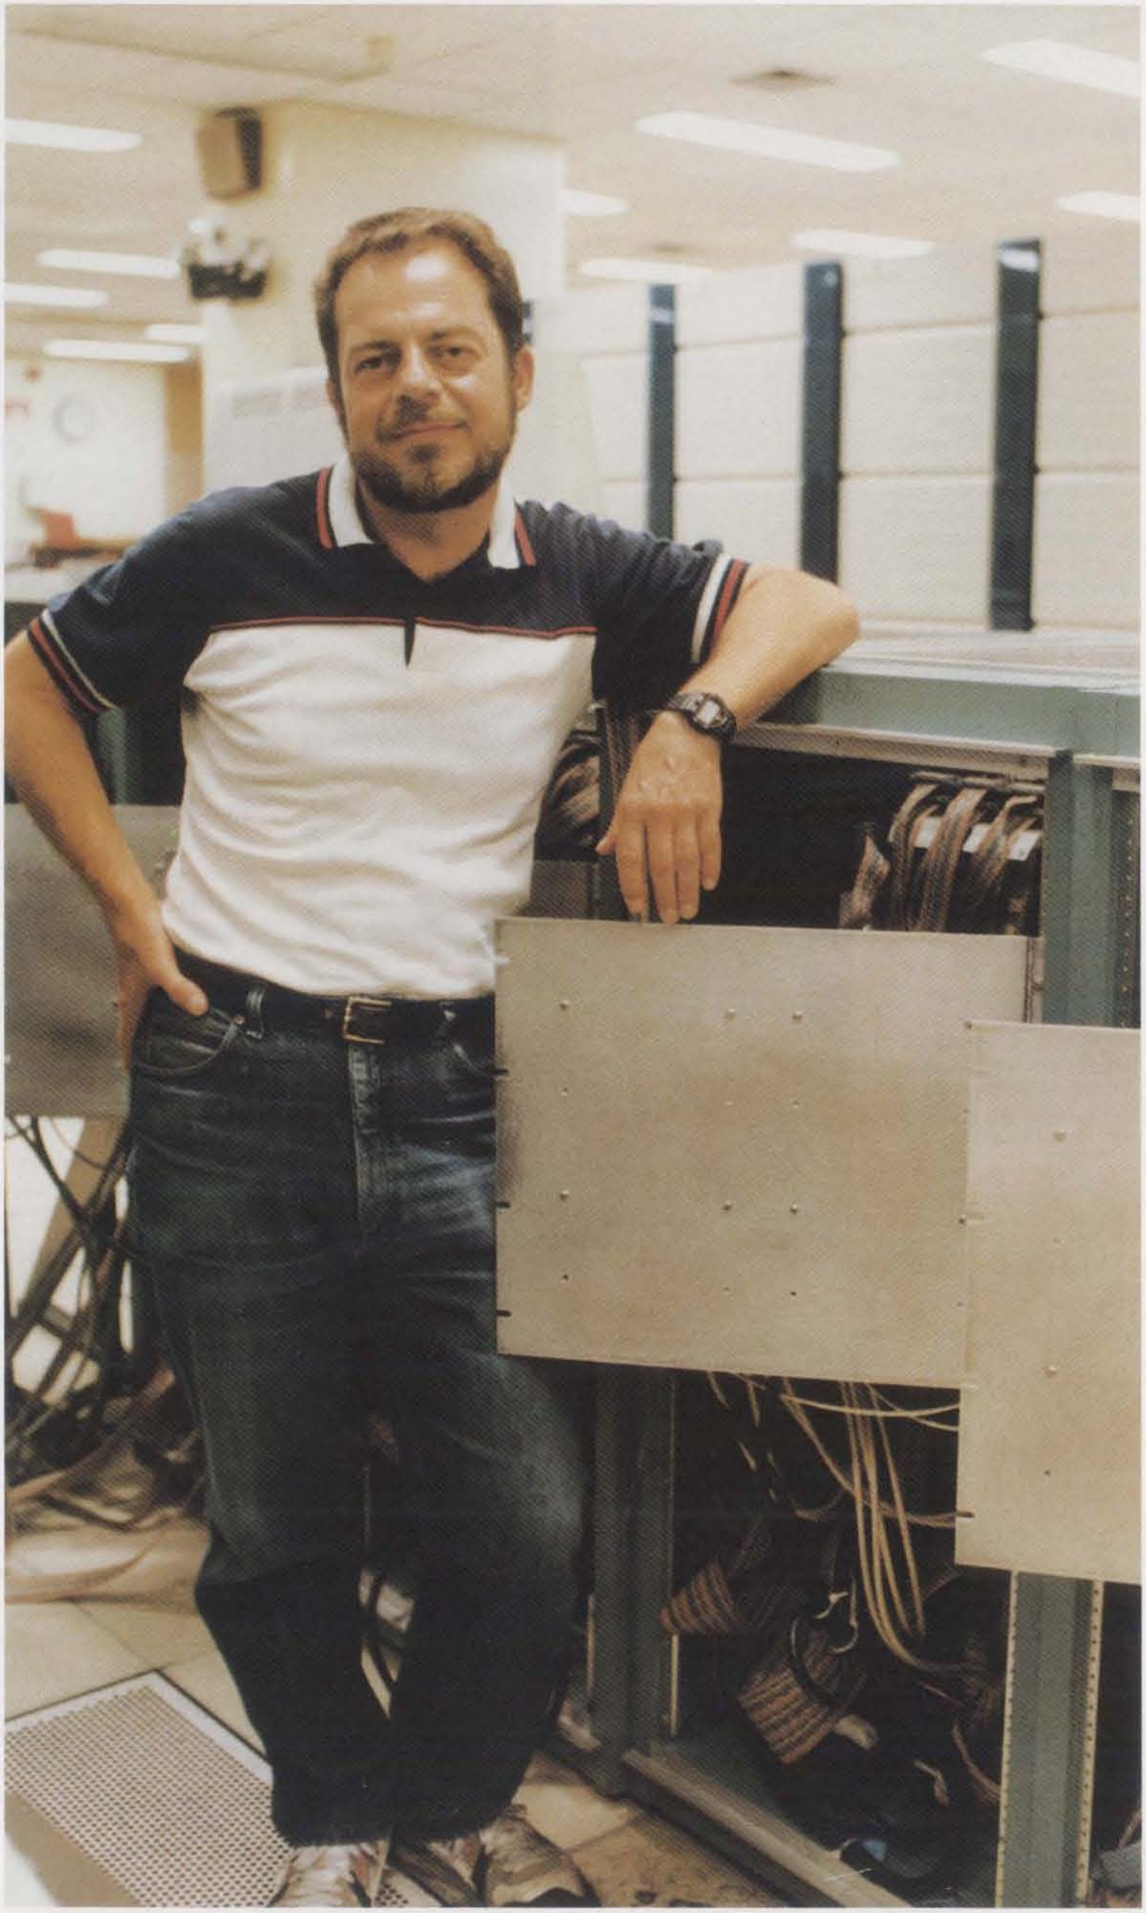

Sharing Supercomputers

The six machines that have been assembled over the past couple of years offer the campus and JPL a truly vast amount of computation power.

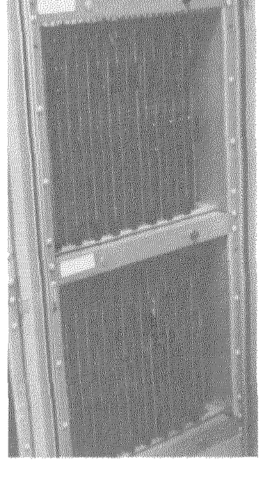

Paul Messina poses with two of the CCSF's supercomputers—in the background, the Ametek/Symult S2010, and in front, the JPL-built Mark IIIfp Hypercube, one cube (16 processors) of which is shown below.

