Mix-and-Match COMPUTING

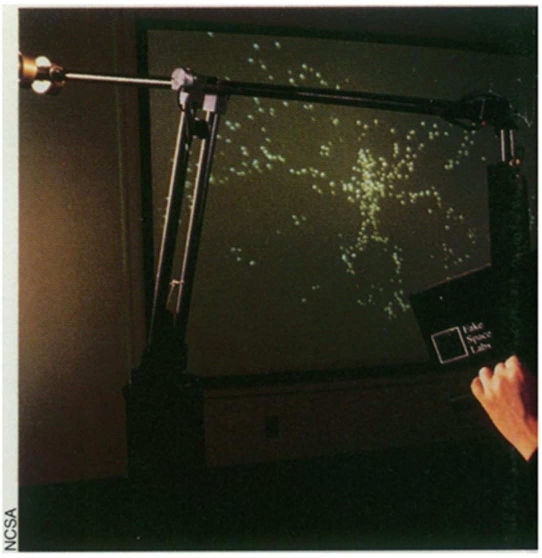
NCSA

Scientific supercomputing without supercomputers

By IVARS PETERSON

"So many galaxies . . . so little time."

Astrophysicist Margaret J. Geller's lament could just as easily have come from other researchers similarly mired in mountains of data. Just replace "galaxies" with such terms as genes, subatomic particles, polymer configurations, ozone readings, or seismic measurements.

To meet data processing and computational challenges, researchers have turned increasingly to high-performance computers. A few years ago, the automatic choice would have been a supercomputer located at a national, regional, or state supercomputing center.

Now, many centers are starting to offer a range of different computers to meet diverse needs, including graphics computers for visualization and multiprocessor machines for heavy-duty calculation. At the same time, a number of research groups are exploring the possibility of using extensive networks of ordinary desktop computers to match or even surpass the performance of a single conventional supercomputer.

To many researchers, the "mix-and-match" mode of computing that results from linking different machines provides an attractive, cost-effective alternative for relieving the work load of the heavily burdened supercomputers.

Over the last decade, Geller and her co-workers at the Harvard-Smithsonian Center for Astrophysics in Cambridge, Mass., have painstakingly and systematically recorded the redshifts of thousands of galaxies. Redshifts are increases in the characteristic wavelengths of light emitted by stars. Caused primarily by the expansion of the universe, they allow researchers to estimate the distances of galaxies from Earth. By combining these distance measurements with a database of galaxy positions in the sky, astronomers can construct step by step a three-dimensional map of the distribution of galaxies in the universe.

Geller and her colleagues have measured the redshifts of galaxies that lie within long, thin strips across the sky. Taken together, these wedge-shaped slices reveal that galaxies tend to clump into thin shells, like the walls of enormous soap bubbles hundreds of millions of light-years across (SN: 11/25/89, p.340).

To obtain these insights, the researchers used computers that provide three-dimensional views of the data. But it took a lot of experience and manipulation of the pictures on the computer screens to pick out the salient features.

As an experiment in alternative methods of visualizing huge amounts of data, Geller recently worked with graphics specialists at the National Center for Supercomputing Applications (NCSA), located at the University of Illinois at Urbana-Champaign, to animate the redshift survey. Using images of real galaxies, the NCSA team created the illusion of a journey through the universe.

This sequence became part of a 40-minute film illustrating how science is done. "I've been showing the film to standing-room-only audiences at various universities," Geller says. "People react to the graphics in an extraordinary way."

The team also converted one slice of the redshift data into a virtual-reality environment (SN: 1/4/92, p.8). By looking through a stereoscopic viewer mounted on a boom, Geller could inspect computer-generated images of the galaxies, and the scene would change as she moved her head or body.

"We were able to navigate through the slice . . . without having to have somebody preprogram the path for us," Geller says. "It certainly was extraordinary to have the sensation of really traveling through [the slice] and being in command."

"Had we had [this kind of capability] when we first obtained the data, there are a lot of things we would have known more quickly," she adds.

Geller's experience at NCSA illustrates one aspect of the changes that have occurred in supercomputing at the four national supercomputer centers, which were established by the National Science Foundation in 1985 (SN: 3/2/85, p.135).

