archaic as it was, it could run simulation programs with frametimes of 500ns, which couldn't be replicated with a cheaper (\$1 million instead of \$20 million for the Cyber) newer supercomputer.

But since supercomputer technology has been implemented on integrated circuitry, raw processing power doesn't seem to be as great an issue as moving the data to and from multiple CPUs. Are data transfer to and communication between parallel banks of computing units the current hurdles to overcome in increasing supercomputer performance?

**David Turek:** Data transfer is one of the major hurdles because we are entering a phase of machine generated data whereas the past was of volumes of data more consistent with human transactions. Machine generated data comes from rfid devices, from video cameras, from gene sequencing machines, from NMR devices and so on. And the volumes can be daunting. A single digital mammogram is around 100MB of data but multiply that by the number of women who need this procedure and the fact that the data needs to be stored through time. So this is a huge issue for the future and is an area that i think the venture capitalists are beginning to pay attention to.

---

**Bethesda, Md.:** A simple but naive question - I am not a computer jock: When dealing with machines

of such capacity, how can you be comfortable with that the assemblage, and functioning is error free?

**David Turek:** The simple answer is you can't and you experience that phenomenon everytime your desktop or lap top freezes up. It is impossible to test all possible pathways in all possible ways on any computer. So the test strategy for all companies is an exercise in approximation (which yields very good results by the way). But testing is an after the fact phenomenon so the better approach is to design recovery and redundancy into the systems from the beginning. That is, expect the systems to fail or have problems, but automate the system in such a way that the system can self-recover from the problem. That was the philosophy that stoked our pursuit of autonomic computing over the last several years and is featured in all of our designs. It is one of the key features that will be highlighted in our PERCS system that is being designed in concert with DARPA as part of their HPCS program.

---

**Washington D.C.:** How much do supercomputers cost and are there restrictions on who can buy them? Does the government control it? I would think we would not want Iran to have one of these, for example.

**David Turek:** Fastest machines in the world can cost upwards of \$100M but the design we employ constructs these systems out of composable units which cost far less and are supercomputers in their own right. Put it this way, you could buy for a few hundred thousand dollars today a machine that would have been the fastest in the world in 1996. That is a huge improvement in economic performance. There are some restrictions on these systems and export controls for these types of computers are managed through the Department of Commerce

---

**Washington, D.C.:** A climate researcher who uses supercomputers once told me: "Hardware is easy. Software is hard."

What recent advances have been made in programming code to make these supercomputers solve problems more efficiently? I understand the temptation to "throw a bunch of processors" at a problem, but some scientific problems don't break down into tiny chunks very easily.

**David Turek:** Your friend is almost right. I think it is all hard but true usefulness only comes with the amalgamation of hardware and software in the right way. DARPA has been at the forefront of supporting research to improve programming on these systems but there is a ways to go. So beware the trap of thinking more cores are better than fewer cores; the real question is: how do you program the system.

---

**Chevy Chase, Md.:** Excellent story, but in stories like this I must rely heavily on the writer because I have only a modest ability to independently assess the accuracy or reasonableness of what is said. When I come across a statement of fact that I can assess, and that appears to be pretty far off the mark, it shakes my confidence. The story says that the computer uses 4 megawatts of power, the equivalent, the story says, of using 10,000 lightbulbs. Wouldn't this be true only if each lightbulb uses 400 watts? Wouldn't it have been better to have said something like "the power consumed by 40,000 100 watt lightbulbs"? Or "67,000 60 watt lightbulbs"? It was a throw-away line, very tangential to the point of the very fine article, but at least for me it is an annoyance and a detraction.

