"Faculty from many disciplines are discovering, pretty much on their own, that many very large, computationally intensive problems are now becoming approachable, but *only* through these big, concurrent machines."

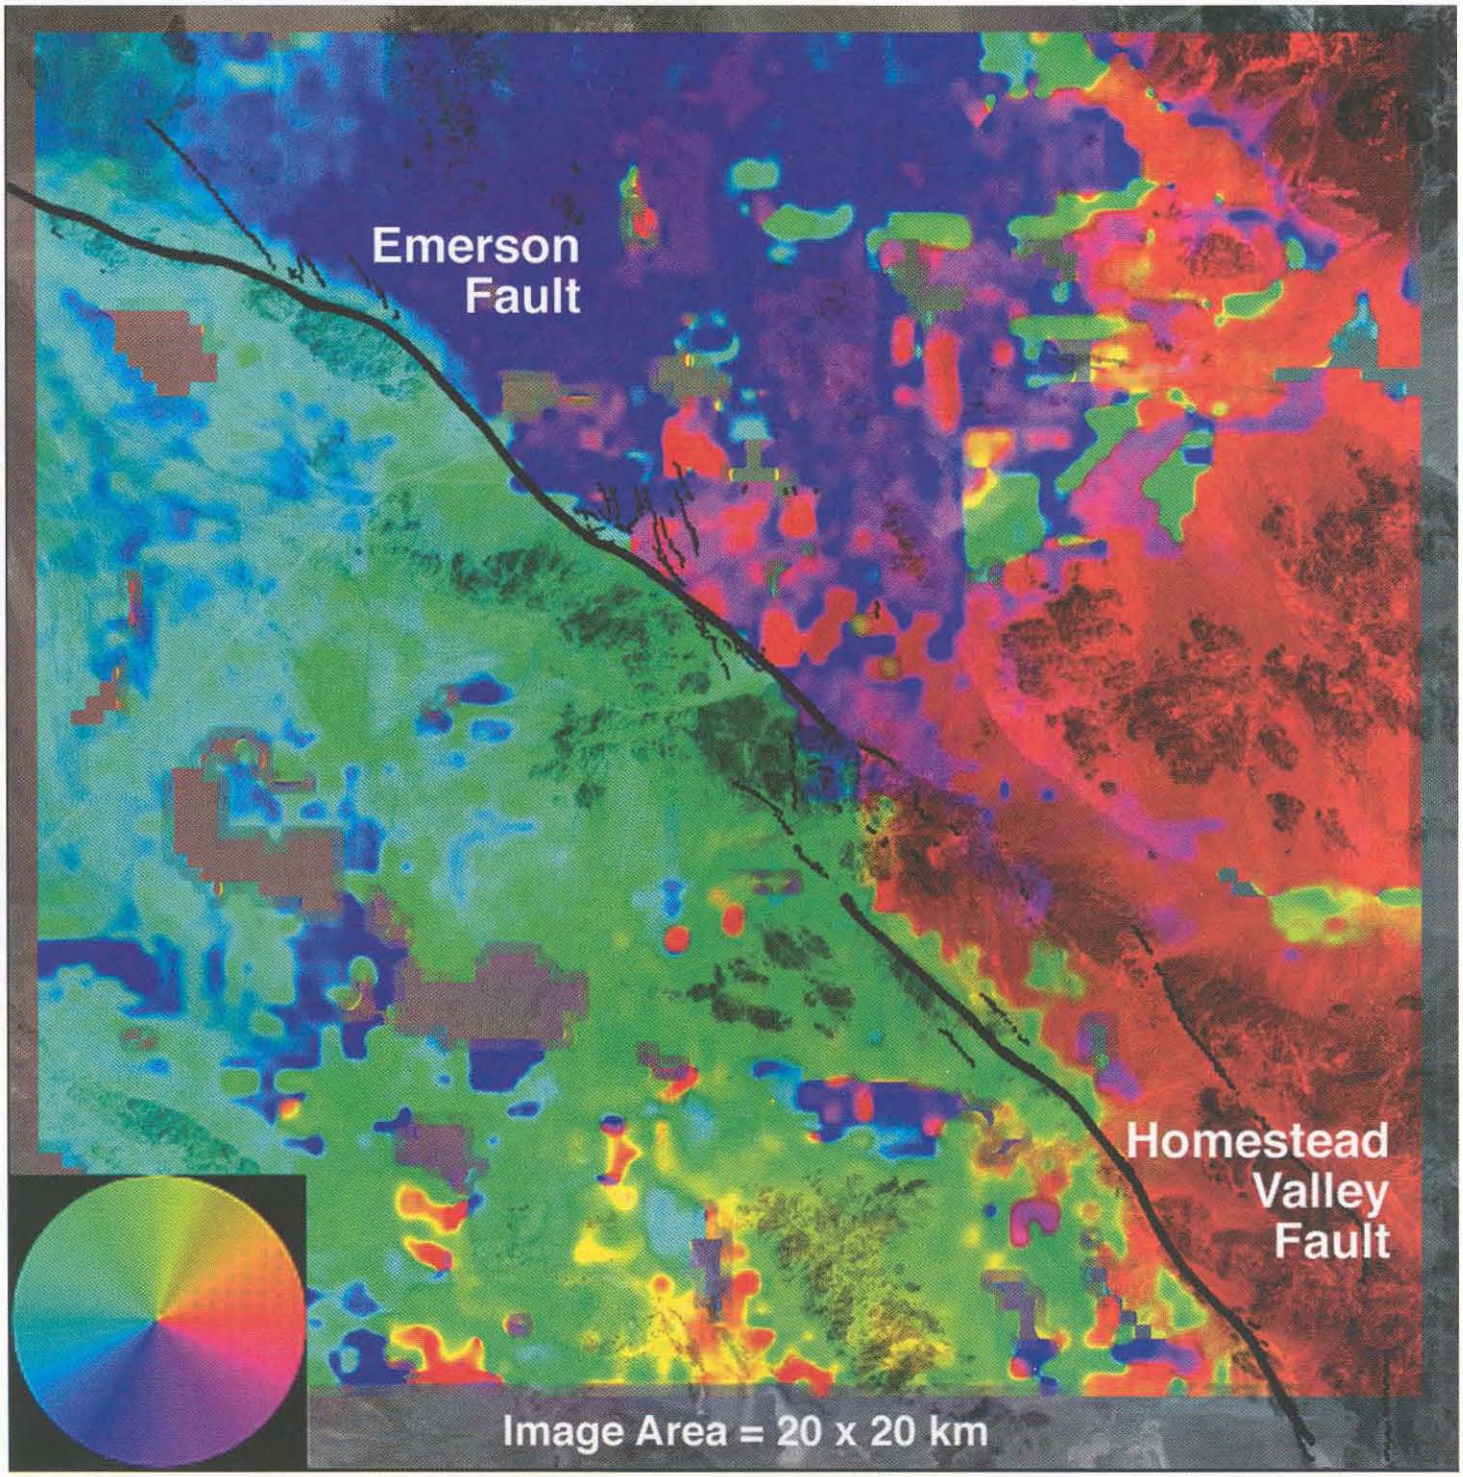

The Nodes Know

by Douglas L. Smith

The colors in this picture show the displacements caused by the Landers earthquake, as revealed by analysis of satellite photographs taken before and after the quake. (North is at the top; the faults have been drawn in as black lines.) The ground moved in the direction shown in the color wheel in the lower left corner—for example, northeast motion would be represented as yellow. A new analytical method, practical only with a supercomputer, can retrieve motions that are much smaller than the size of an individual pixel. The result is pixel-by-pixel readings of ground motions—an impossible feat for a survey team.

In the late 1960s, computers of any size were scarce, and time on the big ones was a precious resource. That's when Professor of Chemical Physics Aron Kuppermann started knocking atoms together in an IBM 360. He was modeling the simplest possible chemical reaction—hydrogen exchange, in which a hydrogen atom slams into a hydrogen molecule and replaces one of its two atoms. Even so, the computer was barely up to the task—it could only handle the model as long as all three atoms were kept in a straight line. Caltech ushered in the 1970s by acquiring a bigger IBM, the 370, and Kuppermann was keen to try hydrogen exchange in three dimensions. But in those days, the Campus Computing Center charged for computer time, and the new machine went for $300 an hour. Kuppermann didn't have that kind of money, but discovered that "there was a place in town that had a duplicate of Caltech's machine—at the time essentially the world's most powerful—and this was the Worldwide Church of God. They used it to keep track of their donors." He also discovered that they didn't use it on weekends and religious holidays. He somehow talked the church leaders into letting him use the machine during that idle time, a feat he calls a minor miracle. But the ground rules were strict—concerned for the privacy of their donors, they wouldn't let him anywhere near the machine. So on Friday afternoons, he would drop off at a little cashier's window the box of punched cards that was the weekend's calculation. On Monday morning a printout—read to be sure it contained no donor information—would be handed back to him through the same window. "If they read in my box of cards and one got mangled, that was it—a weekend lost." Kuppermann logged more than a thousand hours on that computer while finishing the world's first complete three-dimensional quantum-mechanical calculation of a chemical reaction. "Without it, we couldn't have done the work."

Thirtysomething years later, computers are ubiquitous, but the *really* big ones are still relatively rare, and Caltech still has some of the most powerful computers in the world. Unlike those IBM mainframes, which executed one instruction at a time, today's zippiest machines are "concurrent" computers that execute many instructions simultaneously. Rather than having one large, fast brain, concurrent computers have lots and lots of small or medium-sized brains, called nodes, working in parallel. Each node chews on its own piece of the problem independently, passing information back and forth to its