Located at the University of Illinois, Cornell University, the University of Pittsburgh, and the University of California, San Diego, the centers originally were geared toward testing the power and versatility of supercomputers for scientific computation. Over the intervening years, these powerful machines attracted thousands of users — so many that researchers now must sometimes wait days or weeks to run their programs.

At the same time, it became evident that additional, specialized computers were needed to handle the prodigious output of the supercomputers. So the centers gradually added various machines for such tasks as visualization and graphics, and hired the staff required to support these activities. This approach gave researchers like Geller access to graphics and visualization techniques normally affordable only to Hollywood studios or large oil companies.

Now, the primacy of the traditional supercomputer — a single, enormous, multipurpose machine — is itself being challenged. Faced with supercomputer prices ranging from $15 million to $30 million apiece, many groups are looking

Computer-generated images and virtual-reality technology allow a user to view the distribution of galaxies in the universe from many different perspectives. The user can zoom in on individual galaxies or navigate around enormous clusters.

In a recent demonstration of mix-and-match computing, three linked supercomputers at the Minnesota Supercomputer Center, Inc., performed the calculations needed for a study of thermal convection in Earth's mantle. A graphics computer then converted the data into images on a screen.

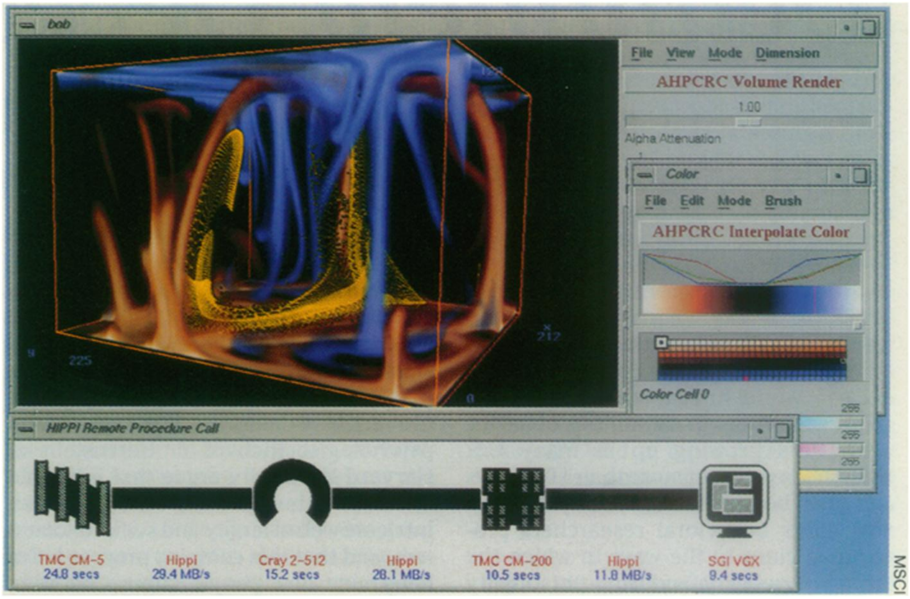

MSCI

High-performance communications links play a crucial role in network computing. This illustration shows the volume of incoming traffic for September 1991 at various points along NSFNET's high-capacity lines. The traffic ranges from zero bytes (purple) to 100 billion bytes (white).

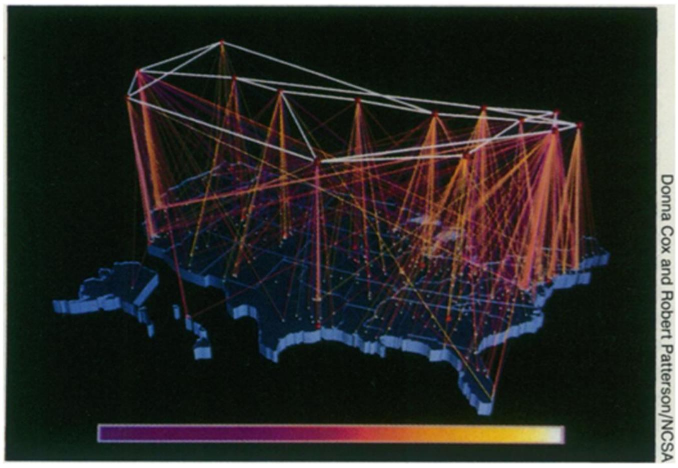

Donna Cox and Robert Patterson/NCSA

for alternative approaches for increasing computational capacity.

