

# Beyond 'Big Iron' in Supercomputing

A National Science Foundation effort is going beyond number-crunching muscle to merge supercomputers and visualization into a new information infrastructure

URBANA-CHAMPAIGN, ILLINOIS—Just an elevator ride from the office of Larry Smarr here at the National Center for Supercomputing Applications is a 3-meter virtual-reality cube. Put on special goggles, step inside, and you will find yourself in a three-dimensional world summoned up on the walls of the cube by computers sitting out of view. When Smarr recently demonstrated the cube for a visitor, the first scene that flashed to life was an animated crayon landscape. You could walk into the crayon house, splash through the crayon pond, or get chased by bees in the crayon forest. Or you could stand in a crayon meadow while Smarr gave an impromptu description—in effect, a crayon sketch—of what may be the future of high-performance computing, networking, and visualization: the National Computational Science Alliance.

The Alliance, which Smarr directs, came to life just a year ago as one of two winning proposals to succeed the National Science Foundation (NSF)-backed Supercomputer Centers, a program that provided high-end computing at four or more sites around the country. The contraction to just two centers is deceptive: The Alliance, which will receive about \$30 million from NSF in fiscal year 1998, involves researchers at over 60 institutions ranging from the University of Illinois, Chicago—where researchers created the virtual-reality cubes and "Crayoland," a whimsical demonstration program—to Caterpillar Inc., which uses the cubes to test prototypes of its new earth-moving equipment. The other winner, the National Partnership for Advanced Computational Infrastructure (NPACI), headquartered at the San Diego Supercomputing Center (SDSC), has a similar number of partners and level of funding.

These partnerships among universities, government labs, and industry have a more ambitious mission than the earlier supercomputing centers, says Robert Borchers, director of NSF's Advanced Computational Infrastructure and Research Division. "The whole community has come around to the idea that 'big iron' [massive hardware] isn't the be-all and end-all of computing these days," he says. One major task for the partnerships is developing software to reap the

benefits of the new "parallel" computer architectures, in which hundreds or thousands of processors work in tandem to speed up computation. Another is developing technologies that can coordinate computers and people linked over thousands of kilometers by high-speed networks like the NSF's new very high performance Backbone Network Service (vBNS), which connects some universities and national labs.

"What the NSF is saying is, 'We want you to actually create a working prototype of the early 21st century information infrastructure,'" explains Smarr. He and his partners expect that infrastructure, which they are calling the National Technology Grid, to include radically new computing and communications technologies such as shared tele-immersion,

cisco. The enthusiasm is shared widely—but not universally. Some researchers who depended on the old centers for supercomputing time have found that the new program isn't meeting their needs, and Alliance members concede that their supercomputers are in some cases oversubscribed by a factor of 4. Borchers and Smarr say they are scrambling to add new big iron to meet the demand.


The original centers program had become a cornerstone for many fields of science and engineering. In the early 1980s, explains John Connolly, director of the Center for Computational Sciences at the University of

Kentucky and a member of the Alliance's executive committee, "supercomputers were so expensive that no university could afford them." The \$65-million-a-year centers program made an arsenal of supercomputers available to the research community, ultimately at the Cornell Theory Center, the Pittsburgh Supercomputing Center,


SDSC, and the National Center for Supercomputing Applications (NCSA).

In 1994, a blue-ribbon panel recommended that the centers program be renewed beyond its 10-year mandate. "Things were going pretty well; we decided to go along with that [recommendation]," says Borchers. But when the agency presented its case to the National Science Board, the presidentially appointed body that oversees NSF, "in essence, we got turned down," he says. The board, Borchers says, believed that the growth of parallel processing and high-speed networking had changed the world of computing. As a result, after yet another review panel and a 2-year extension for the centers, NSF solicited proposals for a new program emphasizing software development and multidisciplinary partnerships.