fellows as needed. This approach, developed in the early 1980s at Caltech and elsewhere (*E&S*, March 1984), is ideal for attacking problems where the same basic calculations have to be applied to a large set of elements—be they stars in a gravitational field, grid points in a wind-tunnel flow, or eigenvalues describing the quantum state of a nuclear particle. Caltech's two largest machines, the Intel Paragon and the Intel (Touchstone) Delta, have 512 nodes apiece. Each node is about 10 times as powerful as that old IBM 370. This bequeaths a huge increase in computational speed—the 370 could execute a million floating-decimal-point calculations per second; the Paragon does ten billion. Meanwhile, says Paul Messina, assistant vice president for scientific computing and director of the Center for Advanced Computing Research (CACR), "faculty from many disciplines are discovering, pretty much on their own, that many very large, computationally intensive problems are now becoming approachable, but *only* through these big, concurrent machines."

These faculty are in the vanguard of a quiet revolution in the way science works. Because concurrent machines can handle complex models jam-packed with minutiae, the monogamous relationship of theory and experiment has become a ménage à trois of theory, model, and experiment, where results from any one can spur developments

in the other two. The availability of so much raw computational power has also led to "data mining," the sifting of huge mounds of information for correlations that would be missed if it weren't possible to crunch all the numbers six ways from Sunday. In the environmental, earth, and planetary sciences, for example, data sets from different satellites can be amalgamated for exploration. (This has been eased by the development of high-speed data links between computers, and Caltech has been involved in that, too, but that's another story.) Here are some dispatches from the front lines of that scientific revolution.

Professor of Physics Thomas Prince and his colleagues are working to integrate a supercomputer into a radio telescope—essentially making the computer just another panel on the instrument rack. Normally, the incoming data is piped through a bank of amps and filters to pump up, say, the blip-blip-blip of a distant pulsar while hopefully squelching the local interference from radars, TV stations, and what have you before being recorded on tape for later analysis. "But you're limited in how much data you can record," says Prince. "So you need to throw out a lot of information. Ideally, you'd like a digital receiver, but this would take as much computational capacity as the Delta." Such a receiver would amplify, filter, and analyze the signal in real time. The analysis requires small slices of time to measure the pulsing of the pulsar, as well as small slices of frequency over a very broad range so you can study the pulse's shape. All this generates data at a prodigious rate. To further complicate things, the signal is smeared out by its passage through the plasma clouds (clouds of charged particles) in our galaxy, and so a pulse's frequency components will arrive on earth at slightly different times. The analysis has to "de-disperse" the pulse to reconstruct its original form, but the average workstation can only handle a few seconds worth of signal at a time. It would be like trying to fill a demitasse cup with a fire hose.

