California Institute of Technology

Engineering & Science

Summer 1990

In this issue

JPL and Caltech partnerships in engineering and science

Caltech–JPL Collaboration

Ever since it was launched by GALCIT's experiments with rocketry in the late 1930s, the Jet Propulsion Laboratory has been closely tied to Caltech. Officially managed by Caltech—first for the Army, and since 1959 for NASA—JPL is the only one of NASA's nine centers to be operated by a university. Both in the fifties and again in the eighties, the Lab's relationship to defense has made some Caltech faculty nervous, and questions about the common ground between the two institutions have occasionally surfaced on both sides. But the connection has endured and, especially of late, prospered.

Separated by only a few miles of freeway, the populations share a significant amount of overlap. JPL has always recruited intensively from its neighbor and employs the largest number of Caltech alumni of any company or university. Several JPL staffers hold joint appointments on campus; grad students are increasingly discovering research opportunities "up the hill"; and SURF (Summer Undergraduate Research Fellowship) students are pursuing more and more projects there. One JPL staff member, who is also a visiting associate at Caltech (and whose husband works at JPL), notes that the Caltech–JPL connection is also cemented in what is probably an unusually large number of marriages.

The incompatibilities that do arise between the two institutions are more of a "cultural" nature. Caltech is a private educational and research institution, and JPL is an engineering organization funded by the government; Carver Mead once called it "the big gray machine that builds things." Building things that scientists need is what JPL does well, and most often those are big things that demand more organization and more layers of management than Caltech scientists are used to. But when interests converge, the different strengths of the two institutions are complementary.

"It is not surprising that when matching the interests of JPL and campus, there's not a complete match," says Lew Allen, JPL's director since 1982, who will retire in December. "There are areas where the interests overlap and areas where they do not." But he believes that the former are much more numerous and intensive than they have been in the past. Allen points to the increased emphasis at JPL on particular areas of basic technology that have relevance to the Lab's primary mission of space exploration, but are also complementary to basic research in science. In particular he mentions that "microelectronics, neural networks, and parallel computing have resulted in developing some strengths at the Laboratory and a number of very productive interactions with campus research." Some of these interactions, many of which began with seed money from the JPL director's discretionary fund, are described in the articles in this issue of *E&S*.

California Institute
of Technology

Engineering & Science

Summer 1990
Volume LIII, Number 4

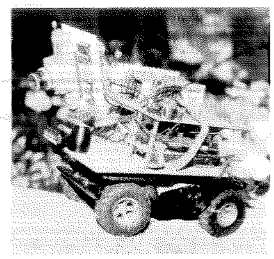

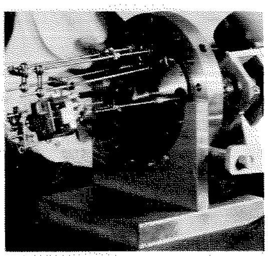

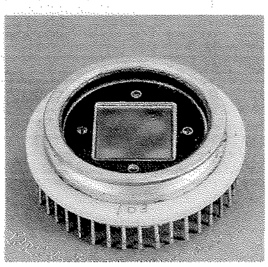

On the cover: A self-steering vehicle built by Timothy Horiuchi (BS '89) placed in Throop Site's sylvan setting. The toy car's neural-net vision system distinguishes contrasts sufficiently to follow a line of black tape on a white floor. Researchers hope to develop a system that can actually see well enough to roam the great outdoors. A story on neural-network research, including artificial vision, begins on page 4.

Engineering & Science (ISSN 0013-7812) is published quarterly, Fall, Winter, Spring, and Summer, at the California Institute of Technology, 1201 East California Boulevard, Pasadena, California 91125. Annual subscription $8.00 domestic; $20.00 foreign air mail; single copies $2.00. Third class postage paid at Pasadena, California. Published by the California Institute of Technology and the Alumni Association. Telephone: 818-356-3630. Postmaster: Send change of address to Caltech 1-71, Pasadena, CA 91125.

PICTURE CREDITS: Cover, 4, 5, 15-17, 27-29, 41, 43, 44 – Robert Paz; 7, 9, 21, 22, 24, 25, 34, 36-40, 42, 43 – JPL; 7, 8, 12 – Tim Brown; 8 – Harry Langenbacher; 10 – Landsat; 10 – Niles Ritter; 13 – David Van Essen, Dan Felleman; 14 – D. Van Essen, Charles Anderson; 14 – Matt Wilson; 15 – Pierre Baldi; 18 – David Brady; 19 – Ken Hsu, Xiang-Guang Gu; 26 – Herb Shoebridge; 26 – John Baldeschwieler, Robert Driscoll, Mike Youngquist; 27 – R. Driscoll; 27 – M. Youngquist; 30, 31 – Paulett Liewer; 32 – Peter Gorham.

E. Micheal Boughton
President of the Alumni Association
Theodore P. Hurwitz
Vice President for Institute Relations
Robert L. O'Rourke
Assistant Vice President for Public Relations

STAFF: *Editor* — Jane Dietrich
Writer — Douglas Smith
Copy Editors — Michael Farquhar, Julie Hakewill, Betsy Hatch
Production Artist — Barbara Wirick
Business Manager — Debbie Bradbury
Circulation Manager — Susan Lee
Photographer — Robert Paz

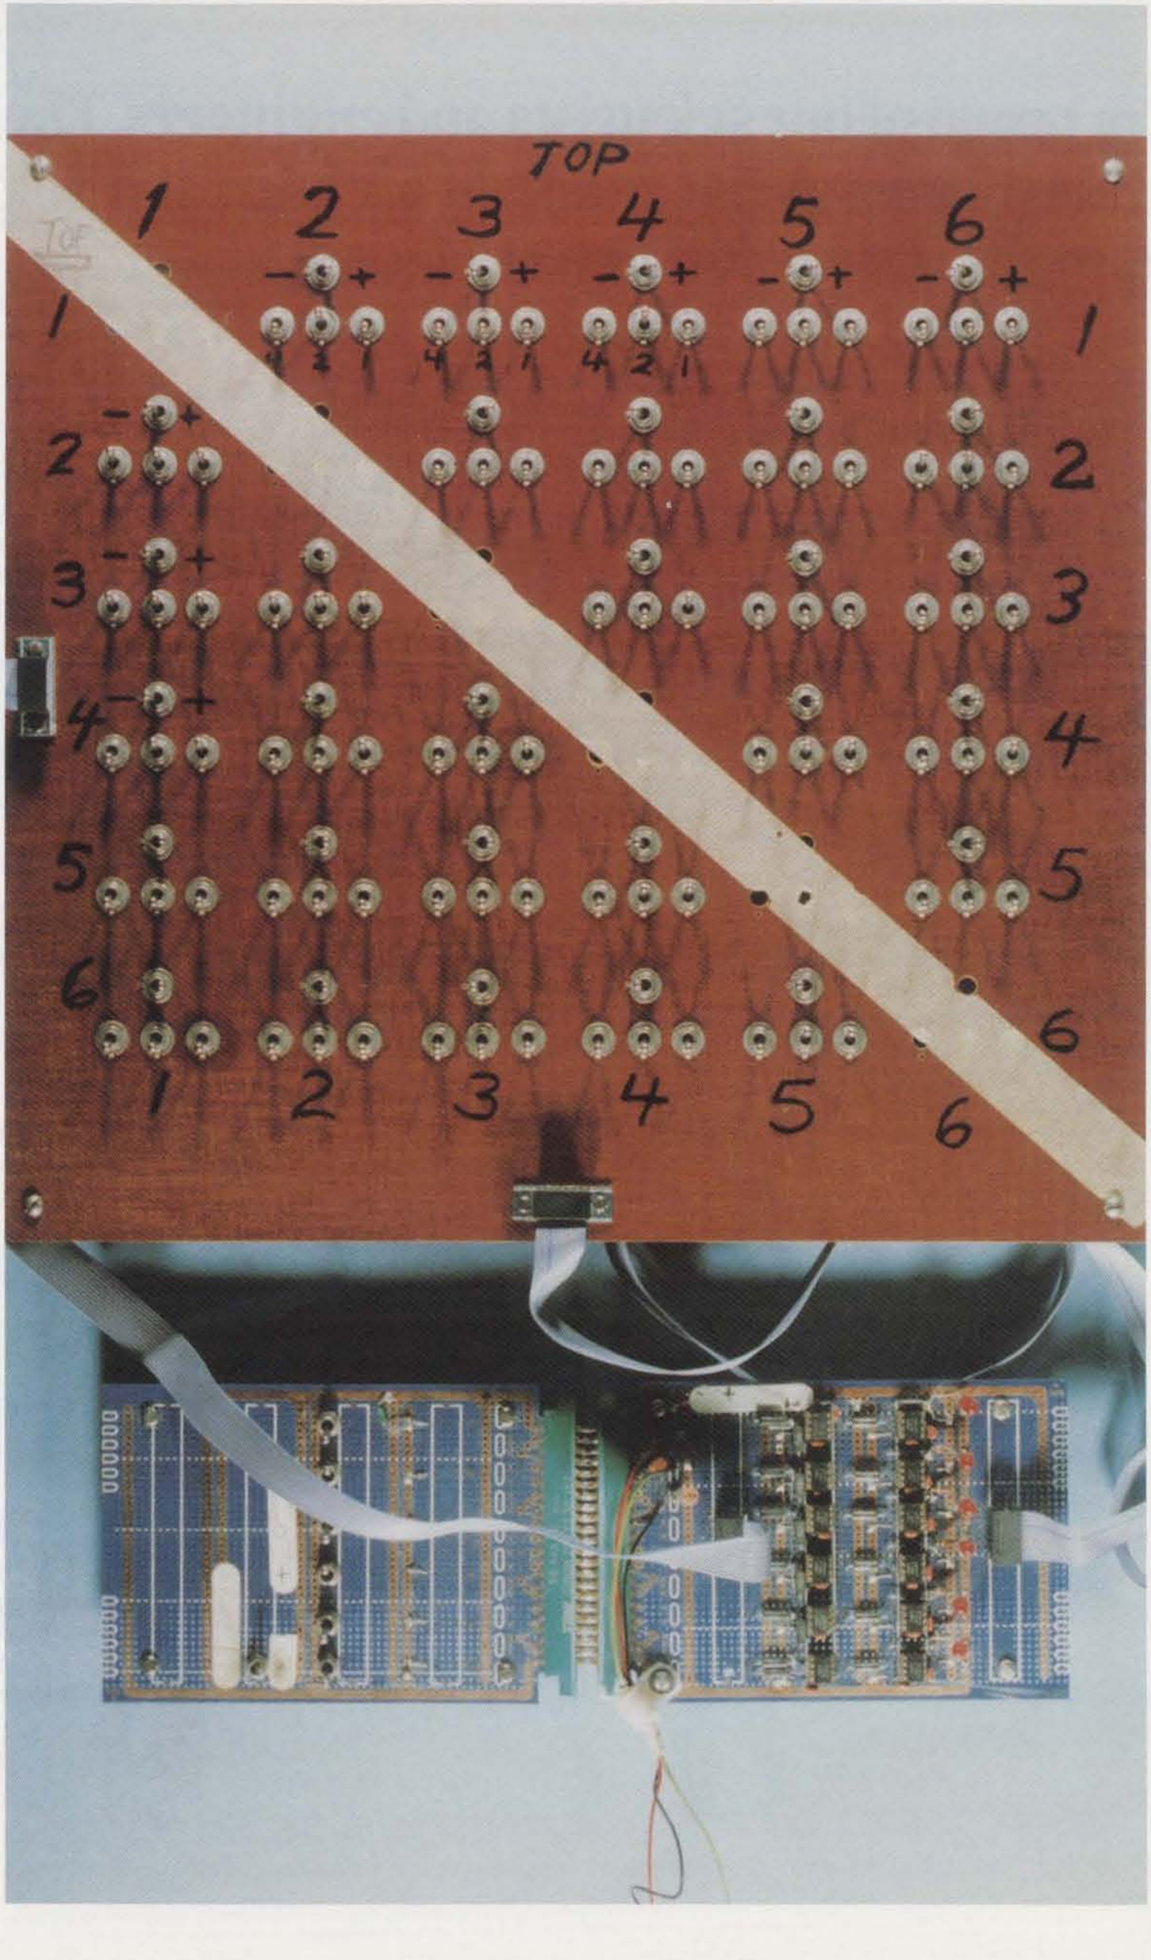
TOP

Neural Networks: The Caltech – JPL Connection

The world's first man-made feedback neural net had six fully interconnected neurons. Each connection used a group of four toggle switches—the upper one chose whether the connection was inhibitory (–) or excitatory (+), and the lower three set the connection's strength. The six toggle switches on the left side of the blue circuit board input a pattern, and the column of light-emitting diodes (LEDs) to the right lit up with the output.