As computers go, the hypercubes and other parallel machines standing modestly about in Booth Computing Center look less impressive than the previous tenants—the big old mainframes. But the six machines that have been assembled over the past couple of years as part of the Caltech Concurrent Supercomputing Facility (CCSF) offer the campus and JPL a truly vast amount of computation power. Some of the CCSF is neither here nor there. "Networks pretty much blur the line as to where the machines are," says Paul Messina, director of the center. "With networks we can use machines off-campus almost as easily as if they were here." The center's most powerful computer is a Connection Machine with 16,384 nodes—at Argonne National Laboratory in Illinois.

"We're connected to everything," says Messina, who spends much of his time on computer networks. Caltech has a high-speed line to the San Diego Supercomputer Center's very large Cray, funded by the National Science Foundation. Caltech has recently leased a Cray for JPL's Supercomputing Project; JPL takes up most of the Cray's time, but campus also shares it. And Caltech is a partner with Rice University in the Center for Research on Parallel Computation, an NSF Science and Technology Center that also includes Los Alamos, Argonne, and Oak Ridge National Laboratories. "We call this a center, but it's spread all over the country," notes Messina.

Concurrent or parallel processing, developed in the last decade, uses bunches of computer processors, called nodes, harnessed together to work simultaneously on parts of a problem, instead of in the sequential, one-calculation-at-a-time manner of conventional computers, such as the Cray. Charles Seitz, professor of computer science, and Geoffrey Fox, former professor of theoretical physics (now at Syracuse University), pioneered Caltech's entry into parallel processing in the early eighties. Their unique collaboration gave Caltech a head start in the field with "real hardware and real software working on a real problem" (*E&S*, March 1984). Fox envisioned a parallel system to solve his massive computation problems, and Seitz and his graduate students came up with a design architecture based on the hypercube, a method of linking processors based on a Boolean *n*-cube. Although a hypercube is a well-known mathematical concept—a geometrical construct in *n*-dimensional space whose vertices are connected in a manner analogous to a cube's vertices in three-dimensional space—no one else had thought of connecting a computer that way. Seitz's 4-node prototype proved so adept at running Fox's programs that the two continued to collaborate and built the 64-node "Cosmic Cube," which Fox immediately set to the task of calculating a previously unsolvable (because of its sheer size) problem in quantum field theory. The Cosmic Cube inspired a number of commercial computers, built by Intel, NCUBE, Floating Point Systems, and Ametek, among others; part of the Connection Machine also derives from it.

JPL also built some descendants (designed jointly with campus) of that original hypercube —the Mark II and Mark III. The most recent