Since both the Delta and the Paragon live on campus, hundreds of miles from the nearest radio telescope, Prince's group designed a custom chip to do the next best thing—digitize as much of the raw signal as possible and record it on a very fast tape deck. The faster the tape, the wider a chunk of the spectrum you get, and the more accurate the time information. Then the Paragon could pretend to be a digital recorder while listening to the tape. Datatape Incorporated donated the tape decks for the telescope and the computer to CACR, and collaborated in their installation. These babies whir at 50 megabytes (which would fill 36 high-density floppy disks) per second—the fastest recorders then available commercially. After being tested on the 40-meter dish at Caltech's Owens Valley Radio Observatory, the system got its first real workout on the 64-meter dish at Australia's Parkes Observatory in July 1995, recording 10 terabytes (the equivalent of all the Library of Congress's printed holdings) of data from an assortment of objects, including globular clusters such as 47 Tucanae. This took 100 cassettes somewhat larger than VHS tapes—a bargain considering that the equivalent stack of floppies would stand 14 miles high.

Prince's group focused on 47 Tucanae, which is known to have at least a dozen very fast pulsars, in hopes of finding a record-breaking rotation rate. (A pulsar is a rapidly spinning neutron star—the cinder left after a dying star explodes into a supernova—that emits a powerful radio beam along the axis of its magnetic field. If the star's magnetic and rotational axes don't line up, the beam sweeps through space like the beacon from a cosmic lighthouse. If the earth lies in the path of the beam, we see a radio pulse—hence the name—

This took 100 cassettes somewhat larger than VHS tapes—a bargain considering that the equivalent stack of floppies would stand 14 miles high.

with every sweep.) The fastest known pulsar spins once every 1.5 milliseconds. This is approaching the speed limit, believed to be somewhat less than one millisecond, where matter on the surface of the neutron star (which is about 10 kilometers in radius) is whirling around the axis at nearly the speed of light. (And you thought the teacups at Disneyland made you dizzy!) The search came up empty, but the tapes contain lots of other interesting data that's still being analyzed.

At the opposite end of the cosmic distance scale, Professor of Theoretical Physics (and vice president and provost) Steven Koonin (BS '72) has been using the Delta, the Paragon, and other machines to look inside atomic nuclei. You can think of the energy levels within a nucleus as a ladder of infinite height. Each rung is wide enough to hold two nucleons—two protons, two neutrons, or one of each—and each particle seeks out the lowest unoccupied rung. But the particles continuously interact with one another through the strong nuclear force, kicking each other up and down the ladder, creating and filling vacancies in the low-energy rungs. Thus, to find the overall configuration of the nucleus, you set up a table, or matrix, that lists all the interactions between every possible configuration. Calculating the strengths of these interactions—diagonalizing the matrix, it's called—gives you all the possible energy states and their probabilities of occurring. This works

fine for light elements such as carbon, with 12 nucleons, but things rapidly get out of hand. The mere 28 nucleons in silicon have about 100,000 possible energy states; zinc, with 60 nucleons, about 30,000,000; and the rare earth dysprosium, bustling with 166 nucleons, has a staggering 10^{21} (one sextillion) possible states—a number that would make even a federal economist blanch.

Diagonalizing a dysprosium-sized matrix in its full glory remains out of the question, but a few years ago Koonin, postdoc Erik Ormand, and grad students Calvin Johnson and Gladys Lang realized that all they really needed was a statistically valid sample of the matrix's energy states. By using this technique, the heaviest nucleus calculated has jumped from 35 nucleons (chlorine) to 76 (germanium)—a very considerable leap, but one that leaves most of the periodic table yet to go. Meanwhile, Koonin's group has applied the method to two long-standing theoretical problems. One has to do with the Gamow-Teller process, in which a proton, upon swallowing an electron, gags, coughs, flips its spin, loses its charge, and turns into a neutron—a critical step in the "neutronization" of iron that leads to the production of heavier elements in the fiery belly of a supernova. The rates of those Gamow-Teller processes that have been measured experimentally are less than 30 percent of what was predicted by previous calculations, but these simulations are right on the money. The other problem concerns double-beta decay—an event so rare that it takes 10^{20} years to happen to the average nucleus. This is unfortunate for people watching for it, as the universe is only about 10^{10} years old, but it's a good thing for the rest of us—if the nuclei were decaying much faster, we wouldn't be here to watch them do it. For reasons too discursive to go into, the rate of double-beta decay is linked to the mass of the neutrino, which is the most evanescent of particles and is generally assumed to be massless. If the neutrino proves to have even an infini-tesimal mass, it will make the universe a much heavier place, and one more likely to eventually recollapse on itself in the so-called Big Crunch. Again, previous calculations of double-beta decay rates were seriously inconsistent with reality, but the calculations by Koonin's group matched the data nicely.

The group is now moving on to the bizarre properties that heavy nuclei exhibit when spun and heated. At room temperature and with a normal nuclear spin of "only" 10^{19} revolutions per minute, a nucleus can be spherical. But heated to several million electron volts (about 10^{10} degrees Centigrade) and cranked up to 10^{21} rpm, it can become football- or even cigar-shaped. It can even flatten, like a cigar that's been run over by a truck. Such nuclei have actually been observed experimentally, but the theory to describe them has been lagging behind. But now, says Koonin, we can rev them up in the computer and see how their structures change.

Paul Stolorz (PhD '87), technical group supervisor of the machine learning systems group at JPL, is also watching structures change, but his are geological. Stolorz and then-undergrad Chris Dean, along with JPL geologists Robert Crippen and Ronald Blom, have refined a method Crippen and Blom invented that compares a pair of before-and-after pictures and computes the motion of each pixel to within 10 percent of its width—a form of data mining. For example, pictures of the Mojave Desert taken by the French SPOT satellite before and after the Landers earthquake show no visible changes, because the quake's maximum offset was only six meters and each pixel covered a 10-by-10-meter chunk of desert. But the collaboration was able to measure displacements to within a meter by taking a block of pixels (100 by 100 in this case) from the "before" picture and sliding them around on the "after" picture to get the best match. This gave the motion of the block's central pixel very accurately, because the individual matching error for any pair of pixels got buried in the statistics of the whole block. The process was repeated for similar blocks centered on each of the 40,000 pixels in the overlapping image area. This comparison, and its preliminary image processing—subtracting out the motion of the spacecraft and compensating for the different look angles—took JPL's Cray T3D, which has 256 nodes, the equivalent of 24 hours of uninterrupted work. It would have taken months to do on a workstation.

The rare earth dysprosium, bustling with 166 nucleons, has a staggering 10^{21} (one sextillion) possible states—a number that would make even a federal economist blanch.

The collaboration is now reexamining old Mariner 9 and Viking Orbiter photos of Mars for possible evidence of sand-dune motion. They will also be comparing Voyager's images of Jupiter's moon Europa with the upcoming Galilean ones (the first of Galileo's three close flybys of Europa will be on December 19, 1996) for evidence of surface motions that would betray the presence of liquid water beneath its crust of fractured ice. (See page 7 of this issue.) But the method is applicable to any image composed of pixels, even medical MRI scans.