Imagine an autocratic chef, who can't stand to have another presence in the kitchen, preparing a seven-course banquet. Taped to the wall is a master instruction list that combines all the individual steps from all the recipes on the evening's menu interleaved in a sequence that, if followed to the letter, will produce the meal, with each dish appearing on the table at its appointed moment. Our chef, although gifted with an excellent set of taste buds, is extremely absentminded and can remember just one instruction at a time. Thus our hero plops a dozen potatoes on the counter, runs back to consult the list, peels the potatoes, dashes back to the list, cuts the potatoes into one-inch cubes, checks the list, and so on. A conventional—"serial"—computer, from the lowliest laptop to the mightiest mainframe, works in exactly the same way. The computer's central processing unit executes the recipe, or program, step by step—pulling data out of storage piece by piece, doing something to each one, and then putting it back before looking at the next instruction. The more powerful the computer, the faster the chef sprints. But that's not how the brain works at all.

The brain—any brain, from a slug's on up—is more like a medieval kitchen in a great lord's palace on the eve of a feast. A multitude of helpers bustles at a variety of tasks, shouting advice and instructions to one another. Everything happens all at once, and the banquet emerges almost spontaneously from the coordinated actions of many individuals. In the brain, these individuals are called neurons, and computer systems based on the notion of a network of simple devices acting collectively are called neural networks. The net's power lies in the interconnections, or synapses, between the neurons. The neuron itself is a threshold device that "fires"—generates an output—whenever its cumulative input exceeds its threshold. But one neuron in the human brain (and there are about 10^{11}—100 trillion—of them) may connect to 10,000 or more other neurons.

This redundant, highly interconnected scheme has other advantages. Returning to the kitchen for a moment, if some prankster crossed a line off of Super Chef's master list, the goose might get cooked unplucked, an error unlikely to happen in the castle kitchen. Or if Super Chef should fall down the wine-cellar stairs, the guests would go hungry that night. But if a few of the castle staff don't show up, no matter—the dinner still comes off. "Fault-tolerant algorithms" (an algorithm is a detailed strategy for attacking a problem) and "graceful performance degradation" as bits of hardware fail are hallmarks of neural nets but not, alas, of serial computers.

Furthermore, serial computers' ability to follow lots of step-by-step instructions very, very rapidly makes them dandy adding machines or tax auditors, but doesn't enable them to recognize Aunt Emma from a photograph, or, having recognized her, to remember that she and Uncle Joe have two children and a cabin in the mountains, and every other detail of their lives. It doesn't help a computer to reach for a pen to write her a note, either. These problems don't break down into cut-and-dried programs because there are just too many variations to list every contingency explicitly. But the connections in a neural network act like unexplicit rules. Patterns of associations—Aunt Emma, Uncle Joe, two kids, mountain cabin—become patterns of connections between neurons. The stronger a connection, the closer the association.

Given an input that matches some part of the pattern, the connections allow the net to retrieve the rest of it—a feat called associative memory. The connections feed back into one another, and signals slosh back and forth through the network along the pathways with the strongest connections. The feedback pulls out all the related information, regardless of which item you begin with—think of the mountains, and you'd still come up with Emma and Joe. This is called content-addressable memory. And the multiple connections can encode multiple memories, keeping Uncle Joe from being confused with your bachelor brother Joe. The connections can even cope with partially wrong input: if someone asked you about Aunt Emma's three children, you could set the record straight. Conventional computer memories, on the other hand, merely tuck away each tidbit of information in pigeonholes that bear no clues to the relationships between their contents.

People have known all along that computers and brains don't think alike, but biology's subtleties remain elusive.

Most things that neural nets do well—recognizing patterns such as faces, making decisions based on fuzzy or incomplete data, or displaying motor skills such as hand-eye coordination—require feedforward circuits as well as feedback. The theory is that information from an input layer of neurons trickles down through one or more layers of "hidden" neurons, again following the pathways of strongest connections, to an output layer. The hidden layers somehow filter the input, recognize the critical features needed to make a decision, and steer the system to the correct output, or stable state.

Neural nets, both man-made and biological, work fast. By exploring their options all at once, rather than scrutinizing each one in turn, they give you a pretty good answer immediately instead of the best possible answer in a week. And if the nets encounter a strange new input, they'll make an educated guess based on the information they have.

But the most remarkable feature of neural nets is that they can *learn* to do these things. A man-made neural network can alter its internal connections—strengthening some, weakening others—while being shown a training set of correct input-output pairs, until its outputs consistently match the right outputs. It may take a few thousand tries to get things right, depending on the problem's complexity, but still. . . .

If neural nets are so all-fired smart, why haven't they taken over? Why fool with serial computers at all? Because it's an awful lot easier for people to design and build something that does one thing at a time. People have known all along that computers and brains don't think alike, but biology's subtleties remain elusive. The first simple neural networks capable of learning were developed in the 1960s. Throughout the 1960s and 1970s, a handful of people labored, with some successes, to develop computational models of how real neurons behave, and to create the theoretical and mathematical underpinnings needed to design machines that could, in the broadest sense, mimic that behavior.

Neural network research began to catch fire in the 1980s for several reasons. Neurobiologists have made great strides in finding out how neurons work; computers have attained enough power to make running serial simulations of parallel processes practical, if time-consuming; and the development of analog VLSI (very-large-scale-integration) chips by Carver Mead (BS '56, MS '57, PhD '60), Caltech's Moore Professor of Computer Science, is enabling powerful synthetic neural nets (albeit puny ones compared to biology) to be graven in silicon. But the conceptual kindling was probably a 1982 paper by John Hopfield, the Dickinson Professor of Chemistry and Biology, showing that the same mathematical tools that physicists routinely use to analyze large physical systems with complex interactions, such as freezing liquids, could be applied to a network of binary (on/off) switches—a simple neural net. Hopfield demonstrated that a set of switches, each of which was wired to every other switch—fully interconnected feedback—acted as an associative memory; they're now called Hopfield memories. This revelation made neural nets accessible to the legion of scientists and engineers whose last contact with biology was probably a frog in a dissecting pan in high school, but who thought that neural nets might be applicable to their own computational problems, if they only knew how to handle them.

Even now, however, neural-net research lingers in the "not yet" stage, as in, "Can you make a neural net that can read a book and summarize its plot?" "Not yet." "Well, all right; literacy *is* a lot to ask. Can you make a mosquito brain that can dodge a midair swat?" "Not yet." In fact, the most successful multilayer neural-net application being sold commercially to date is a serial-computer program to evaluate loan applicants' creditworthiness. The program picked probable defaulters on small

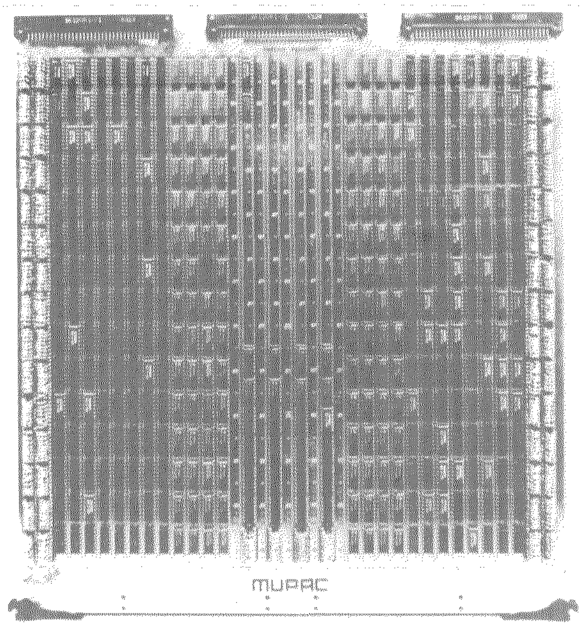

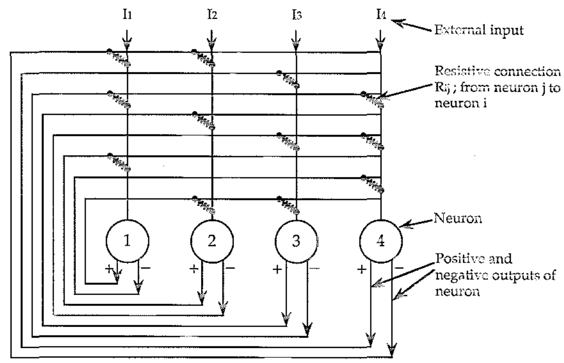

**Right: The DARPA-funded 32-neuron network.
Far right: A schematic drawing of fully interconnected feedback.**

loans as accurately or better than did typical human loan officers. The number of mosquitos employed by S&Ls is unknown.

Caltech and JPL's history of neural-net collaborations goes back to 1981. John Lambe (now retired, but still returning bimonthly to JPL as a Distinguished Visiting Scientist) was visiting campus that spring as a Fairchild Distinguished Scholar on leave from the Ford Motor Company. He attended Hopfield's first Caltech seminar on neural nets, given in the improbable guise of an applied physics talk, and was smitten by their possibilities. Lambe left Ford for JPL that fall, and immediately built the world's first nonbiological feedback neural net—six neurons, interconnected with toggle switches, arrayed on an 18-inch-square Formica board. "The joke used to be that it stored two bits of information per square foot," recalls Anil Thakoor, now the head of JPL's Neuroprocessing and Analog Computing Devices Group. This network, far too small to do anything resembling computation, was nonetheless very useful. Lambe and Hopfield played with the toggle switches and found that each neuron assumed a predictable voltage—the network had settled into a stable state, in other words—instantaneously, allaying fears that any fully parallel man-made network would never settle down, but, a victim of the less-than-precise nature of its analog components and the stray capacitances in its hardware, would oscillate forever instead. This first model was built with JPL Director Lew Allen's blessing in its most tangible form—money from the director's discretionary fund.

This network was big enough to address practical engineering concerns like power dissipation.

Lambe, Thakoor, and Alex Moopenn built a better net the following year, using Defense Advanced Research Projects Agency (DARPA) funds this time. This network was also 18 inches square and used off-the-shelf parts, but it had 32 fully interconnected neurons—1024 synapses. And instead of using toggle switches to make binary connections, this net used resistors to mimic the brain's variable-strength synapses. This network was big enough to address practical engineering concerns like power dissipation. Since the network's product is a distribution of voltages across its output neurons, any internal power loss affects it. Even when the connection between two neurons is supposed to be strong, you don't want a flood of current going through it, burning up power and converting it to heat—a good way to melt components, especially with lots of them active simultaneously. The group found that less was more—they could make good, strong connections out of a network of megohm (million-ohm) resistors, put one volt into the network, and get milliwatt power dissipation, well within the chip's comfort zone. And the output pattern was robust—if the initial voltage or some of the resistors were off by a few percent, the network came to the same stable output state as fast as ever.

Having demonstrated that building neural nets was actually practical, the next step was to try it in VLSI. As it happened, Mead and Hopfield were teaching a joint course, The Physics of Computation, that year—1983. (Mead, Hopfield, and the late Richard Feynman had initiated the course in 1981.) Two of Mead's grad stu-

A neuron is conceptually a very simple device, as the schematic to the far right shows—outputting "+1" if its input exceeds its positive threshold, "−1" if its input exceeds its negative threshold, and "0" if neither of the above is true. The reality is a trifle more complex, as shown by this integrated-circuit diagram of a seven-bit synapse, right.