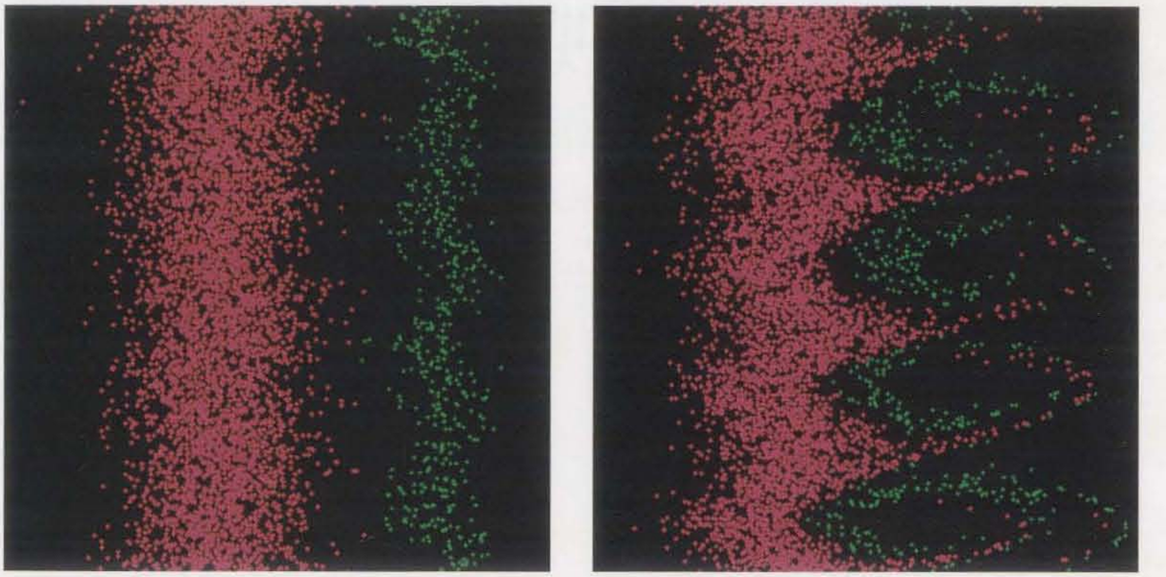

In modeling the behavior of a plasma when a beam of electrons is shot through it, the Mark IIIfp tracked the position (y-axis) and velocity (x-axis) of the electrons. The magenta dots represent the electrons of the background plasma, which tries to eliminate the electron beam by trapping its electrons (green dots) in an expanding wave (far right).

offspring is known as the Mark IIIfp (for floating points) and has 128 nodes. It used to live at JPL, but has been transported to campus over the past year in chunks of 32 nodes at a time. The Mark IIIfp has been up and running (even while split) for about two years, but it's been only recently, according to Messina, that commercial machines are beginning to catch up to it in speed. As a multiple-instruction, multiple-data-stream (MIMD) computer—that is, each processor is doing something different from the others—it's also suited to more general purposes than are some of the faster machines. The Connection Machine, for example, is single-instruction, multiple-data-stream (SIMD) and has to have all its 16,384 processors doing the same thing in lock step. Concurrent machines differ in the size of their processors, so the number of nodes is not necessarily an indication of a computer's capacity. The Connection Machine uses very small, 1-bit processors (with a floating point processor for every 32 of them), while the Mark IIIfp employs fewer but more powerful ones—each of the 128 is the size of a PC board and has its own floating point processor.

The Mark IIIfp has found plenty of users in both of its residences. Paulett Liewer, a staff member at JPL and a visiting associate in applied physics on campus, used the Mark IIIfp for developing parallel algorithms for plasma particle simulation codes. One problem involved modeling the behavior of a plasma (an ionized gas) when a beam of electrons is shot through it—a kind of interaction important in free-electron lasers, in a number of microwave devices, and in radio bursts from the sun. In the illustrations above, which plot the velocities of the electrons along the *x*-axis and their positions along the *y*-axis, the background plasma electrons are shown as magenta dots and those of the electron beam as green dots. In an ordinary gas, collisions slow down the particles, but here that doesn't happen. The plasma becomes unstable and tries to eliminate the beam. A wave evolves by tapping the electron beam's free energy, and grows until it traps some of the electrons in the vortices shown in the right-hand illustration. Modeling the evolution of this wave involves tracking the position and velocity of each of the electrons—a task ideally suited to parallel processing since it can be divided up into a block of space for each processor to follow, dispatching its duties in a minimum of time.

Liewer, working with Viktor Decyk of UCLA, has also used the Mark IIIfp for a theoretical calculation of the earth's bowshock—the magnetosonic shock wave created when the solar-wind plasma hits the earth's magnetic field, analogous to the shock wave produced in air when a plane exceeds the speed of sound. Again, the problem involves tracking the paths of thousands to millions of particles. "We divide them up by dividing space into regions and giving each region to a different processor," says Liewer. "Each processor has to keep updating the trajectories of all particles in its region. But sometimes particles migrate from one region to another. Then each processor must communicate with its neighbors to trade information about which particles are leaving or entering its space."