Stolorz's group, in collaboration with biochem-

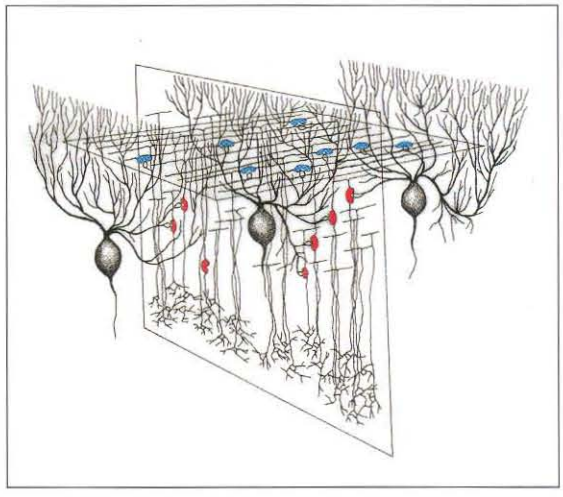

Right: The colors in this model of a single Purkinje cell represent its electrical potential partway through the firing process. The redder a portion of the cell is, the more recently the impulse moved through it. The yellow and green portions of the dendrite tree are returning to their resting state, and the blue parts are fully recovered and ready to transmit the next impulse. Far right: The relationship between the Purkinje cells (the beet-like objects) and the granule cells below them. As the granule cells' axons ascend, they make connections (red) with nearby Purkinje cells. Once the axons split, the parallel fibers make additional connections (blue) to Purkinje cells.

ists at the University of Vienna, Los Alamos National Laboratory, the Santa Fe Institute, and the University of Illinois, is doing some biology already—calculating how a piece of RNA will fold, based on the sequence in which its letters are strung. RNA is a wonderfully versatile molecule. In higher organisms, RNA is the messenger molecule that takes DNA instructions to the cellular machinery, but in some viruses, including HIV, the RNA is all there is—it both stores the genetic instructions and carries them to the host cell's co-opted machinery. And in all organisms, RNA plays a variety of roles in the machinery itself.

Like DNA, there are four letters in the RNA code, and the letters pair up with one another in only one way. So the individual calculations of which letters will willingly pair up are very simple, but the number of overall calculations increases with the cube of the number of letters being considered, and the memory required to juggle the calculations increases with the fourth power. Other machines had hit the ceiling at about 3,000 letters, but the Delta plowed through all 10,000-odd letters in an HIV virus—the first time a predicted structure of such length had ever been computed. Stolorz's group has since folded all of the other 19 known strains of HIV to determine the similarities and differences between them. The structures that recur from mutation to mutation are obviously important to the virus's survival, and analyzing them might suggest how the virus could best be attacked. For example, one of the biggest hurdles to developing an HIV vaccine has been the rapidity with which the virus mutates, making vaccines developed against one strain ineffective against others. Locating features that remain constant from mutation to mutation might give the vaccine designers a fighting chance. And understanding how the other features differ could illuminate the various ways the virus attacks us. "These calculations are just best guesses," Stolorz emphasizes. "It's very difficult to verify these kinds of structures experimentally. But the things people do know are consistent with our calculations." And viruses three times the size of HIV, which would include Ebola's 20,000 letters, appear to be within the Delta's reach.

Associate Professor of Biology James Bower and his colleagues have cast their eyes rather higher up the evolutionary ladder—they are looking at mammalian nervous systems. Says Bower, "These are realistic models, based on anatomy and physiology, that require solving thousands of differential equations at once." The models can reflect many levels of detail: subcellular, single-cell, and multicell. On the subcellular level, for example, the equations describe the way ions diffuse across the cell membrane as the cell fires; how the proteins that make up the ion channels through that membrane are flexing to open and close them; and how the cascade of kinases—the signaling proteins that control the process—are behaving. Zooming the microscope out a bit, Bower's group has constructed detailed models of single neurons that allow the researchers to try to understand how the cell's complex structure relates to its function. And finally, cells can be linked together into networks, and even systems of networks, to explore how the brain's circuits work. These models were created using GENESIS (GEneral NEural SImulation System), a versatile software package Bower and his group developed a decade ago that is now used worldwide.

While GENESIS will run on any machine, even a workstation, only recently have supercomputers become powerful enough to exploit the software's potential. "A closed solution of an analytic problem means you know everything that's going on in the system," says Bower. "But in biology, you have thousands of variables, so all the solutions can't be known. Instead, understanding how nervous systems work will increasingly depend on numerical simulations that explore a complex set

of parameters." In other words, you run the model over and over and over again while you twiddle various parameters to build up a picture of how the system behaves. The models are compact enough to put a copy (sometimes several copies) on each node of the Delta or the Paragon, but running them often enough (say 500 runs) to get meaningful statistics soaks up machine time—the equivalent of years on a workstation.

The models' compactness belies the level of detail they include. For example, the Bower group has been studying a model of a Purkinje cell—a kind of neuron that lives in the middle layer of the cerebellar cortex and is involved in sensorimotor coordination. The Purkinje cell receives inputs through its enormous dendrite, which looks a bit like the Engelmann oak by Millikan Library—a spreading mass of gnarled limbs and branches. The model, created by then-postdocs Dieter Jaeger, now at Emory, and Erik De Schutter, now at the University of Antwerp, reconstructs this tree in GENESIS by linking together 4,588 electrically distinct compartments, each of which has a characteristic set of up to 10 different types of ion channels. When the cell is stimulated, the flow of ion currents along the dendrite can be "seen" directly as each compartment's channels open and close in succession.

The model has suggested that the cell works in a fundamentally different manner than previously assumed. A Purkinje cell has some 175,000 sensory inputs, so it's literally bathed in continuous, random stimuli. Some are excitory and make the cell want to fire; others are inhibitory and tell the cell to mellow out. Excitory stimuli come from the granule cells, which the model suggests may have two very different effects on the Purkinje cells. The granule cells are located beneath the Purkinje cells, and lie in the densest cell layer in the mammalian brain (a staggering 6,000,000 cells per cubic millimeter). The granular cells' output fibers, called axons, ascend out of the granule-cell layer and make up to 100 connections with the immediately overlying Purkinje cells. Each axon then splits, forming two so-called parallel fibers that run in opposite directions to each other, but parallel with every other granule cell's parallel fibers. The parallel fibers pass through the Purkinje cells' dendrites like telegraph wires through overgrown trees along a badly maintained right-of-way, with each fiber making at most one connection with each Purkinje cell it encounters. For 100 years, it has been assumed that the connections between the parallel fibers and the dendrites were the important ones. But the simulations have suggested that the synchronous activation of the connections on the ascending part of the granule cell's axon may have a much more profound effect on the Purkinje cell than the sequentially activated parallel-fiber connections. Within months of Bower's lab obtaining this result, corroborating physiological data was independently published by another lab.