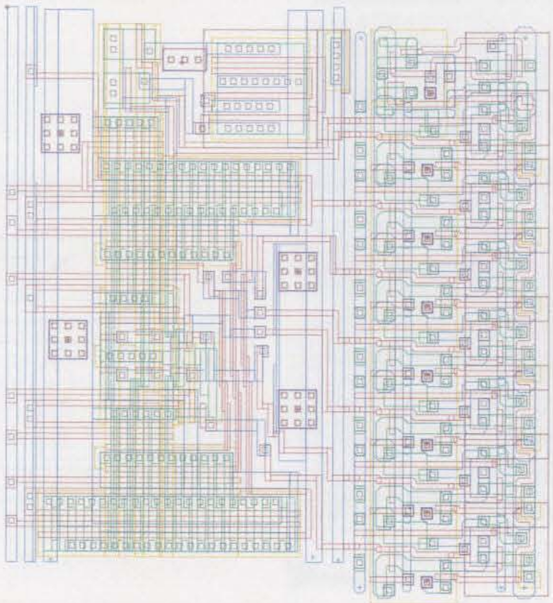

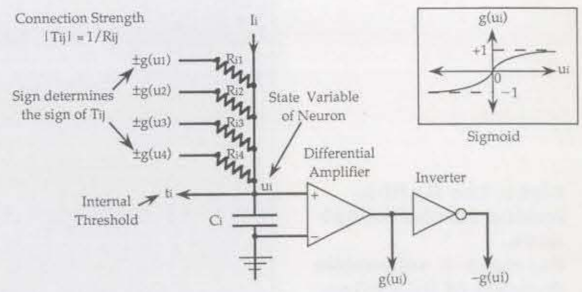

dents, Michael Emerling and Massimo Sivilotti (MS '86), designed a 22-neuron Hopfield memory chip, fully interconnected with adjustable-strength synapses, as a class exercise. A batch of the chips was actually built that December—the world's first single-chip VLSI neural nets. The chip could store three memories and retrieve any one of them from partial input in less than 50 microseconds (millionths of a second). One of the batch was severely defective—40 percent of its connections proved inoperative—yet it still held two memories, and retrieved them as quickly as its more able brethren. A larger chip followed the next year, and silicon neural nets were off and running.

(Mead's VLSI has only partially solved the connectivity problem, however. If each neuron connects to every other neuron, then the number of synapses increases as the square of the number of neurons. A 64-neuron chip—large enough to actually do something useful—needs 4096 interconnections. (By comparison, a slug brain has between 100,000 and 1 million neurons.) The brain has the luxury of making connections in three dimensions, but a microchip is still essentially flat. Fortunately, less-than-full interconnection suffices for many simple applications. It also turns out to be wretchedly difficult to make a continuously variable resistor out of silicon. Digital synapses that use stairways of preset resistances are much easier, although their design is quite elaborate. An 8-bit synapse affords 256 steps, sufficient for most applications.)

A critical mass of people had formed by 1983. The group included Hopfield; Lambe; Mead; Allen; Thakoor; Terry Cole (PhD '58), then a senior member of the technical staff at JPL and a senior research associate on campus and now JPL's chief technologist and a senior faculty associate on campus ("He knew an awful lot of people in both camps and was a very important part of the glue holding the critical mass together," says Hopfield); Demetri Psaltis, who joined the Caltech electrical engineering faculty in 1980; John Pierce (BS '33, MS '34, PhD '36), chief technologist at JPL since 1977 and professor emeritus of electrical engineering since 1980; Edward Posner, a visiting professor of electrical engineering and JPL's chief telecommunications and data acquisition technologist; and Robert McEliece (BS '64, PhD '67), professor of electrical engineering. The group began meeting regularly, if informally, to talk about neural nets, and hosted more formal sessions each October from 1983 through 1985, which drew biologists, engineers, and physicists from Lab and campus. Anyone doing neural-net work was encouraged to speak about it, and bound copies of the proceedings circulated widely up the mountain at JPL and down on campus.

By 1986, JPL's neural hardware group, then under Satish Khanna, was well established, and had launched a program to build modular chips that could be wired together to make feedback, feedforward, or hybrid arrays of any size for specific applications. The group now has a family of neuron-only and synapse-only chips in the square, 84-pin format standardized by integrated-circuit manufacturers. (These chips, and every previous one all the way back to Sivilotti's

Some of JPL's family of neural chips. From the top: analog, binary, and seven-bit versions of a 32 × 32 synapse chip, and a 36-neuron chip.

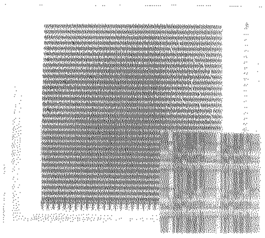

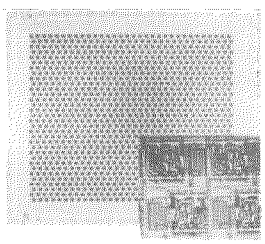

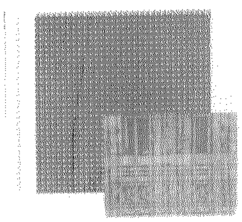

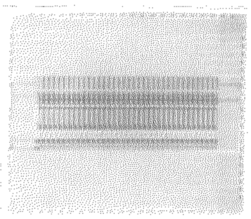

A slug brain has between 100,000 and 1 million neurons.

and Emerling's, are actually built by MOSIS, a government-funded custom-chip broker for defense-related research.) The hardware collaboration continues very closely today, with Mead a frequent visitor up the mountain, often in his capacity as a member of the Center for Space Microelectronics' scientific advisory board. "Even a quick question to Carver often gets us a significantly better solution to a problem in hand," says Thakoor. "He'll point out completely new directions to explore."

Also in 1986, Caltech established its Computation and Neural Systems PhD program, the first and, according to Posner, the most truly cross-disciplinary one of its kind in the world. The program drew faculty from biology, chemistry, engineering and applied science, physics, and mathematics. "It's easier to set these things up here than most places, because of Caltech's small size and fewer layers of bureaucracy," says Posner, "and that's one of our greatest strengths."

That same year, Cole and Allen established a neural-net theory group at JPL. The Neural Computation and Nonlinear Science Group, as it is properly called, is headed by Jacob Barhen, who is also a visiting associate in engineering and applied science on campus. "We have our own critical mass of very good people," says Barhen. "I think that we and Bell Labs are the only two places in the world that have enough nonlinear theorists working together to really make a difference. Some of the papers we have published have really revolutionized the field." The group does basic research in neural-network theory and develops algorithms for specific problems, driving the development of hardware to run them. In general, algorithm design begins with simulations running on an ordinary computer, and many algorithms, still too complex to build into hardware, remain there. The simulation calculates each neuron's output to its mates until the network reaches a stable state. Barhen's group is considerably better off than most, because their simulations run on the Cray X-MP supercomputer recently acquired as a JPL-Caltech joint facility.

The original critical mass lost some of its cohesiveness as more people got involved. Campus folk gravitated to the CNS program, and JPL folk, suddenly blessed with Defense Department money, became involved in their own formal programs. Some collaborations continued, but most people went their separate ways, exploring the terrain. Projects became complementary rather than collaborative. In the last year or so, however, a new generation of collaborations has sprung up to capitalize on the past few years' work. While most have been successful, a few have foundered because, in the words of one campus observer, "You need someone on the JPL end who is absolutely committed to the collaboration for it to work. There's a lot of bureaucratic overhead, mainly because JPL's research is contract-funded, while Caltech grants are usually unrestricted. When you're on contract funding, you have to break everything down into tiny increments, and you spend all your time writing progress reports." With that caveat duly noted, here are some of the successes.

JPL's most ambitious hardware project to

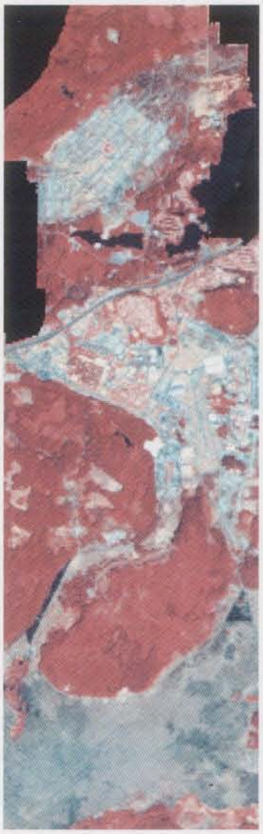

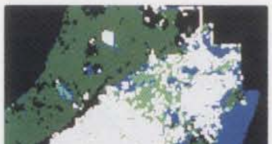

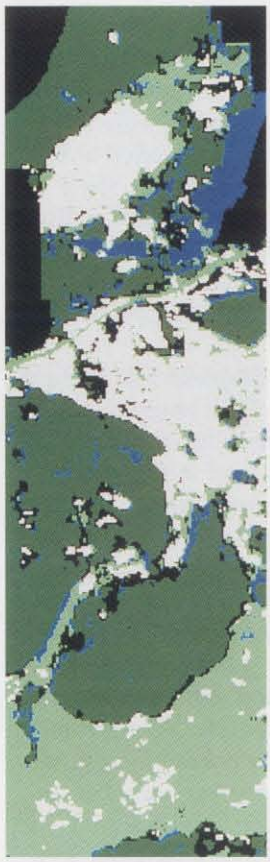

"There's a lot of ad-hockery involved now when you set up a network. We'd like to see how we could make a network grow itself into the appropriate structure as the data arrive."

date could be called an electronic scout. Given a false-color Landsat image of a piece of ground, the system finds the best way to drive a vehicle from one point in the image to another. It's essentially a data-compression problem. Each point, or pixel, of the image consists of 24 bits of information—eight each of red, green, and blue—but we just want one bit of information: can we drive through it or not? Three layers of neurons decide what kind of terrain each of the almost 17 million possible shades of color represents, using algorithms developed by Nevin Bryant, Niles Ritter, and Thomas Logan of JPL's Cartographic Applications Group. The net ultimately determines three components: slope, vegetation type, and load-bearing capacity. These three components reduce to one—the movement cost to pass through that pixel. This output passes to a quasi-neural chip, under development by JPL hardware handyman Silvio Eberhardt with Douglas Kerns, a graduate student of Hopfield's. The chip finds the path from point A to point B that incurs the lowest cumulative movement cost. A serial simulation of the chip works fine. The prototype chip itself, which will process 25 pixels at a crack, is scheduled for fabrication this summer.

On a much more modest scale, Eberhardt, Fernando Pineda (from the Lab's theory group), and Mead-Hopfield grad student Ronald Benson have designed a 4-neuron prototype chip that incorporates the learning algorithm, which normally runs on a digital computer, directly onto the chip. This "recurrent back-propagation" learning algorithm (there are other types) com-

Opposite page: A false-color Landsat Thematic Mapper image (left), combining three visible and one infrared spectral bands, of a portion of Ft. Lewis, Washington. The computer interprets the color, classifying the terrain as either forested (dark green), grassy (light green), urban (white), water (blue), or unclassifiable (black). A conventional program (middle) misread numerous land areas as water or urban, and couldn't figure out several large regions at all. The simulated neural network (right), although given only one-fourth as many correctly classified regions as exemplars, performed considerably better. Both programs had trouble with Nisqually Lake—the seahorse-shaped region in the lower left of each image—whose shallow, luxuriantly vegetated waters gave an ambiguous spectral signature.

pares the network's actual output with the correct output and, with the input still in place, twiddles the connection weights from the output layer on back upstream until the outputs match. These algorithms usually clank through many cycles of software steps, but Pineda saw a way to restate the algorithm as a set of differential equations that could be transcribed directly into silicon. The neurons connect via "floating gate" transistors. The current through a transistor is governed by the amount of charge in the transistor's gate, which would have to be replenished if the gate were connected directly to the rest of the circuit. But a floating gate sits in splendid isolation. Charge injected into it by quantum-mechanical tunneling stays there for months.