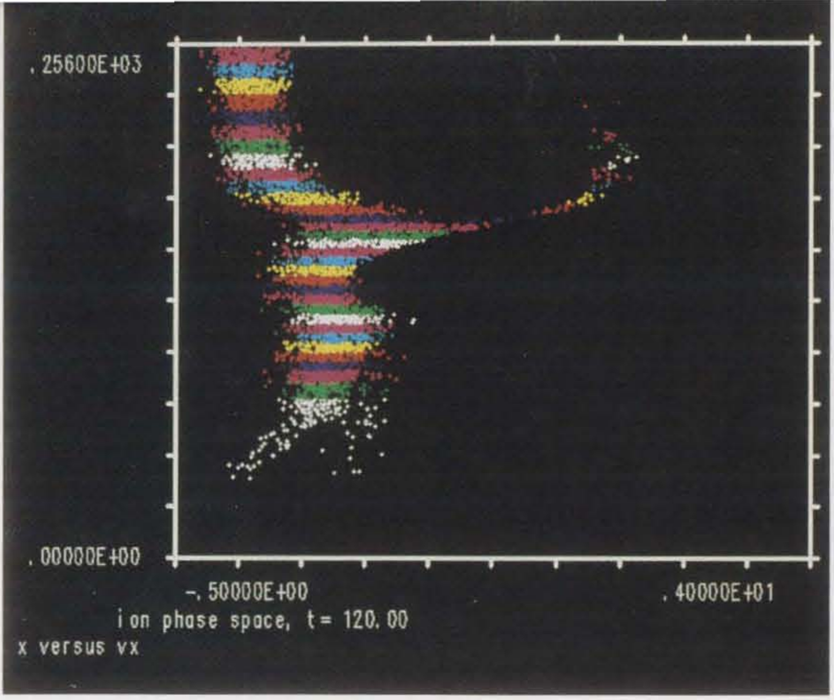

A simulation of the earth's bowshock shows the solar-wind ions streaming in from the top and slowing down as they hit the shock about a third of the way down. The vertical axis represents ion position, and the horizontal axis, velocity. Each color stands for ions in the region of space tracked by one processor of the Mark IIIfp.

In the illustration above each dot represents an ion and each color represents one processor's region of space. The vertical axis represents the position of the ions and the horizontal axis, their velocity. As the solar-wind ions stream in from the top of the graph, they hit the earth's bowshock about a third of the way down the vertical axis, and an abrupt slowdown occurs. A few of the ions don't pass through the shock but bounce off it, producing the straggling tail of particles off to the right. The rest pass through the shock and continue their journey. "You can do this on a serial computer," says Liewer, "but it all depends on the size of the computer and how much time you can get. The amount of computer time we could get on the Mark III let us run circles around others trying to model this situation."

Since there aren't many chances to sample the earth's actual bowshock, such theoretical calculations help to interpret the data that does come back from missions to space. Liewer is currently working on a proposal to calculate the "termination shock" at the edge of the solar system, where the solar wind slows down before meeting the interstellar wind. After encountering Neptune last year, the spacecraft Voyager II is on its way to the termination shock, and is expected to cross it early in the next century.

A couple of 8-node Mark IIIfp hypercubes also contributed to Voyager's spectacular Neptune encounter. Jerry Solomon and JPL's image analysis systems group used them to make animated mosaic sequences of the planet's atmospheric dynamics. In the space of just a few hours the hypercubes were able to perform the intensive computations on the Voyager data to derive the differential rotation rates of the Great Dark Spot and the "Scooter" feature closer to Neptune's south pole. On a normal computer such calculations would have taken days—too long to be of any use. The Mark III's speed made it possible to predict the positions of these features relative to each other, allowing the spacecraft's handlers to give Voyager precise instructions on where to point its cameras as it made its closest approach. At JPL the Mark IIIfp has also performed simulations of an advanced global communications network and has been used to develop algorithms for tracking missiles. Other CCSF machines have found application in synthetic aperture radar processing and in neural networks.