While some folks are modeling processes that occur in biological brains, others are modeling processes that are used to build silicon brains. Professor of Theoretical Chemistry B. Vincent McKoy and collaborators are studying the collisions of electrons and chlorofluorocarbon gases (CFCs), a process critical to manufacturing computer chips. When you hit a CFC molecule with an electron, you can break off a chlorine or fluorine atom—voracious fellows that eat silicon and practically everything else, and which are used to etch the electronic circuits onto the chip. But to keep the process under control and avoid frying the chip, the reaction has to be run at relatively low temperatures (say, 80° C). Nonequilibrium, low-temperature plasmas are well suited for this, because the electrons in the plasma are energetic enough to shatter a CFC molecule upon collision with it. There's surprisingly little experimental data on how CFCs disintegrate under these conditions, so designing the etching systems has relied extensively on intuition and experience. "We wanted to model the process in such a way that it relates electrons and CFCs to the composition of the plasma and to the etching results," says McKoy. "Then you can design the process more intelligently."

Of necessity, the model started from first principles, which in chemistry means the dreaded Schrödinger equation.

Of necessity, the model started from first principles, which in chemistry means the dreaded Schrödinger equation. This equation in principle provides a complete, quantum-mechanical description of the motion of any molecular system, but there's a catch—as soon as more than a few particles are involved, the equation gets fantastically complex and incredibly difficult to solve. There was a way out in this case, however—because the incoming electrons whiz by the atomic nuclei in 0.1 to 0.01 femtoseconds (a femtosecond is 10^{-15} seconds; it takes about one second for light to travel from the earth to the moon, while in one femtosecond light traverses 1/100th the thickness of an eyelash), the nuclei don't have time to react and could thus be treated as stationary objects. The calculation fires electrons at the CFC molecule over and over again at various speeds and angles, and tracks how the molecule's electron cloud quivers under the blows. Where the electrons end up then tells the researchers how the molecule will

Below: At low solvent pressures (top), the reacting molecule (red) can spin unimpeded by the solvent molecules. But as the pressure goes up (bottom), the reacting molecule gets hemmed in.

Right: The experimental version, done with iodine molecules dissolved in helium or argon. The plot shows the iodine–iodine bond stretching and shrinking, with a period of 300 femtoseconds. At high densities (bottom), the vibration dies out rapidly. But at low densities (top), the vibration persists long enough for the individual bonds to drift slowly out of, and back into, phase—just as two notes at very nearly the same pitch will "beat" against each other.

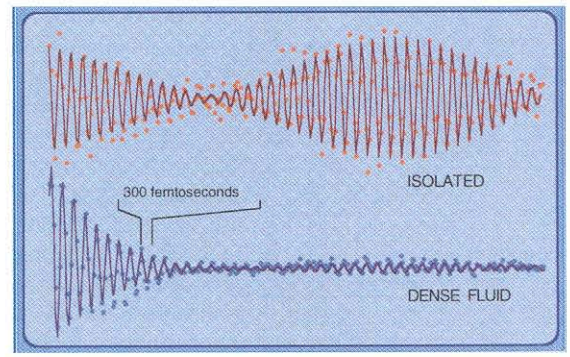

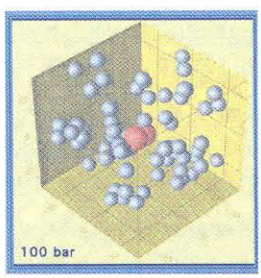

fall apart. Even so, as of a few years ago, the only electron-collision processes that had been modeled involved very simple gases such as nitrogen and carbon monoxide. More complex molecules just had too many electrons and too little symmetry (symmetrical electron clouds are easier to calculate).

But with the advent of the Delta, McKoy, Senior Research Fellow Carl Winstead, and grad student Howard Pritchard (PhD '94) were able to calculate the electron-collision probabilities for boron trichloride, a common etchant, within a few months. The group has now moved on to bigger game, and is almost finished calculating C_2F_6 and C_4F_8 as part of the first year's work in a longer-term project to model the electron-collision probabilities of important etchant gases. This project, sponsored by SEMATECH (the national research consortium of semiconductor manufacturers), is one component of an ambitious plan by McKoy, Assistant Professor of Computer Science Stephen Taylor, and Sadasivan Shankar of Intel to simulate the low-temperature plasma etching process, in which the breakup of the etchant gas is merely the first step. Such simulations will provide the foundation for improved computer-aided design tools that could reduce the cost of developing the next generation of etching equipment.

Moving from the gas to the liquid phase, Ahmed Zewail, Pauling Professor of Chemical Physics and professor of physics, is examining chemical reactions in a solvent. It's long been known that the solvent molecules play an important part in determining the course of the reaction, but there were just too doggone many of them to simulate them in any detail until recently. Zewail, grad student Qianli Liu, and research fellow Chaozhi Wan took up the challenge in an attempt to explain some very odd results that Wan had been getting in experiments he had been doing with the group's femtosecond-spectroscopy apparatus (*E&S*, Spring 1988), which "freezes" a reacting system in slices of time a few femtoseconds wide and allows a "movie" to be made of a reaction as it occurs.

One expects a reaction rate to increase smoothly with the number of atoms per unit volume. After all, most reactions are molecular fender benders, happening when one atom collides with another, so it makes sense that they'd happen more frequently in crowds than out in the wide-open spaces. But Wan was finding instead that the rate remained constant over a pretty wide range of densities—from five to 15 atoms per cubic nanometer, in fact. (By comparison, nitrogen at room temperature and pressure has .05 atoms per cubic nanometer; methanol has 26; water 33.)

So Liu used the Delta to look at the 500 solvent molecules closest to the reaction, and calculated their motions in relation to the reacting atoms as the reaction proceeded. Because these simulations focused on the atomic nuclei, which move much more slowly than the electrons, the calculations could assume that the electrons were keeping up and sidestep the Schrödinger equation by using non-quantum-mechanical approximations for the atoms' motions.