The collaboration had a nice balance of forces. Floating gates—which are routine in some chip designs but had not been used extensively in neural nets—were Benson's specialty. He and Eberhardt came up with a circuit design based on Pineda's algorithm. The work began in November, 1989, and the chip was sent out for fabrication in the spring of 1990. Testing will begin in the fall. "Even if this version doesn't learn successfully, it will show us a lot about how the physics in the chip constrains the algorithms that can be put on it. It betters our odds of success with the next one," Pineda says.

Looking farther ahead, Padhraic Smyth (MS '85, PhD '88), of JPL's Communications Systems Research Group, would like to develop faster learning algorithms. Smyth did his graduate work with Associate Professor of Electrical Engineering Rodney Goodman and has kept in close touch with him since moving up the mountain. The two are now coprincipal investigators on a just-launched two-year project that will use techniques from information theory, probability theory, and statistics to try to discover exactly how neural networks learn. A network learns through trial and error, regardless of what algorithm adjusts the connections. But learning, like life, is a risky business. The learning algorithm somehow has to decide what factors in the input are critical to making the right decision. If the algorithm oversimplifies things, the network's "mental image" may not apply to all circumstances, but if the algorithm retains too much complexity and enshrines a host of extraneous factors in the net, it may not function at all. "We'd like to be able to determine the right network architecture just by looking at the data it's going to handle," says Smyth. "There's a lot of ad-hockery involved now when you set up a network. We'd like to see how we could make a network grow itself into the appropriate structure as the data arrive."

Eberhardt and colleagues Taher Daud and Raoul Tawel are working with Caltech Senior Research Associate in Theoretical Physics Tom Gottschalk and former Professor of Theoretical Physics Geoffrey Fox (now at Syracuse University) on a chip to solve "dynamic assignment" problems. The specific problem is this: if there are 25 missiles coming at you, and you have 25 missiles of your own to shoot back, how do you ensure that each good missile shoots down a different bad guy, and doesn't chase all over the sky in the process? You don't need to find the absolute best solution—the shortest path for each missile—in this situation, but you *really* need an answer fast. The neural net should settle into a reasonably good answer in a few millionths of a second. The chip, now in fabrication, will match 64 object pairs—skimming through a space of 2.2×10^{89} possible combinations to do so. Although the chip was developed for the Strategic Defense Initiative program, the problem is a generic one, appearing in such civilian guises as routing calls through a telephone exchange, or ensuring that all the processing units in a parallel computer are sharing the work equably.

Posner, the communications technologist, is applying neural networks directly to communications problems by designing special-purpose nets whose architecture mimics the problem's structure. One of his grad students, Timothy Brown (MS '87, PhD '90), showed in his thesis that a certain neural circuit with inhibitory feedbacks does, in fact, solve the telephone routing problem quite nicely. All the members of one set of

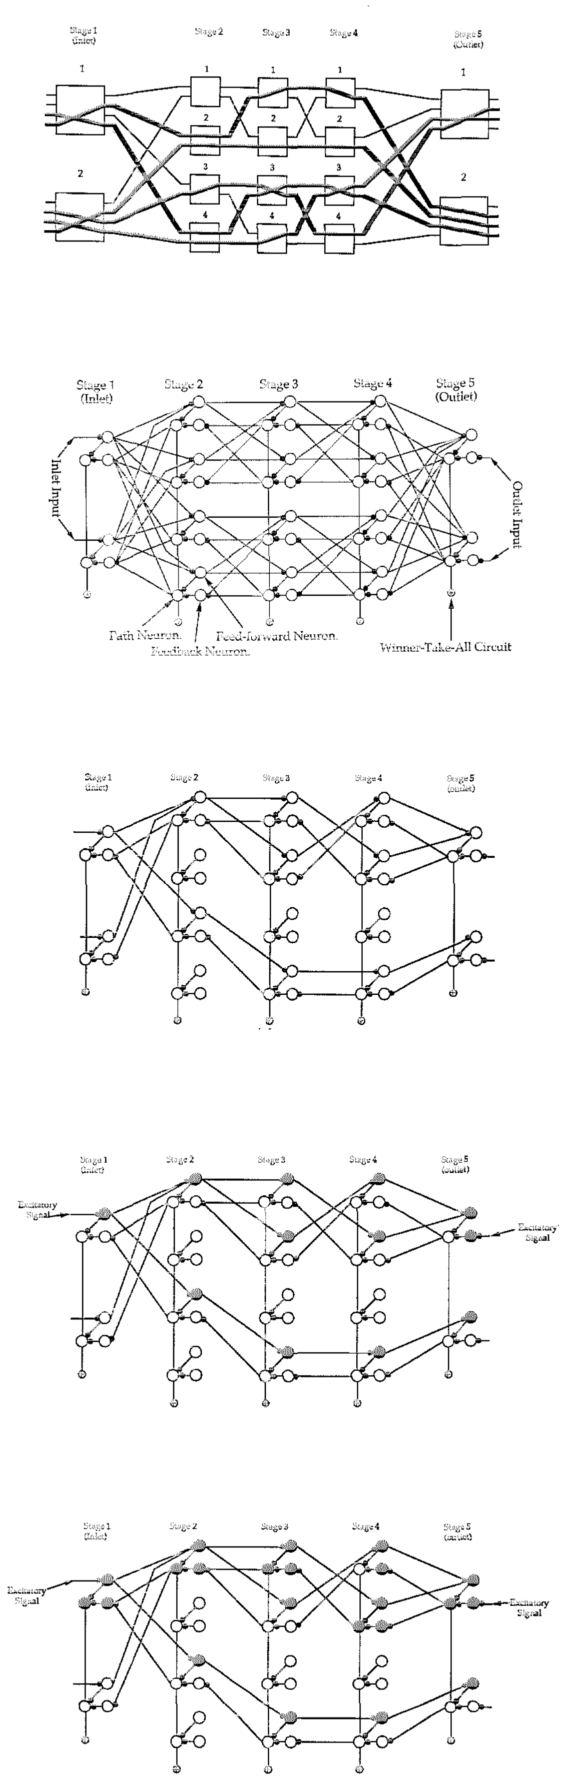

Routing calls through a five-stage telephone exchange. Top: the relays are routing five calls in progress (heavy lines) between inlet-outlet pairs (1,2), (1,2), (2,1), (2,1), and (2,2). Next: A neural-net model of the same relay arrangement. Each neuron is shown as a large open circle with one or more output lines leading from it. A small filled circle is an output connection to an adjacent neuron. "Path Neurons" trace call routes. "Feedforward Neurons" are stimulated by the inlet stage, and in turn stimulate inactive path neurons that could carry the call forward through the network. "Feedback Neurons" are stimulated by the outlet stage, and stimulate inactive path neurons along a route leading back toward the inlet. A path neuron can only become active when stimulated by both a feedforward and a feedback neuron, thus tracing a continuous route through the exchange. "Winner-Take-All Neurons" inhibit competing path neurons, preventing each call from using more than one route. Center: The connections available to route additional calls; these connections correspond to the light lines in the top figure. Next: A call request for (1,1) turns on the feedforward neurons leading from Inlet 1. Active neurons are shown as filled circles. Bottom: The feedback neurons from Outlet 1 turn on, lighting up two available routes for the call. The winner-take-all neuron at Stage 4 arbitrarily chooses to route the call through Relay 2, lighting up one set of path neurons the rest of the way.

neurons are connected to each other in the same way that the exchange's relays are, and set up the call's route. A set of feedback neurons ensures that the routes don't interfere with each other. And a set of "winner-take-all" neurons ensures that each call only gets one route. "Most people are looking to neural nets to solve 'fuzzy' problems, like pattern recognition, where the things in the problem that are critical to solving it aren't well understood," says Brown. "This problem is very well understood. The phone company's computers have been solving it for years. But we can solve it much faster by building the computations right into the hardware. The network doesn't have to learn anything." Brown collaborated with Eberhardt, Daud, and Thakoor over the summer, trying to imbue his algorithm into an assemblage of the hardware group's standard chips. Brown found, as others have before him, that getting an algorithm to "take" in the hardware isn't as easy as it ought to be, but all is not lost—the JPL group was sufficiently impressed to hire him immediately upon his graduation.

And speaking of long-distance calls, there's the Communications Systems Research Section, in charge of developing the coding concepts and hardware and software prototypes that JPL needs to keep in touch with the Voyagers and other far-flung spacecraft. Neural networks, so adept at learning to recognize patterns, could prove useful for the error-correction and data-compression needed to send the data back to Earth.

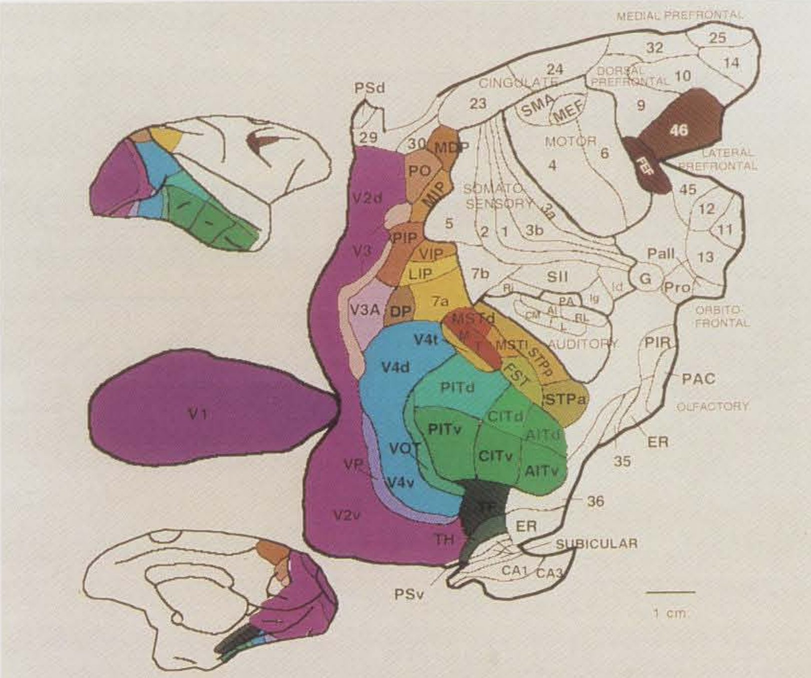

The right hemisphere of a macaque brain, with the visual areas mapped in color, as seen from the right ear (small top figure) and from the left hemisphere (small bottom figure). The main figure shows the cortex unfolded and laid out flat. "V1" is the "primary visual area."

Kar-Ming Cheung (MS '85, PhD '87) and Fabrizio Pollara are working with Goodman on a neural-net data-compressor. JPL engineers used considerable ingenuity to cram all the Voyager data into the narrow communication channel available to it. But Voyager's data stream will be to the torrent of data from the next generation of spacecraft as a dripping faucet is to a fire hose. A single instrument on one of the two Earth Observing System (EOS) craft that will be watching our own planet for signs of global change (see "Observing Earth From Space," *E&S*, Winter '89) will be spewing 300 million bits of information Earthward every second. Voyager's tiny brain could compress data two- or threefold through such stratagems as not transmitting how bright a given pixel in an image was, but rather the difference in brightness from the previous pixel. A neural net might achieve 10 times more compression by handling pixels in rectangular blocks. The network would compare the block to a "code book" of standardized pixel blocks, like matching a wallpaper swatch to a pattern book, and would transmit the index number corresponding to the block that most closely matched the original. A serial computer would take an inordinate length of time to thumb through a code book big enough to guarantee that any input could be matched with minimal distortion, but the answer would tumble right out of a feedforward net whose connection strengths modeled the code book. This project, funded by the same director's discretionary fund that underwrote the Lab's first neural net, got under way last year.

