SUPER PROBLEMS FOR SUPERCOMPUTERS

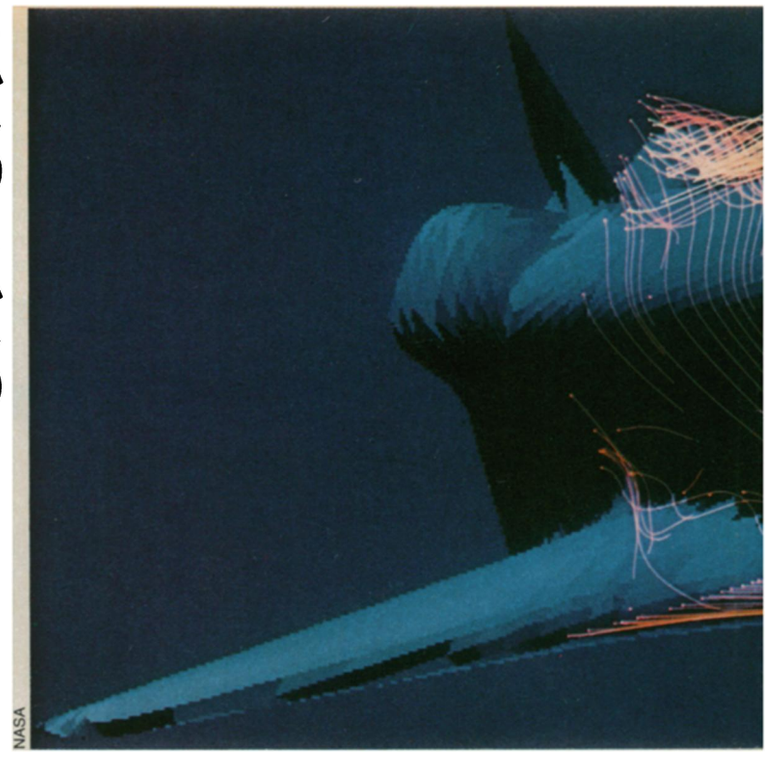
NASA

By IVARS PETERSON

Crystal growing has long been an art. The slightest disturbance — a jiggle, a cooling gust, an impurity, a speck of dust — can alter, sometimes unpredictably, the way atoms or molecules pile up to create a crystal. This sensitivity to apparently small details tends to hide how a substance actually goes from a disordered fluid to an orderly lattice. But by simulating the process on a computer, researchers are beginning to understand exactly what's going on.

Computer crystallizations, however, are slow, even on the fastest computers. It can take a day of computer time to build up 10 layers of 100 atoms each out of a simulated, supercooled liquid made up of digital atoms. "This seems like an extremely minute amount of material," says George H. Gilmer of AT&T Bell Laboratories in Murray Hill, N.J., "but it does give some information about the growth rate and the concentration of defects in the [computer] crystal."

Simulating crystal growth by tracking the motion of individual atoms as they jostle and interact is one example of a wide range of scientific problems that are stretching the memories and speeds of the fastest computers available. These "supercomputers," which range from large, general-purpose machines like the Cray-1 and the CYBER 205 to machines specially designed for a specific type of calculation, are becoming an essential tool for researchers in many fields of science and engineering. Special-purpose machines like Belle, the chess-playing computer, already exist (SN: 10/8/83, p. 236). Many more are being developed for everything from factoring very large numbers (SN: 1/14/84, p. 20) to simulating the crystallization of silicon.

At the same time, an increasing number of scientists from university, government and industry laboratories are scrambling to gain access to supercomputers. Recently, the National Science Foundation (NSF) announced grants totaling $7.3 million to make time available for scientists throughout the nation on supercomputers at the University of Minnesota in Minneapolis, Purdue University in West Lafayette, Ind., and the Boeing Computer Services Co. in Bellevue, Wash. The awards are the first step in a long-range NSF plan, directed by its recently formed Office of Advanced Scientific Computing, to provide easier access to supercomputers for the scientific community.

As further evidence of this growing interest, two recent conferences devoted a considerable amount of time to current and potential applications for large-scale computing. The meetings also highlighted important issues facing supercomputer users.