"We're at a critical moment in supercomputing," says Larry L. Smarr, director of the NCSA.

One possibility being explored is the linking of workstations – the kind of microprocessor-based computers that most researchers have sitting on their desks – into coordinated clusters to perform certain kinds of computations. Although such networks may take longer to solve a particular problem, the total cost of the machines involved is far less than the price of a single conventional supercomputer.

Moreover, because these desktop machines often sit idle for lengthy periods, connecting them into networks so that they can work together on large problems increases their effectiveness. Such arrangements also permit greater flexibility in selecting the number and types of computers required for a particular application.

Last year, a physicist and two computer scientists provided one of the more dramatic examples of what a collection of high-performance computers, scattered around the United States, could accomplish when linked together.

Hisao Nakanishi of Purdue University in West Lafayette, Ind., was interested in the physics underlying what happens to the shape of polymer strands passing through a membrane or trapped in a porous material such as sandstone. Confined to the material's pores, the chains of molecular units that make up polymers bend and twist in ways that differ from those possible in a liquid.

Nakanishi turned to Vernon Rego of Purdue and Vaidy Sunderam of Emory University in Atlanta for help with the computer simulations he needed to investigate this aspect of polymer physics. The team concentrated on the question of how the straight-line, end-to-end length of a polymer increases as the polymer grows into a chain and eventually traverses a cube containing an array of randomly placed obstacles. Of special interest was the "critical" case in which the cube contains just enough obstacles to provide only a single connected region comprising all the open paths along which the polymer chain can grow from one side of the cube to the other.

The researchers realized that doing the simulation on a scale large enough to yield meaningful results on a single Cray supercomputer would require several days to several weeks of computer time. As an alternative, they developed special software that treats a cluster of separate computers as a single machine, with computations divided among the participating computers.

Nakanishi and his collaborators had access to computers at Purdue, Emory, Florida State University, California Institute of Technology, Oak Ridge (Tenn.) National Laboratory, and the University of Tennessee. The most elaborate arrangement they tested combined 48 IBM RS/6000 computers, 80 Sun Sparc workstations, and two Intel i860 hypercube computers. In 10 minutes, this configuration did computations that would take three hours on a Cray Y-MP.

That was good enough for the Purdue-Emory group to earn first prize in the 1992 Gordon Bell competition. This award recognizes significant achievements in the application of high-performance computers to scientific and engineering problems. The judges describe the winning entry in the January issue of COMPUTER.

Continued on p. 284

Zealand girls who attended single-sex schools, whereas stealing, drug use, frequent sexual intercourse, and fighting occurred much more often among the 132 girls enrolled at coed schools.

Older students at coed schools, particularly boys, may demonstrate to younger girls the ways in which delinquency severs childhood apron strings and secures, at least from a teenage perspective, adult privileges, the researchers argue.

Few U.S. students attend single-sex schools, which recruit a select group to expensive private institutions, religious schools, or military facilities. But New Zealand maintains many all-girl public secondary schools, which attract students largely on the basis of academic reputation, location, and prior attendance by other family members, Caspi notes. For all its differences from the United States, he says, New Zealand offers a relatively controlled environment in which to study the ways in which puberty interacts with the social realm of high school.

Related findings emerge from a study of 125 U.S. girls attending coed schools and tracked from age 11 to 15 by Brooks-Gunn and her colleagues. Depression, eating problems, and delinquency appeared most often among those who reached puberty early and who encountered a greater number of stressful events than their peers, the researchers assert in *Stress and Coping in Infancy and Childhood* (1992, Tiffany M. Field *et al.*, editors, Lawrence Erlbaum Associates, Hillsdale, N.J.). Such events covered a broad spectrum, including parental divorce, getting a boyfriend for the first time, losing a school election, and not making an athletic team.

Regardless of their age at puberty or the type of high school they attend, some teenagers may harbor personality traits that prove incompatible both with delinquency and with emotional well-being, Moffitt says.