Out of six proposals—including four from the original centers—the NSF announced a brace of winners in late March of last year: the Alliance and NPACI (*Science*, 4 April 1997, p. 29). "The structure of both partnerships is quite similar fundamentally," and many of their interests are the same, says Sid Karin,



**Wild heart.** A supercomputer simulation of turbulence in a red giant star reveals its hot core (yellow).



IMAGES: WOODWARD ET AL.

in which researchers at different locations can manipulate data within the same virtual environment, and metacomputing, in which widely separated computers work in parallel as a single machine. Although these innovations will at first benefit mainly researchers in universities, government, and industry, ultimately they—like the personal computer and the World Wide Web before them—could touch the life of next century's Everyman. "It's a once-in-a-lifetime opportunity to change the world," says Smarr.

"This initiative is unique both for its scale and its ambition," says Jim Gray, a senior researcher at Microsoft Research in San Fran-



director of the San Diego center and the NPACI executive committee. "The details are different." NPACI will put a special emphasis on maintaining, managing, and transmitting large databases and on very high end computing, including both traditional "vector" and parallel machines. The Alliance, meanwhile, aims more generally at developing software and technology for that future grid.

The model for that effort is the Information Wide Area Year, or I-WAY, a short-lived collaboration staged for the Supercomputing '95 conference in San Diego (*Science*, 26 January 1996, p. 444). Using novel communications protocols over high-speed links, a handful of U.S. supercomputer centers performed computations cooperatively, passing data back and forth to create a giant metacomputer. In some experiments, supercomputers in one location supported collaborations between researchers at two other sites, who inhabited a shared virtual environment.

"I-WAY was the vision for the Alliance," says Rick Stevens, director of the Mathematics and Computer Science Division at Argonne National Laboratory in Illinois. "That really convinced all of us that if we built it, they would come." Adds Smarr, who is the head of NCSA and a professor of physics at the University of Illinois, Urbana-Champaign (UIUC): "For that one shining moment, America had a 21st century information infrastructure."

To create something more persistent, the Alliance has set up six "supernodes" with especially impressive computing power. The machines housed there are some of the most powerful outside the national weapons laboratories, with brand names ranging from IBM to Hewlett-Packard to Silicon Graphics Inc./Cray Research, the makers of NCSA's Origin2000 machines, which now have a total of 512 processors that can work in parallel. Other Alliance sites have one or more specialties within computer science, outreach, and education. They are also headquarters for the Alliance's applications teams, groups collaborating on the specific problems that are driving the software and hardware development.

### Lines of communication

The applications were chosen to span disciplines and stretch the computing technology to its limits; they include some of the biggest problems in astrophysics, materials science, and molecular biology. To spur software development, the Alliance is also taking on computer-hungry problems outside the traditional bounds of academia. Carl Kesselman of the Information Sciences Institute at the University of Southern California in Los Angeles and Ian Foster of Argonne, for example, are working with researchers who devise "synthetic theaters of war" involving tens of thousands of individual elements—tanks,

trucks, planes, helicopters, missiles, radar, and even individual rounds of ammunition. Thousands of processors on dozens of different machines have to work together to keep track of these elements and what happens to them during the ebb and flow of battle.

Staging a war game, says Foster, means simultaneously scheduling time on distributed machines, then logging onto them, and finally getting the computers to operate happily together. In the past, researchers wanting to run a computation of this magnitude "have virtually called up the administrators of all these systems and said, 'I want to do this big applications run at 2 p.m. for a couple of hours,'" says Foster. "It's a very painful process." In order to make it less painful, Foster and Kesselman are developing and testing software they call Globus.

Foster compares Globus to a travel agent who makes sure that plane tickets, rental cars, and hotel rooms are all available for a single business trip: "You say, 'I want to go to this conference; make it happen.'" Globus works through a single password "proxy" that is trusted by all the systems on a grid of computers around the world. "On one computer screen, we can type the equivalent of GO or RUN, and Globus starts the individual machines," says Paul Messina, director of the Center for Advanced Computational Research at the California Institute

of Technology, whose team created the war games. "That alone goes a long way toward making these distributed systems practical."

Enabling separate processors to communicate, however, leaves a larger problem unsolved: deciding when they *ought* to communicate. Even the fastest links, such as the vBNS, can become a bottleneck if the processors need to exchange information and share memory across the network for each step of the computation. The trick is to divide up a problem among processors in a way that doesn't require a constant stream of communication while they are doing their jobs.