At the 1984 annual meeting of the American Association for the Advancement of Science, Peter A. Gregory of Cray Research, Inc., in Minneapolis discussed the ways in which computer simulations

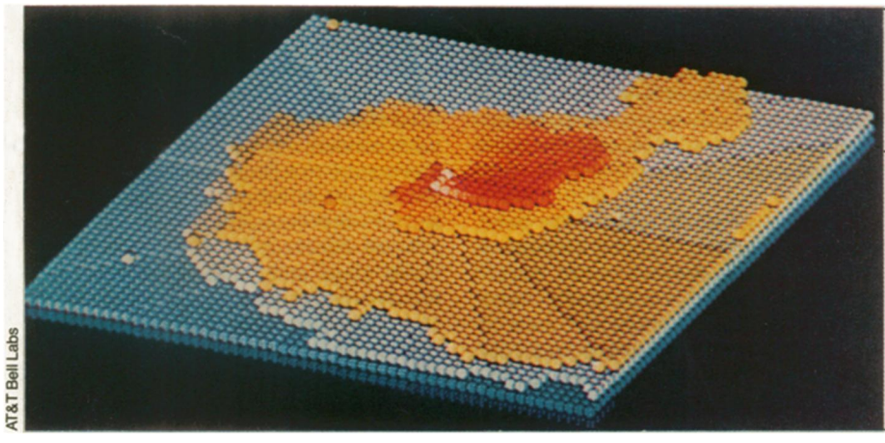
AT&T Bell Labs

A high-speed computer can simulate the particle-by-particle growth of a crystal, in this case the spiral-step pattern that occurs around a particular crystal defect.

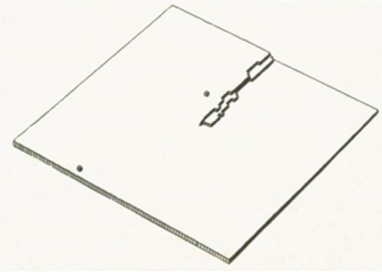

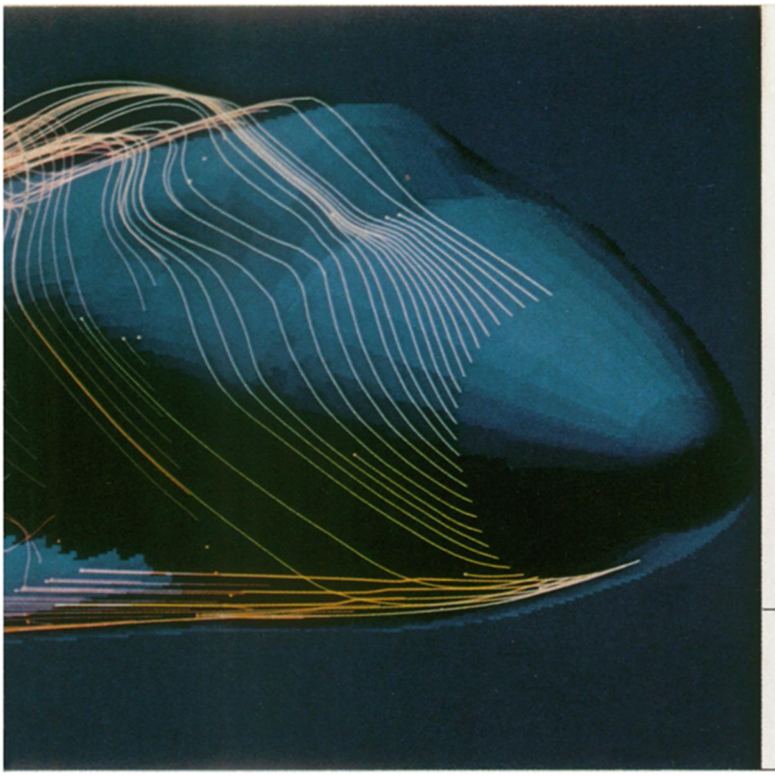

From fluid flow to crystal growth, supercomputers allow researchers to explore the details of complex physical processes

Simulations performed on high-speed computers combined with colorful graphics, in this case representing the flow of air past the shell of a space shuttle orbiter, are beginning to play an increasingly important role in aerodynamic design.

are starting to replace experiments. "This indicates supercomputers will become the wind tunnels, test tracks, materials laboratories and chemical laboratories of the future," he says, "and the adoption of this technology has major implications for U.S. science and industry."

Physicist Kenneth G. Wilson of Cornell University in Ithaca, N.Y., who seems to pop up wherever large-scale computing is being discussed, spoke of the dreams that the computer age has brought. "Imagine being able to design and then simulate the flight of a fighter aircraft, or design and test a new miracle drug at a computer terminal in just hours or days," he says. "This is the promise of supercomputers, a promise that will have to be realized if U.S. industry is to remain competitive in world markets."