Some of Caltech's biologists also use the facility's huge computing power for mapping neural nets and for other applications. Interested users have also shown up from chemistry and aeronautics, and one recent postdoc has been recruited by an investment banking firm for his hypercube expertise. But, with Fox pioneering its applications, concurrent processing has had its longest-term use among the physicists and astrophysicists. One current astrophysics project is a search for pulsars. A pulsar is a spinning neutron star with an extremely strong magnetic field, whose captured electrons and hot plasma produce strong radio waves. These waves flash like the beam from a lighthouse as they sweep out through space. The pulses are also remarkably stable, providing astronomers and physicists with an extremely accurate and precise clock in the sky, useful for investigating such things as the relative motion of stars and, especially with binary pulsars, gravitational radiation.

Neutron stars are also interesting in their own right, representing perhaps the most exotic kind of material observable in the universe, according to Peter Gorham, research fellow in physics. "They're sort of the next step up from a black hole." Formed when a star's nuclear fuel is spent and it explodes and implodes at the same time, a neutron star is only about 10 km in diameter but with a density of something like 10 billion tons per teaspoonful. Its gravity can be the same as a star 1½ times the mass of the sun, and its magnetic field a trillion times that of the earth. With pulsars, astrophysicists are able to study matter in a state in which it doesn't exist on earth (except perhaps for very short times in accelerators).

So astrophysicists would like to find as many different types of these things as they can. The

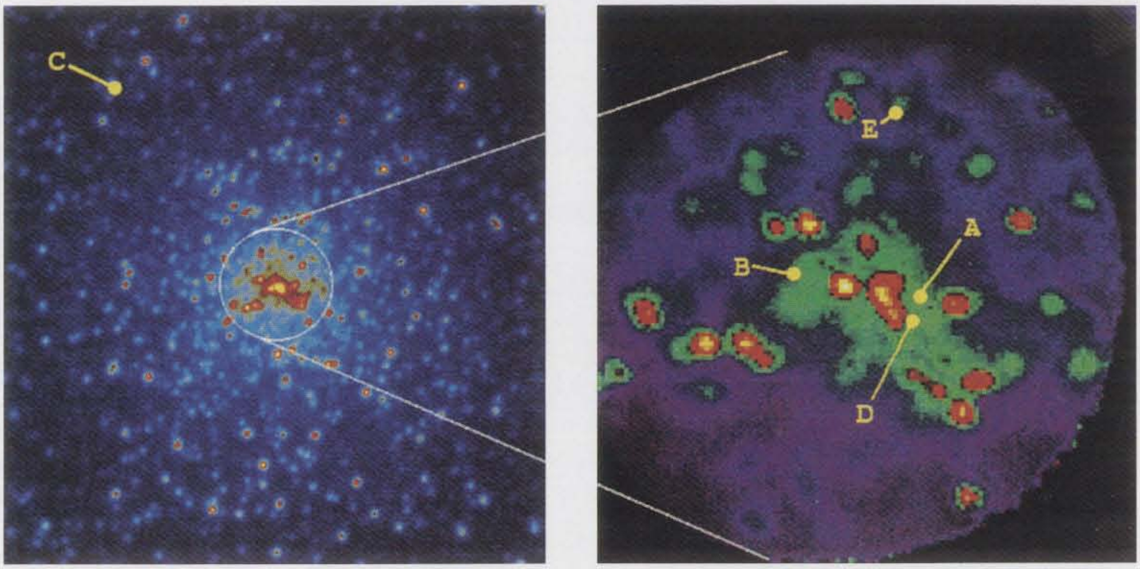

So far 19 such pulsars have been discovered worldwide, and 8 of them have been found using the Caltech supercomputing facility.

trouble is, the radio pulse of the fainter ones is drowned out by the noise in the instruments and the noise in the sky. To find them you can make very long observations and then use Fourier transforms, a mathematical device that converts time into frequency. Because a pulsar's pulse is so precise in frequency, the Fourier transform collects all of the faint pulses into a set of sharp spikes at the exact harmonies of this frequency, and so can distinguish the signature of the pulsar from the random sky noise. Unfortunately, such calculations are very computation-intensive and require supercomputer performance such as that provided by large parallel machines.