Paul Woodward, director of the Laboratory for Computational Science and Engineering at the University of Minnesota and a member of the Alliance's executive committee, has attacked this problem in the context of extremely large-scale simulations of fluid turbulence. The turbulence Woodward and his team are interested in takes place in old, bloated, red giant stars. It plays a major role in determining how old stars spew their

outer layers into space before they die, opening the way to a new round of star formation.

The calculations are immense, because they have to track such a wide range of scales: While a red giant can have a diameter the size of Jupiter's orbit, the fluid computations have to take account of ripples and eddies as small as hundreds of kilometers across. One recent calculation by Woodward's team ran for 14 days on 128 processors of NCSA's Origin2000, generating 2 billion bytes of data. But Woodward says, "We always have an appetite for a larger calculation."

To satisfy it, he would like to be able to yoke together as many processors as he can get, wherever he can find them. So he and his colleagues have developed a "scalable" algorithm, designed to work "whether the many processors are inside one box, or on one

campus, or in different states," says Woodward. That meant restricting the amount of communication between the individual processors and between the processors and shared memory.

He divides the calculations into "chunks" or "bricks," which have to be completed by a single processor before the result can be shared. "You can update a brick without knowing about other bricks," says Woodward. Only then is the result made available to the network. That approach, he says, should allow researchers to perform parallel computations on

many kinds of processors and networks. The only requirement is that the network be fast enough to keep up with the processors as they crunch through each chunk of computation.

### CAVE explorers

Besides finding ways to link computers, the Alliance also aims to link new users to the grid. Take the case of Caterpillar Inc.—not a traditional supercomputer user—and the CAVEs, which were developed principally at the Electronic Visualization Laboratory of the University of Illinois, Chicago.

The CAVEs, like the one in which Smarr unfolded his vision of the Alliance, are room-sized cubes that produce a compelling illusion of a 3D world—the most elaborate virtual-reality machines yet devised. The user wears special glasses that keep track of his or her "location" in a virtual scene; liquid-crystal shutters on the glasses rapidly open and close in synchrony with projectors, which flash stereoscopic images tailored to the viewer's position. Caterpillar engineers have found, says



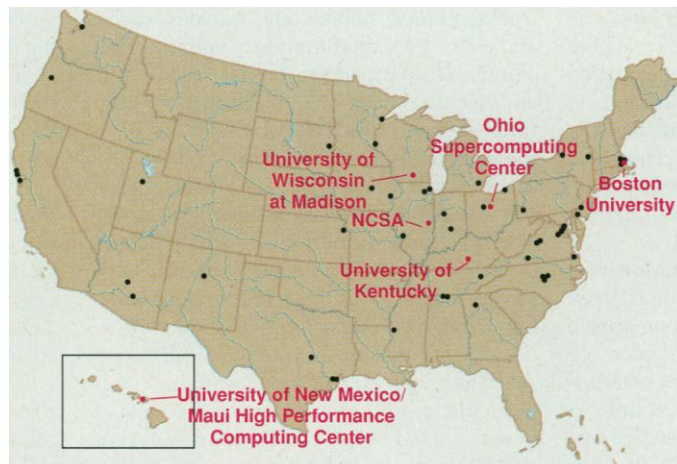
**Visions of the future.** Alliance director Larry Smarr in Crayoland, a virtual reality environment that was created at the University of Illinois, Chicago, by Dave Pape.

THOMPSON-MCCLELLAN

Robert Fenwick, the campus manager for Caterpillar-UIUC relations, that a CAVE can help them quickly test and revise designs for earth-moving equipment while building fewer iron models than in the past. "We don't do it because we think [the CAVE] is a neat toy."

Late last year, Caterpillar took another step when it linked a virtual-reality machine at UIUC with one in Bonn, Germany, allowing engineers to collaborate in testing the operation of a wheel loader—"a great big machine with a big bucket on the front that you load dump trucks with," is Fenwick's definition—in a strip mine. Although the CAVE link was just a demo, using an existing Caterpillar product, "shared tele-immersion" intrigues the global company as a way of including many engineers in the design process while avoiding huge amounts of travel.