But, Wilson adds, "In general, the supercomputer age in industry has not yet arrived." And, in many ways, universities are even father behind. One reason, says Wilson, is an entry-level price tag nearing $15 million to buy and support one supercomputer.

Wilson was one of the organizers of the "Forefronts of Large-Scale Computational Problems" conference, held recently at the National Bureau of Standards in Gaithersburg, Md. The sole connecting thread among the majority of the participants — among the mathematicians, computer scientists, engineers, chemists, physicists and others present — was an interest in using supercomputers to pursue particularly tough research problems.

Several of these problems came out of the chemistry laboratory. David A. Pensak of E.I. Du Pont de Nemours & Co. in Wilmington, Del., for example, highlighted the emerging field of "molecular seismology," which blends aspects of geological prospecting and chemistry. The process involves digitally "banging" one end of a computer-simulated molecule and then seeing how the energy spreads throughout the model. This gives information, for instance, about where holes (such as enzyme receptor sites) may sit within molecules. "In this way, we can look at aspects of molecules not seen before," says Pensak.

The idea originated when Du Pont chemists began thinking about how they could use the supercomputers readily available at Conoco (a recently acquired Du Pont subsidiary). There, computers were already in use for modeling where to find oil based on seismic measurements from underground explosions. "I wish I could blow up some of my molecules," said one chemist, and suddenly a light went on. "Let's blow them up numerically!"

"That's how it happened," says Pensak, and molecular seismology was born.

Pictures — computer graphics — play a very important role in this type of research. "Computers enable us to see things in our data that we have never seen before," says Pensak. "Computer graphics gets us to change how we look at a problem."

Such insights, for example, can come from simply watching on a computer display how a drug molecule actually docks at an enzyme. "You can feel as if you're wallowing around at a receptor site—coming in on a funny-looking runway with a weird airplane," says Pensak. "You can zoom in. You get a feel for what the drug sees as it nears." Understanding the geometric shape of receptor sites may allow the designing of drugs with fewer side effects, "like getting a key rather than just picking the lock" to cure a disease.

In general, he says, "the goal of computing is insight, not numbers. The computer