Meanwhile, down on campus, Professor of Biology David Van Essen (BS '67) is interested in neural nets for what they can tell him about real brains. Van Essen began as a traditional neurobiologist interested in primate vision. In 1985, he met physicist Charles Anderson (BS '57), then at RCA. "Charlie was looking at the same problems we were, but from the point of view of a device designer, and this gave him some novel ideas about how the visual system might work. When he moved to JPL in 1987, our collaboration increased in scope." Anderson, a senior member of the technical staff at JPL, is also a visiting associate in biology on campus. He and Van Essen are trying to discover why the world we see doesn't jump and wobble like a movie about to slip off the sprockets. It should—our eyes never sit still. Even when we stare fixedly at something, our eyeballs jink around in tiny involuntary movements. And in binocular depth perception, we judge the distance to an object whose position may vary by only a few seconds of arc between the right-eye and left-eye views. (The headlights on a Cadillac parked 200 miles away are one arc-second apart.) Yet our eyes misalign by as much as one-fifth of a degree, even stone-cold sober. How does the brain remove the gross errors and preserve the subtle differences?

The pathway from the eye to the visual cortex is generally thought of as hard-wired, with signals passing linearly along parallel columns of neurons in order to preserve relative-position information. But any individual retinal cell flickers on and off as the image dances across it, so if each retinal cell had a direct line to a particular cortical neuron, then the cortical "image" would be correspondingly unstable.

Anderson and Van Essen propose a pathway in which signals shift among columns to compensate for a wandering eye. In their neural-net simulation, the columns are sliced and layered like the pepperonis in a stack of frozen pizzas. Each pepperoni sends its output to a set of pepperonis in the pizza above, but not to the one pepperoni directly overhead. The connections go farther afield with each pizza. A set of inhibitory connections within each pizza suppresses shifts in all directions but one, keeping the parts of the image aligned. By tracing the right path through the pepperonis, the image can be kept in a fixed position in the cortex regardless of which retinal cells sent the signal.

The "shifter circuit" hypothesis is still very much in debate, but it does make testable pre-

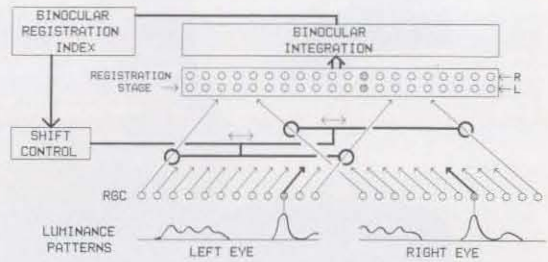

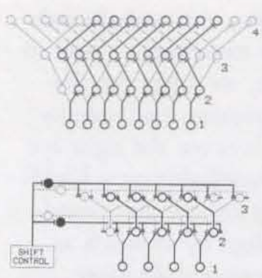

Right: Schematic showing how a shifter circuit could bring misaligned images from the eyes into proper registration in the cortex. The luminance peak from each eye is shifted until both peaks stimulate the same set of cells (hatched). Left: A simple shifter circuit. At every level, each cell stimulates two cells lying in opposite directions in the level above. The shift control suppresses activity along all but one set of paths (heavy lines) to align the final output correctly. Below: A simulated olfactory-cortex oscillation pattern (left to right, top to bottom). Red regions are most active; blue, least. The central trace shows the simulated output from a single neuron.

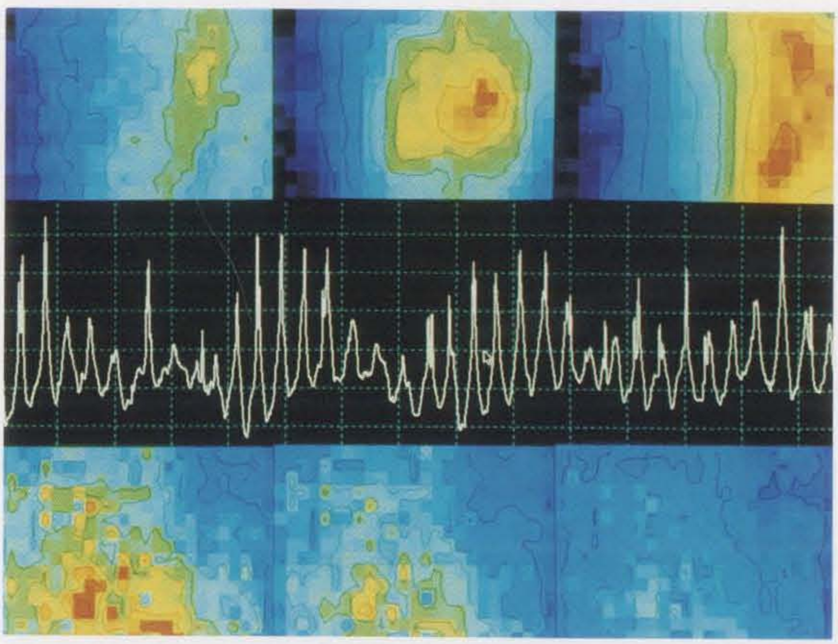

dictions. For example, the image-shifting should occur as early as possible in the visual pathway, and certainly before the images from both eyes are fused for depth perception. Van Essen and grad students James Fox and Tobias Delbrück are studying the "primary visual area," where the first stage of visual processing in the cerebral cortex occurs, using microelectrodes that can localize the source of a nerve impulse to within one-tenth of a millimeter. The predicted shifts should be up to two or three millimeters, and thus readily detectable. Van Essen hopes to have a preliminary result within a year. "We're seeing something interesting going on. It's not exactly what the original theory suggested, but the visual cortex is definitely a more dynamic system than people have heretofore appreciated," he says.

Assistant Professor of Biology James Bower is also "reverse engineering" the nervous system—trying to discover how the brain's complex anatomy actually contributes to its complicated and subtle computations. His group is exploring the olfactory system, which has been mapped in considerable detail and contains elaborate hierarchies of dozens of cell types. The group is particularly interested in the contrasts between the olfactory system and the more elaborate and much more extensively studied visual system, which seems to work quite differently. The visual system reconstructs the three-dimensional world from two two-dimensional images, one on each retina. Every retinal cell responds, sending impulses to the primary visual cortex, where specific cells apparently recognize various attributes. Some cells fire, for example, when they

Right: Coupled random oscillators (below) begin to synchronize their firing patterns (colors) when the textural filters that drive them are shown a pattern (above).
Below: Baldi (holding mouse) and Bhalla use the enclosure in the background for their odor experiments.

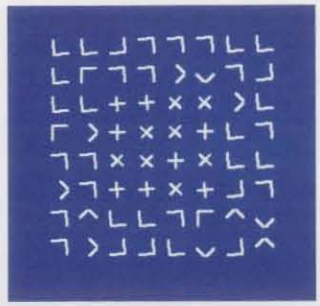

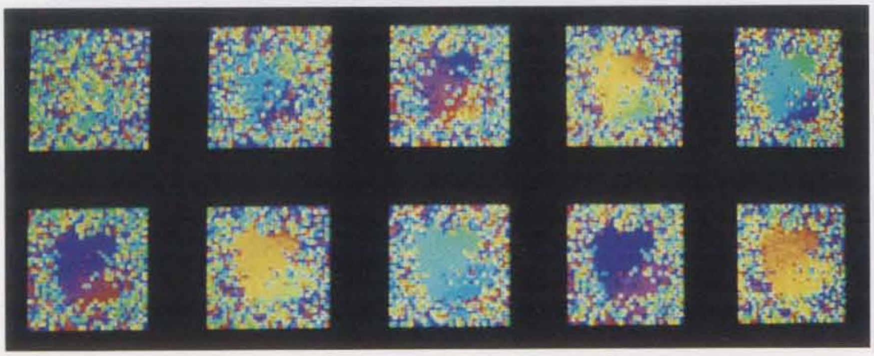

perceive a vertical line, while others are triggered when an object moves from left to right. These attributes are hypothesized to get combined into objects in some complex manner farther on. But an odor has more than three dimensions, says Bower. "If you smell an apple pie fresh from the oven, your nose is sampling a set of volatile chemicals that will be substantially different from the set you'll sample if you smell the same pie after it has been sitting in the fridge for two weeks. But it still smells like apple pie." The membrane lining your nose—the epithelium—contains olfactory receptor cells that recognize and respond uniquely to millions of different volatile chemicals. The impulses travel to the olfactory bulb, where, instead of a particular neuron responding to "lemon" or "pine," many cells respond in some degree to many different inputs.

Pierre Baldi (PhD '86) of the Lab's theory group doubles as a visiting associate in biology, and is working with Bower on the mathematical theory behind a neural-net classification of odors. But Baldi, a theorist with degrees in psychology and mathematics, believes in getting his hands wet, too. "Biological phenomena are too complex for an experimentalist to be able to communicate everything to a theorist," he says. "If you try to be a pure theorist in biology, you'll miss the important details." Thus he, grad student Upinder Bhalla, and Assistant Professor of Biology Kai Zinn have started a set of behavioral experiments with mice. "Rats bite; mice don't," says Bhalla. "That's why we chose them." The trio are looking for a link between the olfactory system and the immune system. (This isn't so farfetched. Both systems recognize and respond to a bewildering variety of foreign substances, so why shouldn't they use similar methods? After all, nature is conservative—a successful stratagem often reappears elsewhere.) The mice learn to push one of two levers, depending on which of two odors wafts into their cage. The experiments will include normal mice and mice with defective T-cell receptors, an immune-system component that recognizes and binds to foreign matter. If the hypothesis is correct, the immunodeficient mice shouldn't be able to recognize as many odors, and an analysis of what they can't smell may reveal how their olfactory neurons are connected.

Baldi's taking a look at vision, too. He and Ron Meir, a postdoc in Hopfield's group, have just published a paper describing how the cortex might use differences in texture to discriminate between an object and its background. Their simulated neural net, which Meir calls "semi-biologically possible," is based on a recent German discovery that groups of neurons in the visual cortex fire simultaneously in "coherent oscillations." These oscillations may be how the cortex defines objects—all the neurons responding to features that are part of a chair would flash at one rate, while the neurons encoding the cat asleep on the chair would flash at a different rate from, or out of phase with, the chair neurons—not unlike having a video game in your head. The model consists of a series of filters tuned to recognize textural elements—vertical bars at a fixed separation, for example—and

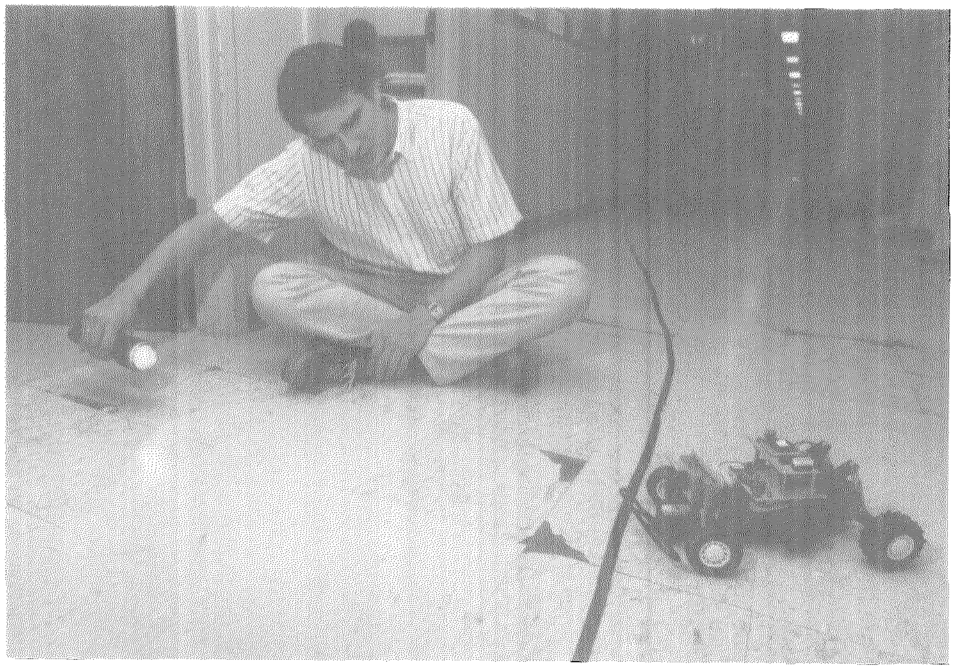

Seated in the hallway outside the CNS lab, Koch tries to entice a light-seeking dune buggy, built by Andrew Hsu (BS '89) as his senior project.

whose outputs drive arrays of coupled oscillators. When a serial-computer simulation of the neural-net model is shown a texture field—a pattern of plus signs on a background of Ls, for example—each oscillator begins to take note of its neighbors, and they spontaneously synchronize over a region corresponding to the pattern. The background remains random.