A couple of years ago a group at Caltech, including Shrinivas Kulkarni, assistant professor of astronomy, Tom Prince, associate professor of physics, graduate student Stuart Anderson, and Gorham, decided to tackle the computational problems of finding faint pulsars with the NCUBE—a 576-node hypercube then recently arrived at the Concurrent Supercomputing Facility. "It turns out that the hypercube is a very good architecture for doing fast Fourier transforms," says Gorham. "In fact, you can't make a better parallel machine for doing FFTs that's still a general purpose computer." One of the reasons is its efficient communication between processors, which leaves more time for actual processing. Another reason is its sheer size—for the long Fourier transforms you need a vast amount of memory. The longest observation the Caltech group can take on the cluster they're interested in is about two hours and amounts to about 16 million sequential samples; this isn't long enough to find the fainter objects, so they make many of these 16-million-sample observations, then do Fourier transforms on one after another and stack them up. They hope eventually to be able to stack up hundreds of such observations and find fainter and fainter pulsars.

Prince's group has been using this technique to search for pulsars in a globular cluster called M15. It has been a surprise to discover globular clusters harboring pulsars; the birth of a neutron star should be accompanied by a kick powerful enough to launch it out of these loosely bound mini-galaxies. So far 19 such pulsars have been discovered worldwide, and 8 of them have been found using the Caltech supercomputing facility. "Caltech is well established as one of the leaders in this particular kind of pulsar," says Gorham.

Five of these eight are in M15, and one of those, at the edge of the cluster, is a binary star. All of them have relatively short pulse periods —between 4.6 and 110 milliseconds. Since these pulsars cannot be young and their periods should have slowed down much more than those observed, the astrophysicists theorize that at some time they had met companion stars able to lend them angular momentum and spin them back up to a very high frequency. But such binary relationships apparently don't last long in the crowded conditions in the cluster's core; close passes result in many companion stars being stolen away.

The pulsar search has been going on for about a year and a half. The NCUBE had an initial problem of not being exactly user friendly. And it also had some data-storage and input-

Left: The Caltech group has found five faint pulsars in globular cluster M15 using a concurrent supercomputer to perform the intensive calculations of Fourier transforms in order to distinguish their frequency signatures. C is a binary pulsar somewhat farther out than the others. Below: The ragged, fuzzy image shows an optical binary star (BS 5747) as it appears from a ground-based telescope, magnified a thousand times. Reconstructing it, using interferometry techniques analyzed by an NCUBE computer to remove the effects of atmospheric turbulence, yields the two distinct points below.

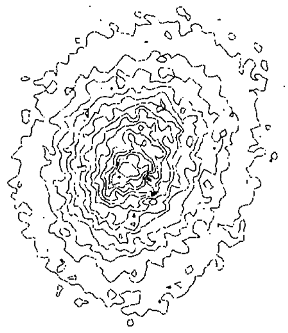

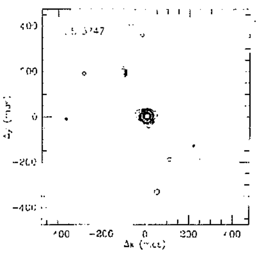

output problems. "We had these gigantic data sets," says Gorham, "and we had a supercomputer, but we had to spoon-feed the machine. It would take us a week to get all the data loaded and then it would finish the search in a few hours." Such problems have largely been solved, but the loading time is still greater than the computing time by a factor of 10.

