

# innovation:

AMERICA'S JOURNAL OF  
TECHNOLOGY COMMERCIALIZATION

## Dumb, but Fast

April/May 2007

*What's next for supercomputing*

By Dave Turek

IBM's legacy of R&D partnership with the federal government and academia stretches back nearly 50 years. Government-sponsored research has been an important crucible for the computer industry. Indeed, some of the most essential technologies and building blocks that underpin the modern IT marketplace can trace their lineage back to government-sponsored research. Today, I believe we are at a familiar crossroad as our R&D work with federal laboratories is poised to energize the commercialization of high-performance computing (HPC), also referred to as supercomputing.

The Department of Energy and the Pentagon's Defense Advanced Research Projects Agency (DARPA) are two of the largest government users and sponsors of supercomputing technology. Developing supercomputing systems that are more widely usable by government and industry is now a key component of DARPA's High Productivity Computing Systems (HPCS) mission. Through HPCS, the U.S. government recognizes that high-performance computing will be a critical enabler for the economy going forward. The ultimate goal of the HPCS program is to simplify the programming and commercialize state-of-the-art supercomputing to the point where it becomes a viable tool that can fuel and speed continual innovation by a broad base of business users, academic researchers and government scientists.

Even as the government was making commercial viability of supercomputers the guiding principle of future systems designs, IBM recognized the limits of current approaches. By 1999 it became clear that current design trends would quickly hit practical limits on physical size, electricity usage and cooling requirements. IBM, working with the DOE and Lawrence Livermore National Laboratory, developed an all-new HPC architecture that would package modestly-priced components with a design that offers favorable price/performance and performance-per-watt characteristics and could scale extremely well. Delivered in 2005, the system, called Blue Gene, is built from PowerPC processors running at a conservative 700 MHz clock speed. The relatively low amount of power required for this chip allowed IBM to pack a lot of them close together. Each Blue Gene rack contains 1,024 dual-processor nodes in a footprint that dramatically saves floor space.

Blue Gene/L at LLNL harnesses more than 100,000 processors and is currently the fastest computer system on earth, capable of peak performance of 367 teraflops (367 trillion floating-point operations per second). Blue Gene's innovative design has allowed it to quickly cement its lead at the fore of scientific research. In the past 18 months, applications supporting astronomy, petroleum exploration, fluid dynamics and biotechnology have been successfully ported to Blue Gene. A Blue Gene ecosystem has begun to mature as more applications are enabled for the architecture, attracting a wider range of users. As prices per "flop" drop, we're seeing more commercial interest in Blue Gene in 2007.

Beyond Blue Gene, the compelling economic forces that shape the computer hardware market are lending a strong tailwind to the government's efforts to spur commercialization of HPC. Manufacturers and producers of supercomputers have leveraged the price/performance improvements of commodity microprocessors with new techniques in interconnect technologies to continually lower the costs of HPC systems. In June of 1997 the so-called "fastest computer in the world," the ASCI Red machine at Sandia National Laboratories, was the first system exceeding a teraflop in compute power. Today, the same amount of compute power could be acquired for around \$200,000, making supercomputing affordable to small companies, single academic departments and, in some cases, even individual researchers.

The ramifications of ever-more-affordable compute power to U.S. businesses can be substantial. For traditional manufacturers and consumer goods companies, for example, the ability to model, simulate and test prototypes with the detail provided by a supercomputer greatly reduces the need to build physical prototypes. This saves considerable development costs and shortens the time it takes to bring new products to market. Nearly every week, my colleagues and I at IBM hear about new applications that our large customers are devising for their supercomputers. A few years ago, many of these applications would have been considered too tangential or esoteric to spend precious supercomputer resources on, even for a Fortune 500 company. Not anymore. Large companies are riding the price/performance curve to drive high-performance computing more deeply into the enterprise, a trend that should be in full gear by 2010.

But perhaps the greatest potential benefit made possible by affordable and accessible supercomputing resources waits to be tapped by America's small and medium manufacturers. The U.S. manufacturing sector currently accounts for roughly 13 percent of GDP, employs over 14 million workers, and ships more than 60 percent of U.S. exports. It is the largest manufacturing sector in the world. And it's comprised mainly of smaller firms—small and medium businesses represent 98 percent of all manufacturing firms and account for half of all the manufacturing jobs in the U.S. Many of these manufacturers are part of larger industry supply chains and will be driven over time to increased use of high-performance computing simulations by their big customers. This is happening now in the automotive sector; auto manufacturers are quietly leading their suppliers into the modern era of simulations and testing done primarily on supercomputers. A few years ago, a maker of exhaust systems, for example, would have provided a Big Three manufacturer with a box of prototype parts ready for testing and fine tuning. Now these suppliers are increasingly delivering prototypes in the form of software code, to be simulated and tested on powerful supercomputers before the first piece of metal is cast.

So, some sectors of America's small-business manufacturing economy will be ushered into high-performance computing by large customers or industry supply-chain networks. But others will need a different pathway to access supercomputing resources that will be required to compete in an increasingly global marketplace. Here too, early efforts to lower the barriers to entry and ease commercial adoption of HPC are encouraging. The Ohio Supercomputer Center's Blue Collar Computing™ program provides an example of public/private partnership that could serve as a model for delivering supercomputing resources to small manufacturing companies. Blue Collar Computing aims to bring supercomputing out of the ivory tower of academic research and make it more broadly available as a local economic development tool.