"These visual-cortex oscillations are very hot right now, because they've just been discovered, but we've known about them in olfaction for about 25 years, and there we think we know what part of the network causes them," says Bower. "We've constructed a biologically realistic simulation of that region with some 200 parameters to it. Pierre is using very abstract models, with three or four parameters, that are more tractable mathematically—exploring the problem unconstrained by biology. The two approaches feed into each other."

"We're essentially trying to put reflexes on chips."

Christof Koch, assistant professor of computation and neural systems, and his group have been working on another way of seeing things for the last four years (see "Computer's Eye View," *E&S*, Winter '88). The group designs "early vision" chips that do such basic jobs as deciding where an object's edges are, or, by calculating how fast those edges are expanding, when a rapidly approaching object will hit. (It's up to other, higher brain centers to identify the objects and figure out what to do about them.) When we look at something, even a boulder with a rough, textured surface, we see a uniform entity with distinct edges. These edges are discontinuities—the different colors of the boulder and the grass; the contrast between the sunlit rock and its shadow on the ground; or (Look out!) the downward motion of the boulder relative to the hillside. Each neuron in the chip corresponds to a pixel in a CCD (charge-coupled device) camera's visual field. The chip recognizes discontinuities—in color, light intensity, or relative motion, depending on the chip—and turns off all the neural connections that span the discontinuities, creating regions on the chip whose sizes and shapes correspond to the objects it "sees."

The group has built three chip generations based on this design, which derives from a retina chip designed by Mead grad student Michelle Mahowald (BS '85) in 1985, and has wired them into little vehicles that began life as radio-controlled toys. These seeing-eye dune buggies do their off-roading in the corridor outside Koch's lab. Although the chips are quite small—from 20 × 20 pixels up to 48 × 48 pixels, compared to some 360,000 in a home video camcorder—they can "see" well enough for the vehicles to zip along a line of black electrical tape on the white-tiled floor, or drive toward a flashlight in a darkened hallway. The group hopes one day to develop a system smart enough to maneuver a vehicle over a three-dimensional landscape. JPL's Brian Wilcox, a member of the vision-system design team for the proposed Mars Rover project, is designing algorithms that recognize and avoid obstacles. "We're essentially trying to put reflexes on chips," says Koch, "decisions that now have to be made by a central processing unit but that should really be made by a

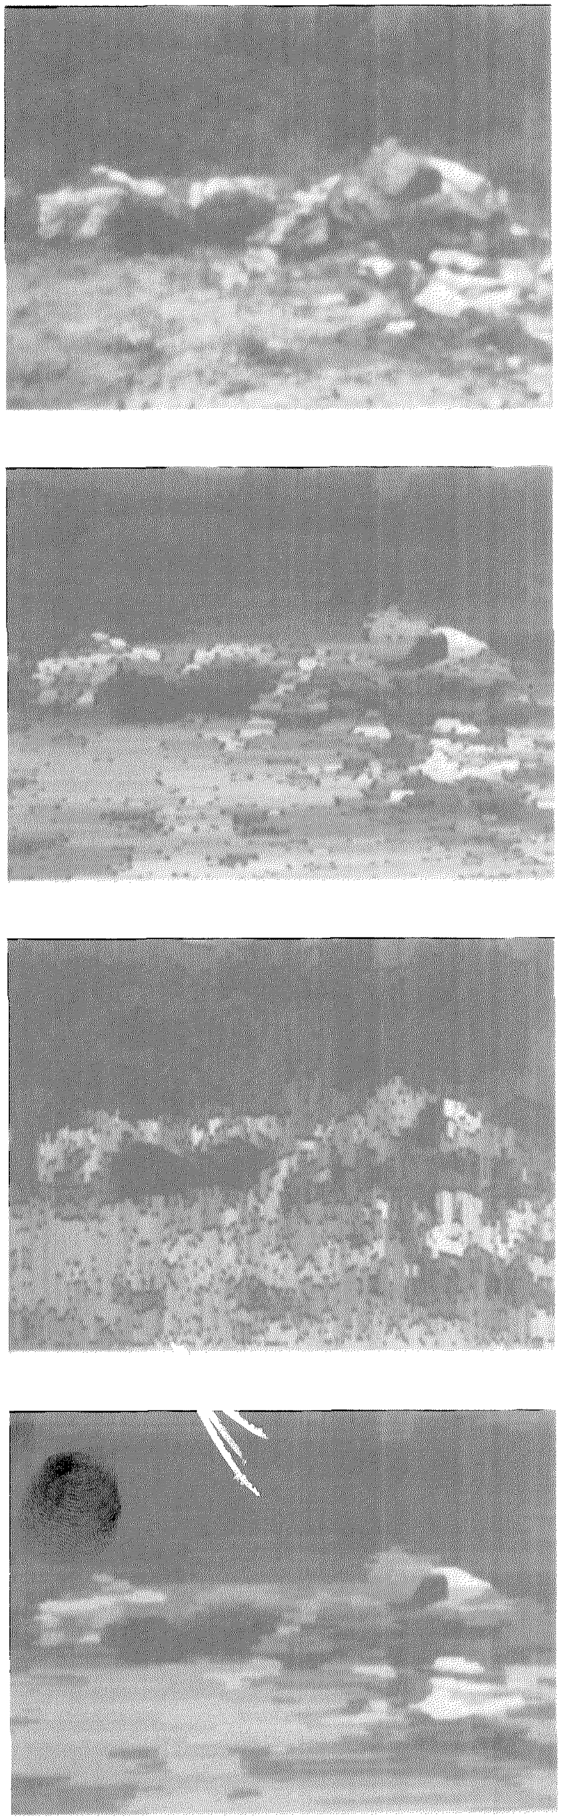

A rocky slope as seen through Lawton's vision algorithms. The original scene is at the top, followed by the set of horizontal line segments, the set of vertical line segments, and then the object map at the bottom.

much lower level in the system."

If the Mars Rover never gets off the ground, there are still plenty of applications closer to home. Edge-detection chips could double-check that a bottling line is really putting two liters of soda in every bottle, or see that toilet paper winds evenly on the roll. And, realistically, "there's a lot bigger market for toilet paper than there is for Mars Rovers," says Koch.

Teri Lawton, of the Lab's theory group and the Mars Rover team, has yet a third perspective on vision. Lawton is using Caltech's CNS lab facilities to design and test "object-oriented" vision algorithms based on biological neural networks. Unlike other, pixel-based approaches, Lawton's algorithms divide a scene into regions with common properties—similar textures and gray-scale values, for example. The algorithms begin by compensating for the jouncing ride over uneven terrain, somewhat as we coordinate our eye and head movements to keep the eyes on one spot as the head moves. (JPL's robotics lab developed these pitch-, heading-, and roll-correction algorithms in the 1970s.) The scenes, now containing just those differences due to the vehicle's real horizontal motion, pass through two sets of filters. One set recognizes horizontal and vertical line segments. The other set registers gray-scale brightness. The gray scale automatically adjusts itself within shadows—which other algorithms perceive as flat, dark objects—to reveal smaller rocks that could wreck a rover. The algorithm then defines and remembers objects as two-dimensional closed loops made of overlapping line segments of roughly the same gray scale, penciling in obscured line segments where necessary. As the rover rolls along, it constantly updates its memory. Parallax—how much an object moves compared to the rest of the scene—gives each object depth, and allows the rover to build a three-dimensional picture, or "object map," of its surroundings. Because this picture includes several sensory dimensions—size, position, texture, illumination, gray scale, and motion parallax—it's more fault-tolerant than object maps depending on one parameter, such as edges, alone; and because the algorithm exploits the relative motion of whole objects in successive images, it handles diverse objects better and operates faster than pixel-based methods that must compute the change at every pixel between successive images. Lawton's simulations take about 12 seconds to generate a depth map, much faster than other algorithms, and the chip she eventually hopes to build with Mead's group should do the same job as quick as a wink. In the meantime, Lawton has been working—first with Brian Fox, a staff member in Koch's group, and now with Aaron Emigh, a senior from UC Santa Cruz on campus for the summer as part of Caltech's Summer Undergraduate Research Fellowship (SURF) program—to optimize the algorithm before committing it to silicon. "Caltech has made it possible for me to do this work," says Lawton. "I couldn't have done it otherwise. And it's a contribution to biology as well—we can use these three-dimensional terrain maps of natural scenes as a test bed to learn how the brain generates a three-dimensional world view from a two-dimensional retinal image."

Lawton's work can help the partially sighted on Earth as well, especially those people—mostly elderly—who can't see fine detail any more. Closed-circuit TV "readers" that magnify and brighten printed matter are already available. Readers modified with one of Lawton's algorithms automatically enhance the text to match the user's remaining contrast sensitivity, and render the text in shades of gray more easily perceivable than black and white. Users have experienced a two- to fourfold increase in reading speed at up to 70 percent less magnification than they needed to read text on the old machine. Caltech has optioned the patents to the gray scale portion of the system, which could be commercially available soon. The next logical step, electronic spectacles (lightweight, wearable units with CCD cameras to look at the world, some simple electronics to process the image, and liquid-crystal displays—LCDs, the screens used in tiny TVs—to present the result) is well within the reach of current technology.

Hybrid neural nets using electronic logic and optical interconnects may be practical shortly.

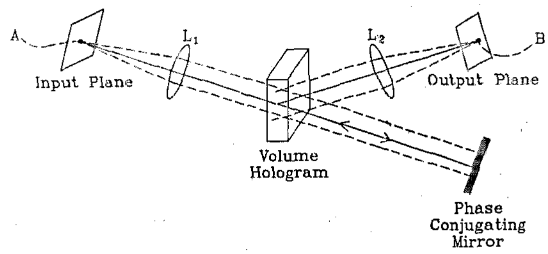

General scheme of a holographic associative memory. L_1 and L_2 are lenses. The "Phase Conjugating Mirror" is used when updating the memory.

Although connections make the net, you can't pass wires through each other. Beams of light, though, can intersect and carry information without interfering with one another. Optics and electronics can mate, with LCDs or LEDs (light-emitting diodes) converting electrons into photons, while photovoltaic or photoconductive devices convert light back into electricity. The optical equivalents of logical processors are still rudimentary, so full-fledged optical computers are still a gleam in the eye, but hybrid neural nets using electronic logic and optical interconnects may be practical shortly. Professor of Electrical Engineering Demetri Psaltis, who maintains an office at JPL and consults there one day a week, thinks that "these hybrid systems will prove very useful in the next few years. Neural nets are ideally suited to coping with the real world, in robotic control, for instance. This way you can have a very highly interconnected net that interprets sensory inputs and passes its conclusions on to a serial controller that decides what to do. And you could have another net taking the controller's output to actually guide the robot."