The simulations, run at pressures ranging from zero up to liquidlike densities, showed that the reaction rate did in fact increase smoothly, but only up to the point where the molecules began to need more elbow room. In a very diffuse solvent, the simulation showed the reacting molecules were rotating very rapidly, as everyone knew. But what hadn't been appreciated was that the centrifugal force from this rotation stretched the bond that was eventually going to break, priming the reaction. The incoming atom didn't have to slam the rotating molecule too hard to knock the outbound atom free. But as the neighborhood got overcrowded with solvent molecules, the reacting molecule kept whacking into the solvent molecules and couldn't spin freely. With-

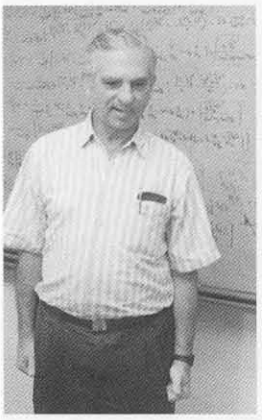

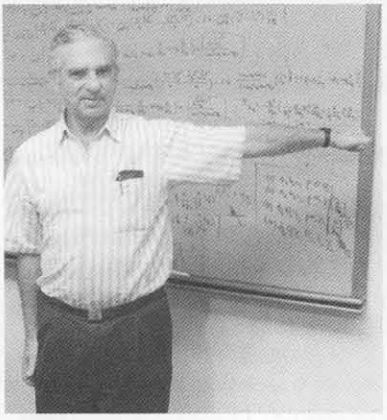

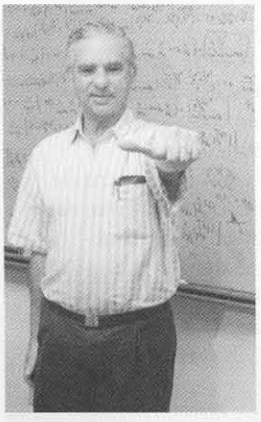

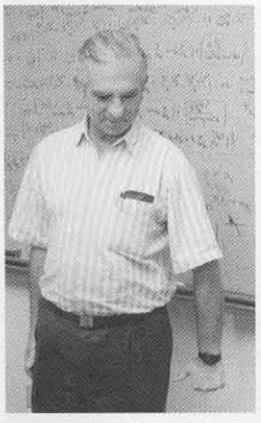

out the centrifugally induced bond stretching, the incoming atom had to hit the reacting molecule harder. These two effects—centrifugal stretching and collisional slowdown—neatly balanced each other over the density range of five to 15 atoms per cubic nanometer.

There's a popular saying among mathematicians that physics is just applied math, chemistry is just applied physics, and biology is just applied chemistry. Professor of Chemical Physics Aron Kuppermann must be a mathematician at heart, as he hopes to eventually derive a large amount of chemistry from the Schrödinger equation. Because molecule–molecule collisions are about 1,000 times slower than electron–electron collisions, Kuppermann can't ignore the nuclei's motions the way McKoy can, but he can still simplify the calculations by treating the electrons and the atomic nuclei independently because the electrons move so much faster. First, Kuppermann "freezes" the nuclei, and calculates the motions of the electrons for every possible nuclear configuration. This gives the so-called potential energy surface, which shows how much energy is stored in those configurations. The system will try to minimize its potential energy, so a second pass with the Schrödinger equation gives the motion of the nuclei when subjected to the forces described by the potential energy surface. Because these forces are more complicated than simple electron–electron repulsion, the equations have to be solved numerically rather than analytically—in other words, you run them over and over again while plugging in lots and lots of numbers. Thus it has taken Kuppermann nearly 25 years to progress from hydrogen exchange ($H + H_2 \rightarrow H_2 + H$) to deuterium exchange ($D + H_2 \rightarrow H + HD$), which he and Member of the Technical Staff Yi-Shuen Mark Wu (PhD '92) published in 1993. (Deuterium, or heavy hydrogen, has a neutron as well as a proton in its nucleus.) Why bother? Says Kuppermann, "If you can do something very, very well—even a system as simple as $D + H_2$—it will tell you things that can be applied to much larger systems."

Both the $H + H_2$ and the $D + H_2$ systems, in fact, revealed the existence of a "geometrical phase effect," which has to do with the path that the atoms take to get to their final positions. Like those awkward moments where someone comes toward you, hand outstretched, and then shakes hands with the person behind you, atom A can approach atom B and then veer off at the last moment to embrace atom C instead. Experiments on the $D + H_2$ system by Richard Zare at Stanford have since confirmed that the effect does indeed exist, and can influence the reaction's outcome by a factor of 10.

By 1995, Kuppermann and Wu had moved on to $H + D_2 \rightarrow HD + D$—a chemically trivial but computationally intense leap. This brought a fresh surprise: something called dynamic resonance. In the transition state midway through the reaction, the three atoms are all stuck together: H–D–D. Normally, one bond shrinks while the other one stretches, and the middle D slides from the end D to the H like a bead on a bit of string. Two of the atoms are snuggling and one is estranged—from there it's but a short step to splitsville. But the bonds can also stretch and shrink in unison, like a bodybuilder flexing the pecs. This vibrational mode is unreactive, because the two outlying atoms are equidistant from the center at all times, and thus one is no more likely to leave than the other. The two modes have different energies, so if the incoming hydrogen happens to have nearly the energy of the symmetric mode, the transition state will tend to vibrate in that mode. The effect is to make the threesome "stickier," and the complex lasts for 150 femtoseconds instead of the usual 10 before the transition state slips into the asymmetric mode and falls apart. For reasons too esoteric to go into, this actually increases the probability of the reaction occurring at that energy.

And now, Stephen Walch (PhD '77) of NASA Ames has finished calculating the potential energy surface for the reaction of oxygen and hydrogen in collaboration with Kuppermann's group. Grad student Stephanie Rogers is using the Paragon to explore the reaction on that surface, a project Kuppermann estimates will take six months to a year to finish. Knocking a molecule of H_2 and an atom of O together to give OH and H is an important combustion reaction in, among other things, the propulsion of the Space Shuttle. However, this reaction doesn't occur just like that. Some encounters are ricochets, some are direct hits; the outcome depends on who hit what where, how hard, and at what angle. Like the chaos of a disturbed anthill, most of this activity goes

A macroscopic example of a geometrical phase effect. Kuppermann's thumb and wrist remain rigid in relation to his forearm—in other words, the thumb's position in a local coordinate system based on the forearm remains unchanged. Yet as Kuppermann rotates his arm, the local coordinate system rotates in relation to a global coordinate system, and the thumb rotates in relation to Kuppermann's body.