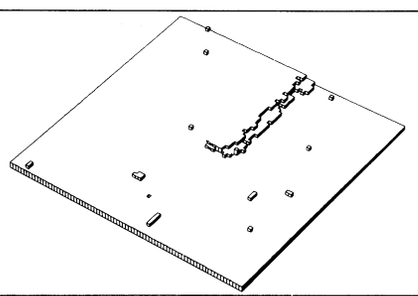

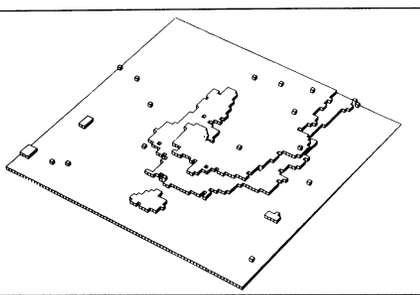

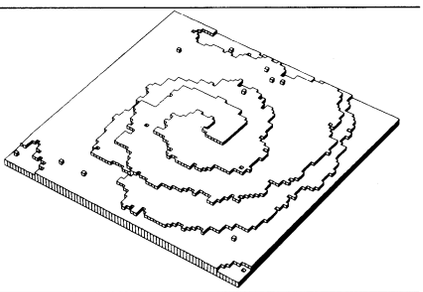

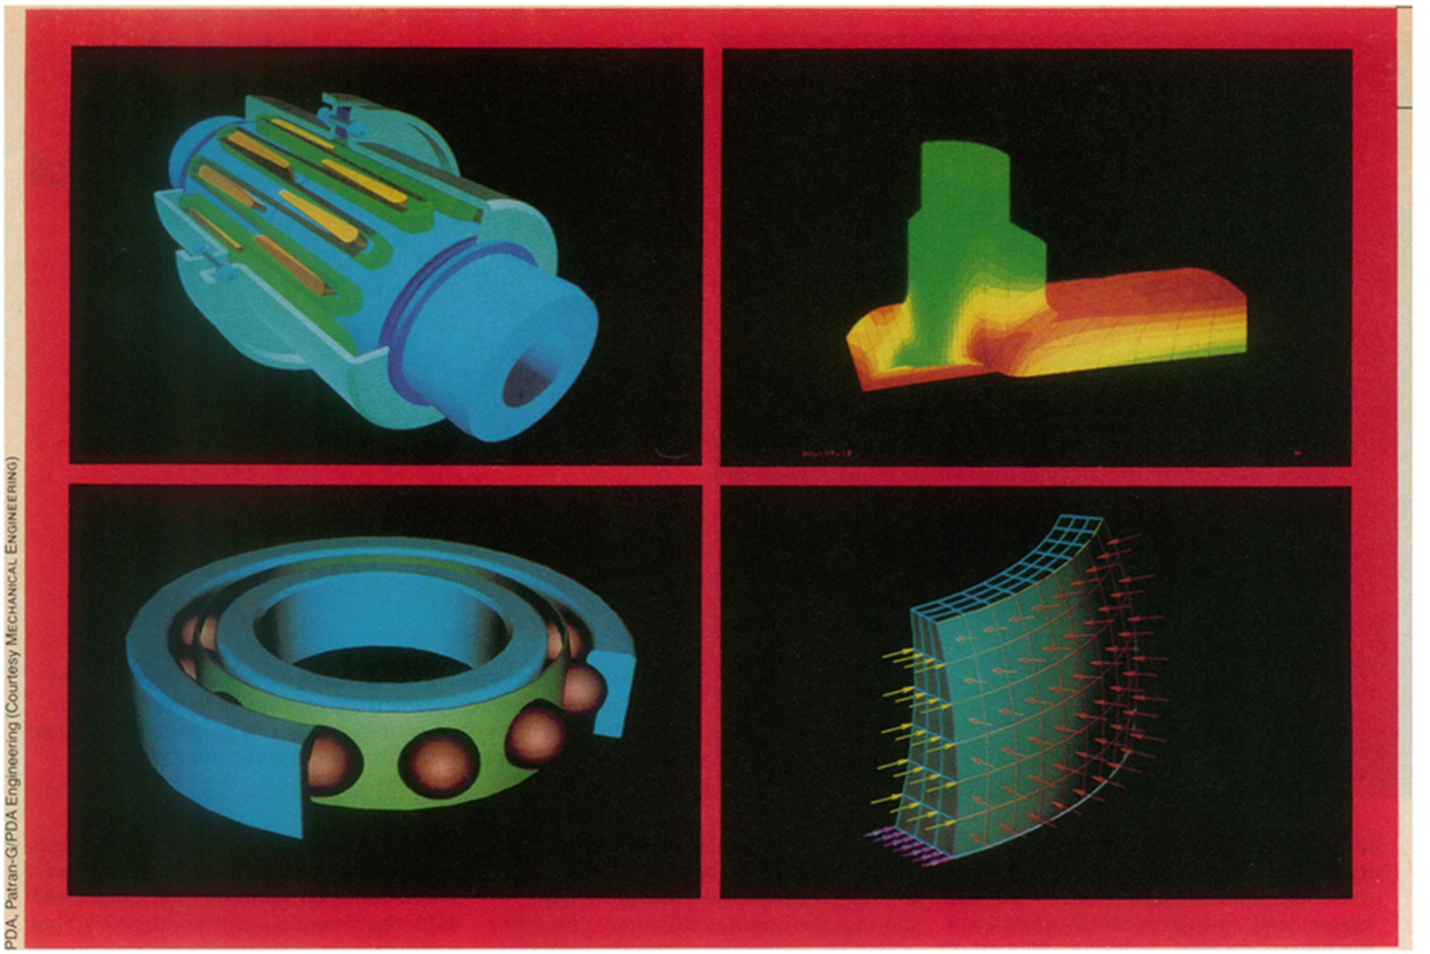

PDA, Patran-G/PDA Engineering (Courtesy MECHANICAL ENGINEERING)

is just a tool to help us." But there are limitations, too, both in the capabilities of computers and in the chemical models that are used. "We don't know enough chemistry yet," says Pensak. "We need to do a lot of basic experimental work to go with the theory."

In time, Pensak predicts, it may be possible to use computers to design membranes that can separate specific substances, such as alcohol from water. Such substance-specific membranes could save billions of dollars in energy costs, but at the moment, not even the fastest Cray computer can handle the simulation involved. Similarly, modeling completely the metabolism of a simple one-celled creature like *E. coli* requires the solution of about 15,000 simultaneous differential equations.

In quantum mechanical calculations, which attempt to derive the properties of materials from the fundamental equations governing the behavior of atoms and molecules, the largest computers can accurately handle chemical systems with not more than about 30 electrons, says biochemist Enrico Clementi of the IBM Corp. in Kingston, N.Y. As the number of electrons increases, the results become more and more qualitative. In statistical mechanics, these computers, at best, can follow the antics of about 500 water molecules (with perhaps a few other molecules thrown in to stir things up a bit). "That's really just a drop in the bucket," says Clementi.