For instance, in a study of 101 boys and girls living in the San Francisco Bay Area and followed from age 3 to 18, Jonathan Shedler and Jack Block – both psychologists at the University of California, Berkeley – found that those who occasionally experimented with drugs as teenagers exhibited the best psychological adjustment. Experimenters used marijuana no more than once a month and had tried no more than one other illicit drug. Frequent drug users, who smoked marijuana once a week or more and had tried at least one other illegal drug, displayed few friendships with their peers, poor self-control, and emotional distress. Teenagers who had never tried marijuana or any other illegal drug exhibited anxiety, difficulty expressing emotions, and few social skills.

In each of the three groups, individuals often carried over their prominent personality traits from childhood, Shedler and Block report in the May 1990 AMERICAN PSYCHOLOGIST.

Drug use certainly does not improve mental health, and it proves highly destructive for frequent users, Shedler and Block point out. "But for adolescents more generally, some drug experimentation apparently does not have psychologically catastrophic implications," they conclude.

Moffitt considers these data congenial to her theory of dual developmental paths leading to teenage delinquency. Drug experimentation wreaks further emotional havoc on "life-course persistent" youths, she argues, but it fails to drag down basically healthy teens navigating the maturity gap. Meanwhile, some adolescent abstainers indulge in their lifelong habit of social isolation and anxiously shrink from the defiance of drug use.

Of course, teenage delinquency responds to other influences, including poverty, unemployment, rising numbers of divorces, lack of parental supervision, and violence displayed through the media. But Moffitt asserts that these factors cannot explain why crime rates charted in the United States and England, as well as in New Zealand, dramatically rise during adolescence and then reverse course as teens reach young adulthood.

Jack Block says participants in the California longitudinal study, which he organized with his late wife, Jeanne H. Block, bear some similarity to the New Zealand youngsters. A small group of Bay Area teens parlayed childhood behavioral and emotional problems into serious adolescent delinquency, he says, while many well-adjusted teenagers briefly sampled various misbehaviors as part of a search to define themselves.

Block disputes Caspi and Moffitt's theory, however, that deep-seated personality traits intensify during adolescence and other ambiguous transitions. Instead, he argues, a teenager's underlying personality may emerge from the shadows for the first time as parents and other adults begin to relax the behavioral guidelines and restrictions of childhood.

Moreover, some unstructured situations evoke a highly consistent response – caution and noncommittal behaviors – from a wide range of people, he says.

The resolution of debates over the relationship of personality development to delinquency depends on continued longitudinal research, such as the New Zealand and California studies, Block says. Few such projects exist, he asserts, because they are logistically complex, costly, and incapable of generating "fast" data for scientists anxious to publish papers and gain academic tenure.

Moffitt notes another obstacle to unraveling the development of delinquency: Researchers possess a "woeful" lack of knowledge about the meanings that teenagers themselves attach to puberty, high school, and various types of delinquency.

"We can't understand adolescence-limited delinquency without first understanding adolescents," she remarks. □

Continued from p. 281

Although the Purdue-Emory scheme represents an important first step, the logistics of handling such a network of computers remains exceedingly complicated. Indeed, the software required for binding the system together represents the main bottleneck. In many instances, software deficiencies keep these systems from running as efficiently as possible.

Nonetheless, researchers are optimistic that such problems will eventually be solved. Smarr envisions the development of a national "metacomputer" – an array of different types of computers linked by a high-speed, high-capacity network to act as a single computer.

In a sense, each national supercomputing center already acts as a metacomputer, invisibly shuffling programs and files from supercomputer to massively parallel machine to graphics computer to mass-storage device to workstation. Ordinarily, users need specify only what they would like done, and the center's software takes care of the details of when, where, and how.

Smarr would like to see this concept extended to networks of computers on a national scale. By automatically adjusting to the power and speed required for solving a particular problem, such systems would provide greater flexibility for scientists working on a wide range of applications.

"But we're not there yet," Smarr cautions.

As one step toward "scalable supercomputing" and the development of a national information infrastructure, the four national supercomputer centers last year announced the formation of a national MetaCenter (SN: 11/28/92, p.374). Center staffs are now working together to establish standards so that people can use any computer, or set of computers, at any center.

"This also allows the centers to specialize, rather than trying to be everything to everybody," Smarr says.

In response to the rapid changes in computer technology, the National Science Foundation is reviewing the role of high-performance computing in scientific research and reevaluating the rationale for the national supercomputing centers. Chaired by Lewis Branscomb of Harvard University, the panel charged with the review expects to present its report and recommendations later this month. □