Right: A slice through a 512 × 512 × 512 three-dimensional "turbulence in a box" simulation by Mei-Jiau Huang (PhD '94). The colors indicate vorticity, with blue being the lowest and red the highest.

Below: Henderson's three-dimensional, spectral-element simulation of the wake behind a cylinder (the black bar at left). The sequence of images shows a slice through the wake along the cylinder's axis. The colors represent velocities, with yellow close to zero. As the color changes from green to blue, the fluid is moving with increasing speed into the plane of the page; similarly, red and white indicate the fluid is coming out of the page. At low speeds (top panel), the wake is very regular, but at a critical speed (second panel), a mode of instability (the scalloped pattern) develops. As the flow speed increases further (third panel), additional modes of instability with shorter wavelengths can be seen, particularly in the two sets of bands closest to the cylinder. At still larger flow speeds (bottom panel) there are so many unstable modes that the bands themselves begin to break up.

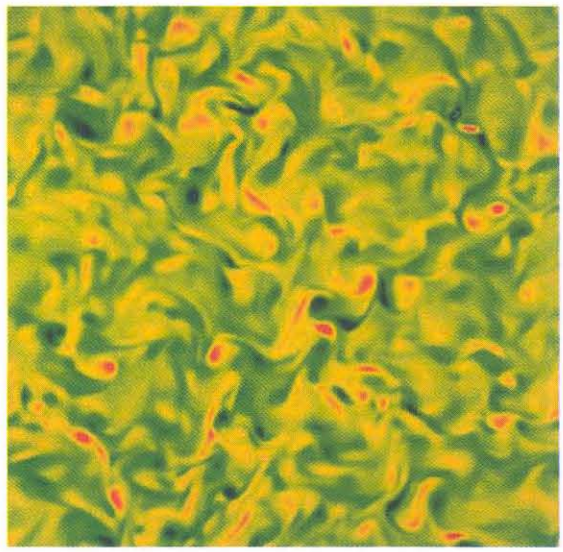

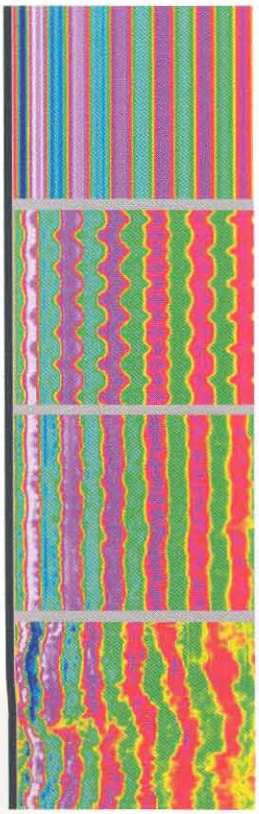

nowhere or is counterproductive. But the calculation has to include every possible collision and breakup mode, because it's the differing rates at which all these processes operate that determine the rate of the reaction. This puts such a load on the computer, says Kuppermann, that "the brute-force approach doesn't work. These are problems you can just barely solve by being clever."

Kuppermann plans to model a number of three-atom reactions to see what the range of behavior is. And, with postdoc Desheng Wang, he's preparing for the arrival of the next generation of concurrent computers by developing methods to do the same with four-atom reactions. This, he hopes, will provide enough information to create a generalized model that will predict the course of almost any reaction, because "most chemistry, including biology—even when you're dealing with proteins—really boils down to three or four atomic centers with one or two bonds breaking and forming at once. The rest is mainly flexible scaffolding."

Computational fluid dynamics is another field where the problems can barely be solved by being clever. In the late 1940s, John von Neumann, the father of the modern computer, proposed simulating turbulent flows on a three-dimensional grid. He noted, however, that 20 points on a side would require a whopping 8,000 data points, and that no machine "in sight for several years to come" could do the job. This "turbulence-in-a-box" school of modeling grew with the computers, and today's Delta or the Paragon can handle a 512 × 512 × 512 grid—134 million data points. But the faster an object moves, the smaller the tiniest eddies behind it are. A 512^3 grid suffices for walking-speed turbulence, says Professor of Aeronautics Anthony Leonard (BS '59), but the smallest swirls around a golf ball in flight would fall through the mesh. To calculate the microscopic moil around a jet aircraft in full detail, Leonard says we need a grid that's 2,500 points on a side and a machine he reckons is about 15 years away.

But you can wring more detail from existing machines by using cleverer approaches. One, called the spectral-element method, was developed by Senior Research Fellow Ron Henderson while a graduate student at Princeton. The method approximates the equations of fluid flow with polynomials and can simulate turbulent flows on irregular grids. This makes it possible to handle a much wider range of scales than would a uniform grid with the same number of grid points. The method is being widely used for very detailed simulations of turbulent flow past simple bodies like cylinders, but the flow speeds are still modest—Henderson's simulation of a fully turbulent flow at a speed equivalent to a six-inch putt is the best anyone has yet managed.

Another approach, called the vortex-particle method, is also making fast progress. Here each grid point encodes vorticity—the rate and direction of spin of a fluid particle at that point. Thus the grid only covers the regions where something is happening, which saves the enormous amounts of computation that would be wasted on the smooth-flowing parts of the flow. Even so, there's room for improvement. The original method, developed in the 1970s by Leonard while at NASA's Ames Research Laboratory, calculated every particle's effect on every other particle. The number of computations increased with the square of the number of particles, severely limiting the grid size. But in the late 1980s, then–grad student John Salmon (PhD '91) figured out how to make the calculation grow more slowly by reorganizing it into groups of distant elements that didn't affect each other strongly, and could therefore be approximated, and pairs of nearby elements that still had to be calculated in detail. (Salmon actually studied how sheets and clumps of galaxies coalesce under their own gravitation, a conceptually similar problem because every galaxy attracts every other galaxy.) Leonard and Grégoire Winckelmans (MS '85, PhD