Psaltis's group is working on hybrids. Grad student Steven Lin is collaborating with Jae Kim of JPL's Microdevices Lab to build neural arrays of gallium arsenide, a semiconductor faster than silicon that can also emit light. (Both semiconductors can be photodetectors.) Each neuron incorporates a light detector and a light emitter. The emitters shine up from the chip's surface into a hologram. The hologram can be two-dimensional—an optical disk, like a CD—or

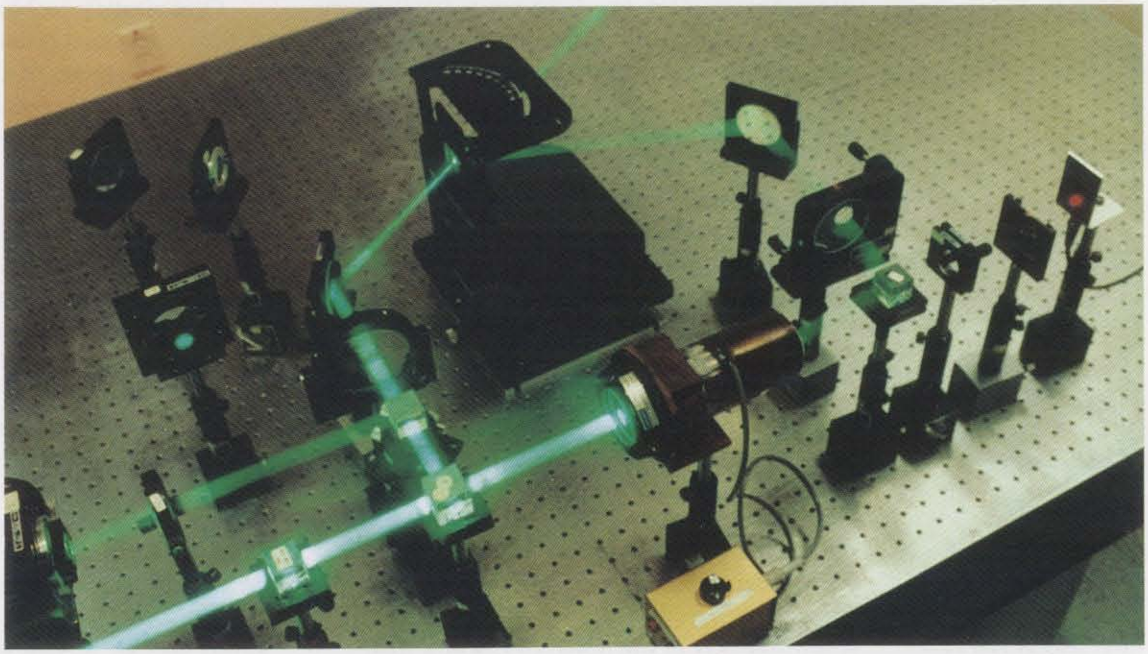

Above: A holographic-memory loop. Input comes from the far right, where the red light illuminates a transparency, projecting its pattern into one end of a liquid-crystal light valve—the flashlight-shaped object at center. A laser beam from the lower left is reflected off the valve's other end according to the pattern. A cube-shaped beam splitter diverts the patterned beam back to a holographic medium in the angle-calibrated mounting at rear. The hologram's output emerges at an angle and goes back to the light valve's input side to complete the feedback loop. Thus the light valve acts as the set of neurons, using an external input and the product of its own interconnections to generate an output. The output registers on a CCD camera at lower left, behind the incoming laser beam. The rest of the setup is used for training the memory.
Left: Arabic and Chinese numeral input-output pairs stored holographically.

a three-dimensional photorefractive crystal. (Such a crystal's refractive index—the degree to which it bends light—is itself light-sensitive. A powerful beam of the right frequency alters the crystal's electronic structure, and thus its refractive index. The change persists after the beam is gone.) The chip-produced hologram channels light from each emitter to each detector in proportion to the connection strength between those two neurons.

Disks are easier to work with at the moment, the technology being more mature, but they aren't really reprogrammable yet. Grad student Alan Yamamura is using disks to make a single layer of neurons act like a multilayer network. The disk stores each layer's connection strengths sequentially and spins in sync with the information flow from layer to layer. JPLer Jeffrey Yu (BS '83, MS '84, PhD '88, and a former student of Psaltis's) is working with Psaltis to apply this technique to image recognition.

The crystals are fully reprogrammable and, being 3-D, can store information more compactly. A crystal can be loaded holographically, for example, so that shining an Arabic numeral onto one face causes the corresponding Chinese numeral to shine out from another face. Scientists elsewhere have recorded more than 1000 such associations on a crystal. In theory, a crystal can store as many as several thousand images per cubic centimeter, versus the tens of thousands of images that would cover a five-inch disk. Even when crystal technology matures, however, it may not displace disks altogether. Crystal memories can fade as new memories are

"Our current models don't allow you to remove just one association without affecting all the others, but animals do it all the time. It's the only way to cope with a complex, constantly changing world."

stored because each new light beam irradiates the entire crystal, partially obliterating its predecessors' traces. But a tightly focused laser writes memories on a disk with plenty of elbow room between them.

Meanwhile, the connection problem may have been solved by Senior Research Fellow in Applied Physics Aharon Agranat and grad student Charles Neugebauer (BS '88) in the group led by Amnon Yariv, the Myers Professor of Electrical Engineering and professor of applied physics. They have a chip that uses a CCD to store connection weights—a radical departure from its designed use as a light sensor. A row and a column of neurons adjoin the CCD, each pixel of which contains a dollop of electrons proportional to the connection strength between the corresponding row neuron and column neuron. (See "Photographic Memory," *E&S*, Winter '88.) The current version has 256 neurons, each of which connects to the other 255, and a thousand-neuron chip is well within reach of standard CCD technology. Agranat and Neugebauer are now building a computer board that will carry the chip and that can be plugged into any IBM-PC-compatible computer. Real neural nets, instead of just software simulations, will become accessible to thousands of researchers.

Yariv's group began collaborating with Barhen's group this year to see how easily their hardware and algorithms integrate. Their first project will be an algorithm to calculate discrete Fourier and Hartley transforms—the two most important (and, coincidentally, most computationally intensive) mathematical tools used in signal processing. The system could be used to process seismic data or hunt for gravitational waves, and might also come in handy in JPL's image-processing work.

And then there's robotics. JPL has been doing robotics all along, of course—strictly speaking, any autonomous spacecraft is a robot—but the Lab is also working on more traditional robotics problems. Joel Burdick, assistant professor of mechanical engineering, is starting several collaborations between his graduate students and various robotics groups on Lab. One student, Bedri Cetin, is working with Barhen to apply neural nets to optimization problems such as making a robot arm move efficiently. Cetin developed a new approach to the problem, based on recent work by Barhen and fellow group members Nikzad Toomarian and Michail Zak, that Barhen calls "a major breakthrough in optimization theory. Everything eventually becomes an optimization problem, so the payoff will be tremendous."

Optimization problems can be thought of as rugged landscapes of hills and valleys. Whatever the physical aspects of the problem—moving a robot arm around obstacles, finding the shortest route through 11 cities—the problem can be cast as a mathematical landscape (in more than three dimensions, if need be), wherein the lowest point in the deepest depression is the optimum answer. Finding this nadir is the mathematical equivalent of setting a boulder loose and waiting for it to come to rest. Various strategies, such as dropping several boulders all across the landscape, have been developed to ensure that you really do find the very deepest point. The breakthrough incorporates ideas from quantum mechanics—the boulder can "tunnel" through a mathematical hillside to escape from an exitless valley—and from nonlinear dynamic systems theory, wherein a newly discovered entity called a "terminal repeller" can suddenly give the boulder a shove strong enough to send it skittering to anyplace in the landscape. "This method has solved some standard optimization problems 100 to 1,000 times faster than the best competing methods," says Barhen. "And applying the terminal repeller concept to man-made neural nets allows them to do things they couldn't do before, like selectively forgetting old associations, or spontaneously creating new ones without extensive training. Our current models don't let you remove just one association without affecting all the others, but animals do it all the time. It's the only way to cope with a complex, constantly changing world."

A lot of people are trying to help robots cope with the real world. Robots to date have been pretty simple-minded creatures. Today's state-of-the-art industrial robot—or spacecraft, for that matter—is really more like a complex machine tool. It has to have nearly every gesture spelled out for it explicitly, and must work in a simple environment in which a few known objects occupy predetermined locations and everything else stays out of the way. But future NASA robots, the ones that will go day-tripping across other worlds or work on the space station, will have to think for themselves and adapt to a complex, changeable environment.

Carl Ruoff, a longtime member of the Lab's Robotics and Automation Section and now a graduate student at Caltech as well, is working with Professor of Mechanical Engineering Fred Culick on a rudimentary robotic Little Leaguer that can learn basic motor skills on its own in a simplified version of such an environment. The device will acquire hand-eye coordination: it will learn to catch (or hit) any ball—from a golf ball

Hollywood epics notwithstanding, day labor at 380 miles up is difficult, dangerous, and time-consuming.

to a beach ball—no matter how the ball is thrown. The robot will have to integrate artificial-vision data about its surroundings with tactile-, force-, and body-sensor data about itself. It will learn coordination the way kids do—first learning how to control its arm, and how the arm should "feel" in particular situations, then learning how the ball behaves. Once the robot has watched enough tosses to be able to predict their trajectories, it learns to intercept them. The team expects to have a simulation running on the JPL–Caltech Cray come August. It's a long way from a ball-catching machine to RoboCub. But the Little Leaguer, when operational, will embody some of the basic attributes that autonomous space robots will need.

Burdick, Barhen, Sandeep Gulati (also of the theory group), and Robotics and Automation Section members Charles Weisbin, Subramanian Venkataraman, Guillermo Rodriguez, and Hamayoun Seraji (who is also a lecturer in mechanical engineering on campus), have started an informal collaboration to think about the rest of these attributes. "The whole task of integrating all these functions—motor skills, sensory processing, memory, and a host of other things—into a system that can learn on the job and make the internal changes it needs to complete its mission, is extremely ambitious," says Gulati. Adds Venkataraman, "It's learning to adapt a mastered skill to different environments, unlike today's robots that would have to start from scratch with every new situation."

The group has chosen to work on an astronaut's apprentice as their demonstration project. Astronauts will be spending a lot of time outdoors in the next century, working on the space station and making service calls on satellites. Hollywood epics notwithstanding, day labor at 380 miles up is difficult, dangerous, and time-consuming. A buddy can't just toss you a Phillips-head screwdriver, for one thing. So JPL envisions self-propelled, voice-controlled robot gofers to fetch tools, maneuver bulky parts and hold them in position, and rescue free-floating objects (including astronauts) before they drift away. A helper taking orders from a human in this situation actually faces an environment more complex than does a solitary explorer picking its way among Martian crevasses to take rock samples, because the helper has to be aware of many objects in three dimensions traveling in all directions at once, including unpredictable humans that will blunder into its way.

Such a robot will need all the neural-network attributes described in this article and then some. It will need pattern-recognition skills and a flexible memory to understand spoken commands issued by many voices, acute vision and deft limbs to execute those commands, and a sophisticated "brain" that can plan complex tasks in a free-form environment.

Many years will pass before such a system can be built, but the group is planning to take the first step. Over the next two years, they propose to develop a system that can deal with uncertainty in a limited environment. The device, initially two robot arms bolted to the floor, will grasp one end of some large, perhaps flexible, object. A person would hold the other end, and a tug-of-war would ensue. The human would push and pull on the object, shift grips, and sometimes let go altogether. The robot would try to keep its end level at all times, and would have to adjust its response constantly to compensate for the human's actions.

There's a long way to go before an autonomous, adaptive robot's gray matter can be trusted in space. "Real biological networks have much complex internal structure that we don't understand," says Ruoff. "Large, complicated systems are really qualitatively different," adds Culick. "Building lots of little pieces and having them all work separately is one thing, and putting an integrated system together and making it work is quite another. It is, however, something that JPL has learned to do very well." The neural net or hybrid neural-serial system that ultimately results—if one does—may finally be the mosquito's intellectual equal. Then it may fairly be said that the Caltech–JPL connection will have helped neural nets come of age.□—*DS*

Tunnel Vision

"We think of this little beam of ballistic electrons as a tiny searchlight. It goes down through the metal . . . and we can use it to illuminate the interface structure."

Ballistic electron emission microscopy illuminates the interface between gold and a gallium arsenide substrate. (From the JPL Digital Image Animation Lab's production of "Gallium Arsenide: The Movie".)

In an article in the February 1960 *E&S*, which has been massively photocopied and recirculated in the intervening decades, the late Richard Feynman prophesied a new field of physics—"manipulating and controlling things on a small scale." He meant a very small scale—angstroms (ten-billionths of a meter). The article, "There's Plenty of Room at the Bottom," described this new field as "not quite the same as the others in that it will not tell us much of fundamental physics (in the sense of, 'What are the strange particles?'), but it is more like solid-state physics in the sense that it might tell us much of great interest about the strange phenomena that occur in complex situations. Furthermore, a point that is most important is that it would have an enormous number of technical applications."