To overcome some of these limits, Clementi is helping to design a flexible computer facility that may speed up processing considerably. The system combines conventional computer equipment in a way that allows operations to be divided into parallel streams so that several calculations can be done simultaneously. Such a "quick and dirty, home-brew" system, although expensive, is still cheaper than a new CYBER or Cray computer, Clementi says.

Despite the development of more advanced computers and computer systems, engineer John E. Bussoletti of the Boeing Co. in Seattle says, "We will always be limited by computer power." But as one conference participant remarked, "Whenever somebody thinks of a faster computer, someone comes up with a way to use it."

In Boeing's case, the problem is calculating the flow of air over airplane surfaces. Although today's supercomputers can handle very complicated shapes, the equations describing the air's motion over those objects must be simplified for the calculations to be completed within a reasonable time.

Nevertheless, "computational fluid dynamics is now a key element in airplane design," Bussoletti says, and was used extensively during the design of the Boeing 757 and 767 jet airliners. In the case of one design change, Boeing engineers had enough confidence in their computer models to order the cutting of the necessary metal parts even before wind tunnel testing had started. Nevertheless, nervous company officials were relieved when later wind tunnel tests showed that the design would really work.

In general, because wind tunnel tests are time-consuming and expensive, computer simulations provide great savings that justify the use of elaborate supercomputer facilities, says Bussoletti. A recent modification to the Boeing 737-300 took only three days to design and test on a computer versus the months it would have taken by using a wind tunnel to evaluate each possible change.

William F. Ballhaus Jr., director of the NASA-Ames Research Center at Moffett Field, Calif., says computational simulations will not completely replace wind tunnel testing in the foreseeable future. "Their roles, instead, are complementary," he says. "The wind tunnel is superior in providing detailed performance data once a final configuration is selected." Computer models help sort through a wide range of candidates, develop new design concepts such as forward-swept wings (SN: 1/8/83, p. 26) and shed light on the details of fluid flow physics. They also provide data for situations, like the high-speed entry of a space vehicle (the Galileo probe) into Jupiter's atmosphere, that are outside the range of present-day laboratories.

The emergence of computational aerodynamics during the last decade has paralleled significant increases in computer power and efficiency, says Ballhaus.

Sophisticated computer graphics can model objects and track physical processes.

Twenty-five years ago, a simulation of the viscous flow of air over an airfoil would have cost $10 million and required 30 years of time on an IBM 704 computer. Today the calculation takes 15 minutes and costs about $1,000. By 1988, new computers like the Cray-2 will allow designers to tackle a complete aircraft or a helicopter rotor by tracking what happens on a million-point grid, Ballhaus predicts. This year, NASA itself has started putting together an advanced "digital wind tunnel" as part of its new Numerical Aerodynamic Simulation program.

While NASA aims at simulating the three-dimensional flow of air around a complete aircraft in sufficient detail to permit early detection of design deficiencies, other researchers are still probing the details of how a steadily flowing viscous stream moves past a smooth cylinder. "This is the most simple flow you can imagine past the simplest body you can imagine," says Bengt Fornberg of Exxon Research & Engineering Co. in Annandale, N.J. Yet there is still much to learn about this problem.

It's hard to do experimental work, Fornberg says, partly because measuring probes often disturb the flow, especially when a detailed picture is needed. Computer simulations, despite the complexity of the equations and the amount of computer time needed, seem to be the answer. In some cases, "we are looking into a regime for which there is no experimental work," says Fornberg. "The computer is suggesting things that look completely different from what was expected."

That raises the importance of checking results when computers move into hitherto untouched research areas. "Computing shouldn't be used all by itself," says Fornberg. "Now, computers are pushing experimental work. If the flow fields we find [by doing computer simulations] look attractive, then someone may be willing to go to the expense of testing them experimentally."

"Why does every materials scientist need his own dedicated supercomputer?" asks physicist James D. Gunton of Temple University in Philadelphia. The answer lies in the need to understand the details of how materials behave. For example, it's handy to know how quickly the average electron bumps and squeezes its way through the microscopic circuits on a silicon chip. For features far thinner than a hair, even modest voltages generate large electric fields. This gives rise to nonlinear phenomena — behavior that can't be described by simple linear equations. Instead, the nonlinear equations involved are mathematically much more difficult to solve. The basic method for getting at a numerical answer is iteration: Guess a possible answer and then follow a given mathematical scheme step by step to calculate better and better approximations. Of course, this consumes a great deal of computer time.