The Alliance is also working toward a different kind of inclusion—making all of this new technology as accessible as possible to minorities, women, and people with disabilities. "There is a real danger that [the grid will] end up amplifying inequality," says Roscoe Giles, director of the Center for Computational Science at Boston University and an Alliance executive committee member. In one effort to head off that danger, the Alliance includes members such as the American Indian Higher Education Consortium (AIHEC). The AIHEC is already exploring possible ben-



**Covering the map.** The Alliance includes 60 universities, companies, and laboratories. Its six "supernodes" (red) are supercomputing powerhouses.

efits such as collaborating with the Alliance applications group in environmental hydrology to calculate the impact of tearing down a dam on the yield of wild rice, or securing online educational materials for tribal schools. The Alliance has also been "very proactive" in seeking advice on how to make the grid as accessible as possible to vision- and hearing-impaired people, says Gregg Vanderheiden, director of the Trace Research and Development Center at the University of Wisconsin, which has longstanding research efforts in these areas.

In the workaday task of satisfying the demand for computing cycles by the research community, however, both the Alliance and NPACI have fallen behind. "The capacity at

the two centers that have been terminated has been shut off to the academic community," says Kentucky's Connolly, who is involved in scheduling computer time for the Alliance. A major setback came when negotiations to fold the Pittsburgh Supercomputing Center into the Alliance broke down, says Borchers of NSF. The Pittsburgh Center is now turning to programs such as the Department of Energy's defense-oriented computing program for funding. "It's caused a little bit of pain," says Connolly.

Some members of the supercomputing community are harsher: "People don't always view [providing computing time] as glamorous," says one. "The result, in my book, is that those folks haven't focused on [routine] high-end computing." Still, this source expects the problem to "heal itself" as new hardware gets funneled into the Alliance sites.

Meanwhile, Smarr is dreaming within the crayon landscape of the Alliance about finding the next "killer app," the software application that will transform 21st century society. "E-mail was the killer app of the ARPAnet"—the prototype of the Internet. "The Web for the modern Internet. What's gonna be the killer app of the vBNS?"

—James Glanz


## RADIO ASTRONOMY

### Motorola to Limit Signals Over Arecibo

The world's largest radio observatory will not be drowned out by round-the-clock cell phone chitchat. The National Astronomy and Ionosphere Center (NAIC), which operates the Arecibo Radio Telescope in Puerto Rico, announced last week that Motorola Inc. has agreed to limit satellite transmissions above the observatory in a critical part of the radio spectrum for 8 hours a day. The pact, together with a previous umbrella agreement that grants the nation's other 14 radio observatories a shorter observing window, clears the way for Motorola to connect United States customers to a new, worldwide cellular phone service.

The heart of Motorola's new service, called Iridium, will be a \$5 billion network of 66 satellites that will transmit phone calls to subscribers by radio signals tuned to about 1621 megahertz. However, just as the signal from a strong radio station may drown out distant stations that broadcast at nearby frequencies, the satellite transmissions will mask some of the faint whispers from space that radio telescopes are built to detect. The

agreement focuses on a narrow part of the radio spectrum, from 1610.6 to 1613.8 megahertz. At those frequencies, astronomers can detect charged molecules called hydroxyl radi-



**Hold the calls.** Arecibo's new reflector should have clear listening 8 hours a day.

cal signals that are found, for example, around comets and in the expanding shells of red giant stars.

The agreement does not stipulate a particular method for Motorola to silence its transmissions, in part because the actual amount of

interference has not been measured yet. Instead, the negotiators agreed on a target power level for the satellites' emissions in the hydroxyl band when they are within view of the Arecibo telescope. The interference will be reduced by a factor of roughly 30 during the nighttime hours, but Arecibo astronomers will still have to take significantly longer on their observations to compensate for the increased static. "This is a compromise solution," notes one of the negotiators, former NAIC director Michael Davis, "not a perfect solution from the astronomer's viewpoint."

But the Motorola-Arecibo agreement may become a model for future time-sharing arrangements between astronomers and communications companies. With more private satellite networks on the way and more radio bands to worry about, astronomers will have to keep fighting to protect their access to radio waves. "We'll have to be as creative as we can to protect radio astronomy into the foreseeable future," Davis says.

—Dana Mackenzie

Dana Mackenzie is a writer in Santa Cruz, California.