NCUBE computation power has also enabled Caltech physicists and astronomers to dramatically increase the resolution of ground-based telescopes—a technique that now may loom even larger in importance with the failure of the Hubble Space Telescope's mirrors to achieve high resolution. The Hubble was to gain its superior resolution from orbiting above the earth's turbulent atmosphere, whose slightly positive refractive index distorts incoming light waves. This phenomenon is called "speckle," and results in a point looking more like a fuzzy blob. This wavefront distortion is such that the resolution of the 5-meter Hale Telescope at Palomar Observatory, when making highly magnified images of distant objects, is actually a factor of 30 worse than the telescope is theoretically capable of, that is, with no atmospheric interference. Even if the Hubble Telescope were functioning properly, its 2.4-meter mirror could not achieve the theoretical resolution of the Hale or the 10-meter W. M. Keck Telescope in Hawaii, which will see first light this fall. And bigger telescopes collect more light and so can see fainter objects. So people are still very interested in ways to achieve the limiting resolution from the ground rather than trying to fly a 10-meter telescope.

Prince, Gorham, and several graduate students are also involved in this project, as is a large contingent of astronomers (including Kulkarni and Gerry Neugebauer, the Howard Hughes Professor, director of the Palomar Observatory, and chairman of the Division of Physics, Mathematics and Astronomy). The group adapted a technique from radio astronomy—very long baseline interferometry (VLBI), developed to counter the wavefront corrugation that the ionosphere imposes on radio waves, similar to what the atmosphere does to light. By combining the signals of three or more telescopes arrayed in a triangle, you get a measure of the wavefront that is the sum of that seen by each of the three telescopes. With enough triangles you can solve for the correct wavefront and derive more information about the source. VLBI usually makes use of 12 telescopes around the world, thereby increasing its "aperture" to the size of the earth.

Optical telescopes can borrow this technology because the key measure is the size of the aperture relative to that of the wavelength; VLBI creates an aperture of 15,000 km, and a radio wavelength is measured in meters, roughly a ratio of 10^7 to 1. In optical astronomy with the Hale Telescope's 5-meter aperture and a wavelength of about half a micron (10^{-6} meters), you also have a ratio of about 10^7 to 1.

What the group did was to divide the 5-meter telescope into an array of 10-cm cells or "telescopes," all combining their light onto a single detector. This makes about 2,000 "telescopes" and a possible 10 million triangles. You can do the same trick as with VLBI—sum up the wavefront effect around all the triangles and the effects will cancel out. Unfortunately, every 10 milliseconds or so the waves change pattern, so you have to take snapshots of the source every 10 milliseconds—up to perhaps 100,000 snapshots. "What we would do then," says Gorham, "is distribute all the snapshots around to the processors of the NCUBE; each processor was responsible for a certain set of triangles that he would check out; he would get a frame [one of the 100,000] and do all of his triangles and pass it to the next guy, who would do all his triangles on it and pass it on Each of the frames got passed around from processor to processor until all 100,000 were done. And then you have this set of measurements of all these triangles, and you can solve all the equations to get back what the actual wavefront was doing. So it's a good NCUBE problem in that respect."

The biggest problem is that the technique can handle only fairly bright objects—if you can't get enough light in 10 milliseconds, you're out of luck. Probably its best use will be in infrared astronomy, where, according to Gorham, "compared to the optical there are still lots of bright things that are interesting astronomically."

Besides these research projects, there are many more, both on campus and at JPL. And the applications will likely continue to grow as they have since 1984 when parallel computing at Caltech consisted of the one-of-a-kind Cosmic Cube. Messina intends to keep pace to make the Concurrent Supercomputing Facility capable of world-class scientific computation. Future plans include acquiring a *really* large-scale commercial parallel computer—an order of magnitude faster than the Cray. Messina is also establishing a very-high-speed network (Caltech, JPL, Los Alamos, San Diego) of big concurrent machines to be used simultaneously over large distances. This is a research project in itself ("can you do it and profit from it?"), but this network too will be solving "real" problems. □ —*JD*