Computers can greatly enhance understanding of nonlinear behavior often found even in simple physical systems, says mathematician Norman J. Zabusky of the University of Pittsburgh. In the July PHYSICS TODAY, he writes: "Almost everyone who has used a computer has experienced instances where computational results have sparked new insights: uncovering mistakes in derivations or calculations; suggesting when to try a new... analytic approach; or, occasionally, shining the light of inspiration into areas that had been thought devoid of new concepts or fundamental truths."

In addition, "video displays can greatly enhance the perceptibility of unexpected phenomena," says Zabusky. "The added perceptual dimensions enhance the mind's ability to recall important features and to correlate old and new results."

Such pictures play an important role in an engineering technique called "finite-element analysis." In this method, an object, such as a building's steel framework, a turbine blade or an entire engine, is subdivided into small pieces. Each piece is represented by a list of numbers that summarize its key properties. When a stress — the buffeting of a gusty wind, the force caused by spinning or the effect of heat — is applied, the computer calculates its effect on the structure. Colorful images track the resulting deformations or flows.

"Graphics is the window onto analysis," says John A. Swanson of Swanson Analysis Systems in Houston. Without graphics, there is generally too much data to see trends, he says. "Pages of printout are obsolete."

Although supercomputers, in their various guises, have potentially a wide range of applications, William Buzbee of the Los Alamos National Laboratory in New Mexico warns users that they must be careful about how they formulate and state their physics, chemistry or engineering problems. "There are many different ways of getting at a problem," he says. Not all of them need a supercomputer, and if the problem is stated badly, even the fastest computer will be of little help.

On the other hand, computer simulations can suggest possible explanations for puzzling experimental results. Chemist Raoul Kopelman of the University of Michigan in Ann Arbor, for one, is studying the rates of reactions that take place on surfaces, including reactions involving solid catalysts. Recently, Kopelman found that a quantity called the "rate constant," which characterizes how rapidly a particular reaction occurs, is, in some cases, not a constant. Instead, the "constant" gradually decreases as the reaction proceeds. For some reason, the probability of molecules reacting goes down as time passes, even when all other factors (such as concentration and temperature) remain the same. Usually such an anomalous result would have been ignored and thrown out, says Kopelman.

By studying computer simulations of some important aspects of these reactions, Kopelman was able to show that the decay in the rate "constant" is probably a natural consequence of the fractal geometry of the surfaces (SN: 1/21/84, p. 42) of the materials involved. It has nothing to do, for example, with the presence of impurities, as other chemists have suggested. Without computer simulations to model the intricate details of fractal surfaces and their effects on reactions, Kopelman says, such an insight would have been impossible.

Kopelman's initial breakthrough came when one of his former students, now a professor at the University of Crete, found that a VAX computer on the island was not being used at all during the night. Kopelman shipped computer tapes and disks to his colleague, who ran the preliminary simulations — a random walker wandering in a maze — that established the potential effect of fractal geometry on reaction rates. Appropriately, the Crete computer isn't very far from the mythical site of the labyrinth in Knossos where the bull-headed Minotaur once roamed.

However, Kopelman has had trouble finding enough computer time to do his simulations and to pursue his insight further. The university computer system is slow and overloaded, he says, and time on the few supercomputers available is scarce. Kopelman has been able to scrounge some time on a CYBER computer at the University of Colorado in Boulder, and he is now hoping to obtain some computer time through the NSF supercomputer access program.

A few hours on a CYBER really isn't enough, either, Kopelman says, because it takes time to write and test the computer program. And if he finds time available on another computer model, like a Cray, he would probably have to modify the program to fit the new computer. Such frustrations among researchers are becoming increasingly common.

"It is time to make supercomputing fully available to the theoretical science and engineering community in universities," Wilson says. "Large-scale computing is now essential for continued progress in many areas of theoretical basic research." In addition, universities should be training students to recognize opportunities for the application of supercomputers, he says. The Forefronts conference was an attempt to demonstrate the national importance of supercomputers, says Wilson, and to bring out issues involving the use of and access to these machines.

More and more researchers are looking toward the day when using supercomputers at their desks will be a routine part of their work. For Roland W. Schmitt of the General Electric Research and Development Center in Schenectady, N.Y., the use of supercomputers and the resulting simulations and graphics already represent a "third way of doing science," alongside theory and experiment. □