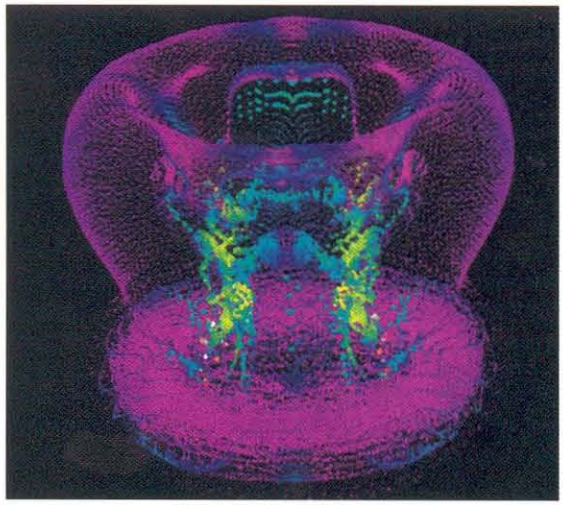

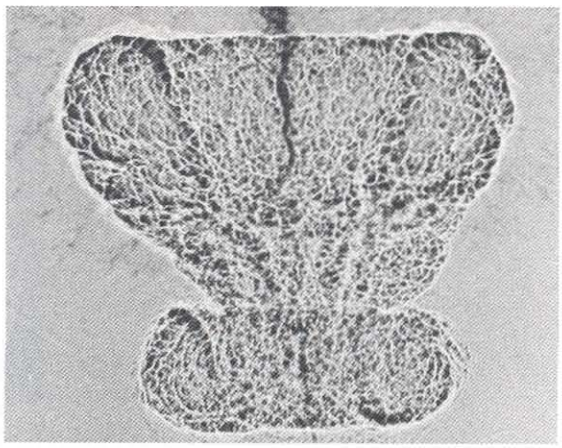

Right: In this three-dimensional vortex-particle simulation by Salmon and Winckelmans (top), the colors represent the strength of the vorticity. The simulation bears a striking resemblence to Haas and Sturtevant's shadowgraph of a Mach 1.25 shock wave striking a helium bubble (bottom). Below: Shiels and Koumoutsakos's two-dimensional vortex-particle simulation of the wake from an oscillating cylinder, which appears in two dimensions as a circle. Red and yellow show increasingly negative (clockwise) vorticity; blue and white show increasingly positive (counter-clockwise) vorticity.

'89) of the Catholic University of Louvain, Belgium, have adapted Salmon's method to three-dimensional turbulence.

Salmon and Winckelmans have done a three-dimensional computation in which they immersed a sphere in a smooth flow and calculated the infinitely thin initial vorticity layer around the sphere. Then they deleted the sphere to see what the vorticity layer, left suddenly unsupported, would do. This simulation bears a strong mathematical resemblance to an experiment done in the early '80s by then–grad student Jean-François Luc Haas (MS '76, PhD '84) and Bradford Sturtevant (MS '56, PhD '60), Liepmann Professor of Aeronautics, in which a gas bubble was hit by a shock wave transmitted through the surrounding gas—a process used in laser-induced fusion studies to compress the plasma fuel. The results were strikingly similar—the bubble collapsed on itself and imploded, spitting out a braided vortex ring on the side opposite the shock wave.

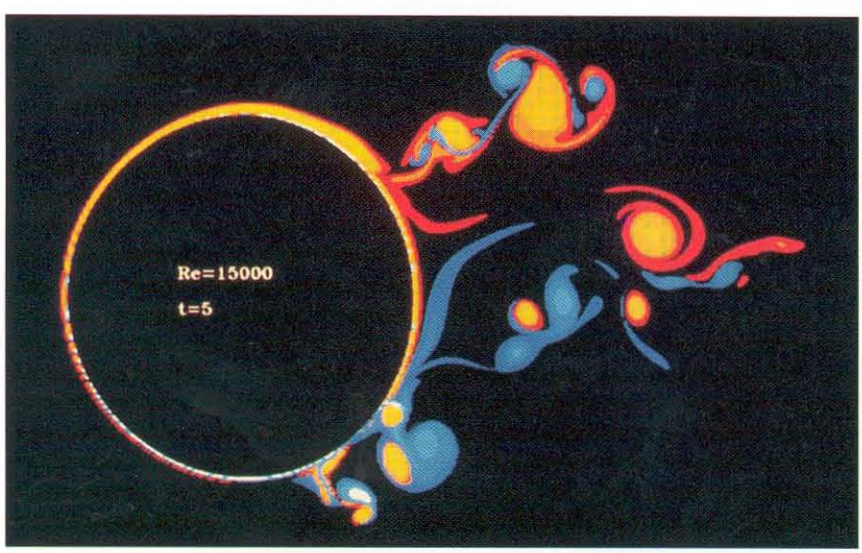

And last year, armed with a similar, two-dimensional method, grad student Douglas Shiels (BS '93, MS '94) and postdoc Petros Koumoutsakos (MS '88, PhD '93) revisited a phenomenon discovered by then–grad student Phillip Tokumaru (MS '86, PhD '91) and Northrop Professor of Aeronautics and Professor of Applied Physics Paul Dimotakis (BS '68, MS '69, PhD '73): that a cylinder strongly oscillating around its axis can have a greatly reduced drag (*E&S*, Winter 1990). The work, done on the Cray T3D at JPL, has shown that the cylinder's wake is dominated by pairs of vortices that spin in opposite directions. The vortices form spontaneously from instabilities in the boundary layer surrounding the cylinder, and radically alter the wake's flow.

Such discoveries fall naturally out of realistic models. The earlier history of computational science (as opposed to computer science, which is the construction and programming of the machines on which computational science is done) had been a lot like the anonymous broadside posted in the hallway outside more than one Caltech lab. This two-column "Guide to Effective Scientific Communication," which purports to translate phrases commonly found in the literature, lists the English equivalent of "Correct to within an order of magnitude" as "Wrong!" In order to make a model simple enough to actually run, the modelers would have to guess which details could safely be eliminated. "As you simplify the model," says CACR director Paul Messina, "you start throwing out phenomena. Then you wind up not matching the experimental results, because things were left out. Kuppermann's work is an example—who would have thought that a geometrical phase effect would be important?" But now that computers are beginning to reach a level of power where no detail is too small to include, models can be made that replicate the real-world data exactly. And once a model does that reliably, you can begin to take any unexpected results it generates as manifestations of fundamentally new phenomena that are being revealed by, and that are not artifacts of, the model. In the words of Steve Koonin, now wearing his provost's hat, "Supercomputing has been a great enabler for science all across campus. You just can't do science without it." □