To quote from the group's web site: "Blue Collar Computing (BCC) is an approach to supplying computational modeling and simulation solutions to companies without requiring large up-front investment. The BCC model is service-oriented, relying on amortizing the cost of expertise, software and hardware across many firms to reduce individual expenditures. The BCC model will enable the majority of U.S. firms to exploit numerical models and simulations."

BCC is the model for a current bill winding its way through congress. The Blue Collar Computing and Business Assistance Act seeks federal funding to establish up to five such centers across the U.S. It's possible that legislation such as this act—which ultimately provides a way for U.S. business to reach the

fruits of computational research—will become a logical adjunct to federal research funding going forward.

From a hardware point of view it's fairly obvious that advances in supercomputing speed, energy-use and cooling will come in foreseeable, incremental stages. The latest DOE HPC project, known as Roadrunner, will underwrite more research and investment into the next generation of ever-more-powerful, commercially viable supercomputer designs. For Roadrunner, IBM plans to leverage existing investment in its Cell videogame processor and off-the-shelf Opteron chips from AMD to develop a hybrid supercomputer. Such hybrid designs promise to yield powerful computers that run a diverse set of applications at reasonable prices.

The weak spot for full commercialization of HPC is software. Innovation in hardware has far outstripped innovation in software over the past 15 years. Substantial effort will be required to re-write essential code that takes full advantage of multiple processors. Here, market forces are holding back the tide of HPC commercialization—many software vendors are too small and their customer bases too thin to support the substantial investments required. Government efforts are beginning to address this issue, but for the foreseeable future, we would do well to concentrate our collective efforts on software deployment if the U.S. is to reap the benefits of fully commercialized supercomputing.

Dave Turek is the vice president for deep computing at IBM. He has business responsibility for high-performance computing solutions including Power, Intel and AMD based servers and workstations, Blue Gene systems, visualization solutions and future technologies.

\*\*\*\*\*

### Wooden Terms for Wondrous Technology

For non-engineers, the lexicon of supercomputing seems a jumble of acronyms and terminology culled straight from the dull innards of a physics textbook. With the exception of the term, “supercomputer,” the computer industry seems to excel at creating and using wooden and cryptic descriptors for some of the world's most wondrous technology. Unfortunately, for laymen, that trend shows no signs of abating anytime soon. Herewith, then, a short guide to some of the most commonly used phrases in supercomputing.

What exactly is a supercomputer?

Roughly speaking, a supercomputer is any computer that, upon its introduction, provides leading computational capacity and performance, and is deployed to run the most challenging, complex problems at the time. Often, supercomputers become more basic computers over time, as newer, faster systems are introduced. IBM's RS/6000 family of scientific supercomputers (now called pSeries) quickly morphed into more general-purpose business computers a couple of years after they were introduced in the late 1990s.

### Floating point operations

A floating point operation (flop) is an operation, such as addition and multiplication, between numbers with decimal points. Think of flops as the computer's version of scientific notation. Floating point calculations play a central role in all of the science and engineering disciplines. Since supercomputers were originally designed to solve problems in these disciplines, the number of flops per second that a computer could execute became the benchmark for performance.

The Flop sisters— Mega, Giga, Tera and Peta

Unlike your home PC, which rates performance based on clock speed (how fast the processor is capable of running), supercomputers are rated according to the number of floating-point operations they can churn out in one second. For a rough, automobile analogy, think of a PC as being measured in RPMs, while a

supercomputer is measured in horsepower. As supercomputers are capable of doing enormous numbers of flops, speed is measured in millions (mega); billions (giga); trillion (tera) and quadrillions (peta) flops. The industry uses a standard software benchmark tool called Linpack to measure flops. In the real world, more flops generally means that demanding scientific and engineering applications can offer results in hours or days instead of weeks and years. The fastest supercomputer in the world, IBM's Blue Gene/L at Lawrence Livermore National Laboratory, is rated at peak performance of 367 teraflops—367 trillion floating point operations every second. Exactly the kind of horsepower required to handle the sophisticated calculations that would otherwise keep DOE scientists waiting months for the results of a single simulation. In starker terms, it would take one scientist with a calculator 177,000 years to do what Blue Gene can do in one second. And the poor guy or gal would be working 24/7—and that's 743,400 person years.

#### Clusters & Nodes

Some supercomputers are assembled from nodes and clusters of smaller computers harnessed together. A node is a computer—it's independently capable of doing computations. Clusters are collections of nodes that are networked by a high-speed interconnection and managed from a single point. A cluster appears to the end user as a single computer. Many systems designs use clusters of commodity processors to create supercomputer platforms.

#### Blue Gene

Blue Gene marks a radical departure from how supercomputers are designed or "architected." In 1999, IBM engineers realized that if supercomputer architecture stayed the current course, the machines would soon require their own football-field sized buildings to house them. They would eat enough electricity in one year to power a mid-size town and they would require yet more power to cool them and prevent them from overheating. Enter Blue Gene and a \$100 million, five-year development effort by IBM. Designed to harness thousands of low-power, cooler-running processors that would mitigate the need for extraordinary power and cooling systems, Blue Gene was originally built to help biologists observe the invisible processes of protein folding and gene development. Hence the name—a play on IBM's nickname, Big Blue, and the application it was designed for. By 2006, Blue Gene systems held a dominant 28 spots on the industry's yearly ranking of the Top 500 fastest supercomputers, including the No. 1 position for the 367 teraflop system at LLNL, which lashes together more than 131,000 Power PC processors. —IBM

[Close Window](#)