**David Turek:** point well taken. THE 4 megawatts is correct. I use flourescent light bulbs so for me it would be a lot more than what the story says. Sorry for the confusion

---

**Sunnyvale, Calif.:** Good afternoon, Dave. One trend we've noticed in the industry is that HPC is increasingly dominated by "leveraged" architectures -- systems primarily designed for business computing, repackaged or repurposed for HPC. IBM Blue Gene (along with Cray and SGI Altix, arguably) is a notable exception. What value does IBM get out of designing systems for the supercomputing crowd? -- Addison, Tabor Research

**David Turek:** We get value in a variety of ways but one of the key ways is exposure to problems in the near term that seem quite isolated but which typically turn out to be a foreshadowing of the problem becoming more general over time. I remember back in the mid-nineties the first time a client came to my lab and talked about the horrific problem of the terabyte files they were about to encounter. Laughable in hindsight but leading edge back then. So we see stuff coming out of this community that we can react to earlier than if we waited for the complex problems to manifest themselves in the general market. It's the simplicity of the early bird parable. And of course, a good chunk of the technology we design cascades into our more general product lines over time. A perfect example is the Blue Gene system. We announced that as a research project at the tale end of 1999 and when people saw the design principles behind it they thought we were crazy. That was because the conventional wisdom of the period was that you gained performance by using fast micros. we saw the iimplication of that on power usage and went massively parallel with low power processors. now everone wants to play the green computing card but we've had a pretty good head start on the field.

---

**Atlanta, Ga.:** Hey there! Could you talk a little bit about how the computers are actually ranked for speed? How do you determine who's computer is fastest? Is it done independently somehow? That's always left me curious.

**David Turek:** The accepted ranking criteria used by many is the results of the linpack benchmark. It is well known and understood that this is not a perfect benchmark but the fact that it boils down to a single number has made it very consumable. There are other measures as well (the HPC Challenge benchmarks for example) but when it comes down to a customer buying a system they all want to run their particular application on the system to see what happens. It's the same as the EPA saying your mileage will vary depending on how you drive.

---

**Silver Spring, Md.:** Rensselaer Polytechnic just installed a super computer, reported to be the largest on any college campus. How does that one compare to the one's in the article?

**David Turek:** The RPI system is one of the most powerful systems in the world today, not to mention in academia. It differs by virtue of already being installed and used and it is a different architecture (Blue Gene). It is also highly focused on research problems in nanotechnology.

---

**Christopher Lee:** Well, time to pull the plug on this one, folks. Many thanks to Dave Turek for leading a fascinating discussion about supercomputing and it's implications for all kinds of scientific inquiry. Thanks, too, for all the great questions from our readers; the brainpower out there is pretty impressive. See you next time.

---

**Marlboro, N.Y.:** Why should the average person care about supercomputers?

**David Turek:** It turns out that the average person derives the benefit from supercomputing in products and services used every day. The drugs you use, the car you drive, the gas you pump, the movie you

watch, the soap you pour, all have been improved through the use of supercomputers in those and other respective industries. The ubiquity of supercomputing (and there are tens of thousands of these things deployed world wide) are in integral part of commerce and science today. Take a look at some of the publications from the US Council on Competitiveness for High Performance Computing for other examples.

---

**David Turek:** Thanks everyone for the stimulating questions. I haven't had to type this fast since my high school typing class. I hope the quality of my answers was on par with the quality of your questions. I almost got to them all but sadly enough we are out of time

---

Editor's Note: washingtonpost.com moderators retain editorial control over Discussions and choose the most relevant questions for guests and hosts; guests and hosts can decline to answer questions. washingtonpost.com is not responsible for any content posted by third parties.

### Post a Comment

[View all comments](#) that have been posted about this article.

You must be logged in to leave a comment. [Login](#) | [Register](#)

**Submit**

Comments that include profanity or personal attacks or other inappropriate comments or material will be removed from the site. Additionally, entries that are unsigned or contain "signatures" by someone other than the actual author will be removed. Finally, we will take steps to block users who violate any of our posting standards, terms of use or privacy policies or any other policies governing this site. Please review the [full rules](#) governing commentaries and discussions. You are fully responsible for the content that you post.

© 2007 Washingtonpost.Newsweek Interactive

#### Ads by Google

##### [Cloud computing](#)

Run your own virtual data center in the cloud today

[www.3tera.com](http://www.3tera.com)

##### [Fast Super Computer](#)

NVIDIA® Tesla™ Technology Brings Supercomputing Power to Any PC

[www.NVIDIA.com](http://www.NVIDIA.com)

##### [Computer Science Degree](#)

Earn Online Computer Science BS/AS Classes Forming Right Now at MSU!

[www.mountainstate.edu](http://www.mountainstate.edu)