Feynman's predictions did come true; nanotechnology is very big, so to speak, these days, and is indeed finding "an enormous number of technical applications." But nowhere is this dramatic downsizing more pertinent, perhaps, than in space—at least in the things human beings send into space. To exploit these possibilities JPL's Center for Space Microelectronics Technology (CSMT) was established in January 1987, under the directorship of Carl Kukkonen.

CSMT (pronounced Kismet) is divided into four research areas: photonics, custom microcircuits, computer architecture, and solid-state devices. Photonics (optoelectronics) involves devices that marry laser and integrated-circuit technologies; such devices can generate, detect, modulate, and switch electromagnetic radiation and in space can be used for communication, guidance and control, and robotic vision. The custom microcircuits program develops specialized chips for communication and for image and signal processing; it's also investigating what happens to microcircuits subjected to the ionizing radiation in space. CSMT's efforts in computer architecture include groups working on neural networks (see the article beginning on page 4) and in parallel computing (see page 28).

Research on solid-state devices (such as sensors for the very-far-infrared and submillimeter portions of the electromagnetic spectrum, and other electronics for space) is housed, along with the photonics group, in the Microdevices Laboratory, built just last year. The NASA-funded laboratory contains state-of-the-art equipment for fabricating and characterizing both semiconductors and superconductors. Its laboratories and various categories of clean rooms (located on the vibration-isolated floors necessary for submicron device fabrication) are home to a bevy of advanced instruments and techniques, among them scanning tunneling microscopy.

Scanning tunneling microscopy (STM), which won the 1986 Nobel Prize for its inventors, Gerd Binnig and Heinrich Rohrer, is scarcely a decade old, although the phenomenon was predicted by quantum mechanics and the idea for building such an instrument has been around since the thirties. Its realization had to wait until technology was capable of bringing two surfaces together and holding them at a constant separation of just a few atom diameters. According to

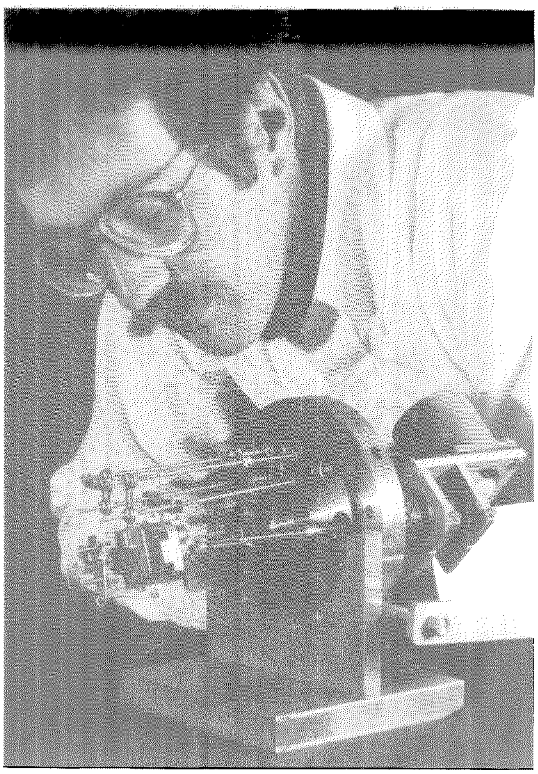

The ballistic electron emission microscope (BEEM), shown here with Bill Kaiser, provides a new view below the surface. Below: The surface of a gold/gallium arsenide structure is revealed by STM (top) and the interface of the two materials by BEEM (bottom).

"A major fraction of the world economy depends upon the silicon-silicon dioxide interface."

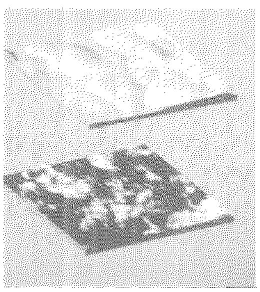

quantum mechanics, when you do that with two metal electrodes in a vacuum, there is some probability that a few of the electrons that are propagating back and forth at high velocity in the electrodes will leak out into the vacuum; they won't all be trapped inside as classical physics would have it. Of course, they don't leak very far. The electron distribution in the vacuum is a function of distance and it decays to zero very quickly—about a factor of 10 for every angstrom out from the surface. But if you have these little clouds of electrons floating above two surfaces that are, say, 10 angstroms apart (about three atom diameters), some of those electron probabilities will overlap, allowing electrons to "tunnel" through the vacuum and across to the other surface. Although tunneling would normally occur in both directions, applying a voltage difference to one of the electrodes unbalances this symmetry, giving more of the electrons on one side the kick to jump the potential-energy gap between the metal and the vacuum, and prompting a one-way flow. So an electrical current is created, even though classical physics would say the circuit is open.

In a scanning tunneling microscope, one of these electrodes is a very sharp metal tip; the other is a conducting surface, say a semiconductor, to be studied. A piezoelectric servomechanism (one that can change its dimensions in response to voltage) uses feedback-controlled tunnel current to keep the tip hovering at a constant altitude of a few angstroms above the sample. Because the tunneling phenomenon decays so quickly, it's extraordinarily sensitive to changes in distance of even a fraction of an angstrom; so measuring the tunneling current while scanning with the tunnel tip can provide an atomic-resolution image. The current flows only through the atom on the very end of the tip, making the STM an extremely sensitive position detector. When the tip scans across the layer of atoms on the sample surface, its trajectory follows the contours of the individual atoms and reveals the electronic structure of the surface.

Bill Kaiser, now a senior research scientist at JPL, did some of the early work in studying surfaces by STM. But even more interesting than surfaces are the inaccessible interfaces below the surface, for example, where a semiconductor and a metal (or two semiconductor materials) meet. "A major fraction of the world economy depends upon the silicon-silicon dioxide interface," says Kaiser.

To get down to this interface, known as the Schottky barrier, Kaiser (along with Doug Bell and Michael Hecht at JPL) invented BEEM—ballistic electron emission microscopy. BEEM gets a little more mileage (angstromage?) out of the tunneling electrons. To accomplish this, Kaiser has made use of an unusual phenomenon that is "almost always neglected in tunneling, but it's always operating." That is, the electrons don't stop when they hit the other electrode surface—some of them scatter, losing energy in collisions, but some of them continue to propagate through the electrode with all the energy they arrived with. They can actually travel a relatively long distance in microscopic terms—up to several hundred angstroms—before petering

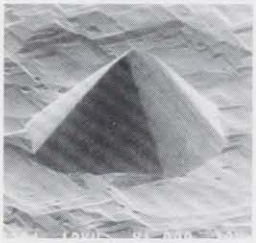

The tunnel sensor consists of three tiny (the scale is 1 inch) silicon plates sandwiched together (but shown separately here). The bottom plate incorporates the folded gold springs on either side of the gold pad, and the tunnel tip, barely visible as a black dot on the horizontal arm above the pad (and magnified at left). Any slight vibration will jiggle the spring-supported plate and disturb the current from the tunnel tip to the opposite electrode (upper left). The third plate is a spacer between them.

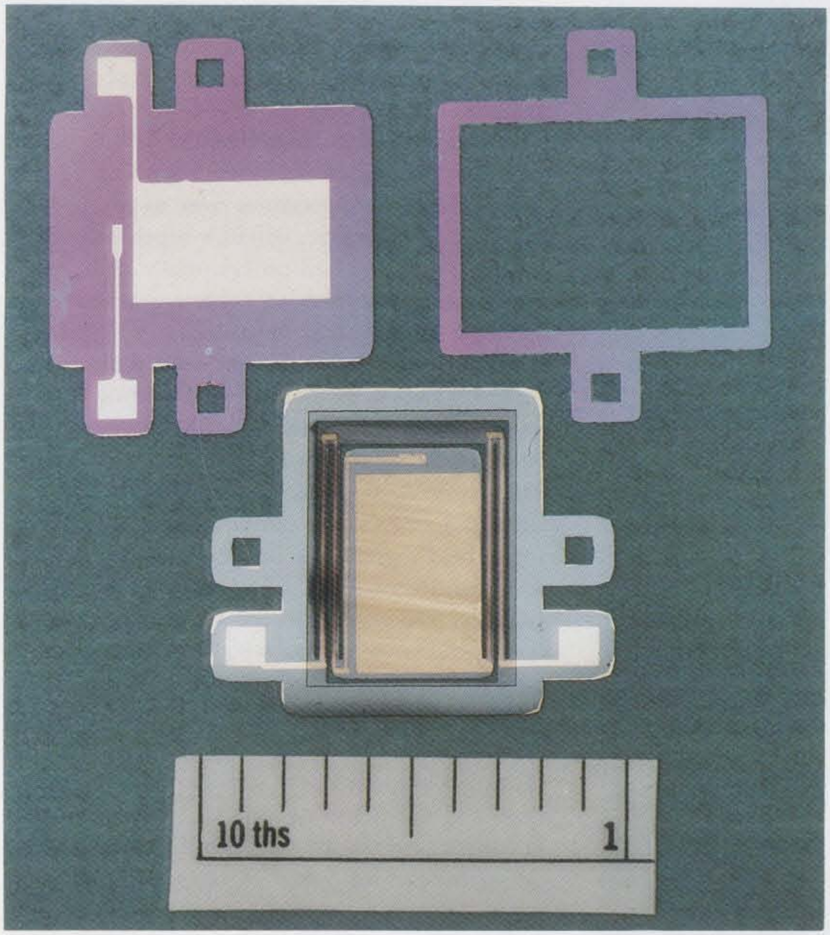

out. Since the top layer of a semiconductor is 100–200 angstroms thick, the electrons can quite easily make it to the buried interface. A little extra voltage difference injects these electrons "ballistically," giving them enough oomph to shoot straight through to the Schottky barrier and surmount the potential-energy step, similar to the one between electrode and vacuum, that exists between metal and semiconductor. To do this the BEEM device requires a third electrode (where an STM has only two)—a thin metal film deposited on the back of the semiconductor target to collect the electrons that make it through the Schottky barrier. Varying the tunnel voltage controls the energy of the ballistic electrons, which allows detailed spectroscopic analysis of the fundamental interface properties as a function of electron energy.

"We think of this little beam of ballistic electrons as a tiny searchlight," says Kaiser. "It goes down through the metal even though the metal is normally opaque, and we can use it to illuminate the interface structure." What they have seen has also been illuminating, and scientists who have been studying invisible Schottky barriers for 20 years are doing some revising now that they can actually see them. One colleague claimed that BEEM has set the field back ten years, because it revealed that the semiconductor interface was not nearly as simple as had been assumed. Beyond semiconductors, BEEM makes it possible to bring to light the properties of all kinds of buried interfaces that have hitherto been hidden from direct study. BEEM research has grown rapidly and is now being pursued at many laboratories in the United States, Europe, and Japan; the first international BEEM workshop was held at JPL in March with more than 70 participants.

Kaiser, Steven Waltman, and Tom Kenny have also adapted the tunneling phenomenon to a very sensitive motion detector—a tiny silicon sandwich that functions as an accelerometer. The tunnel sensor is accurate to 1 nano-g (10^{-9} g) and can be fabricated as a 50 × 50-micron square on a silicon chip. "Typically a nano-g accelerometer is something the size of several shoe boxes," says Kaiser, "and weighs tens of pounds and has expensive electronics associated with it."

Three parallel plates—micro-machined from single crystals of silicon—form the guts of the tunnel sensor. The bottom plate contains a microscopic gold STM tip and tiny, folded, cantilever springs, also of gold. This plate levitates at the bottom of the sandwich, and when the plates move within tunneling range, current flows


John Baldeschwieler's group made this STM image of DNA (far left), magnified approximately 25 million times, the sharpest image ever obtained of the molecule. It's compared to a computer model of DNA. Below: Baldeschwieler (right) and Mike Youngquist examine an STM instrument designed and built by Youngquist to operate immersed in liquid helium.

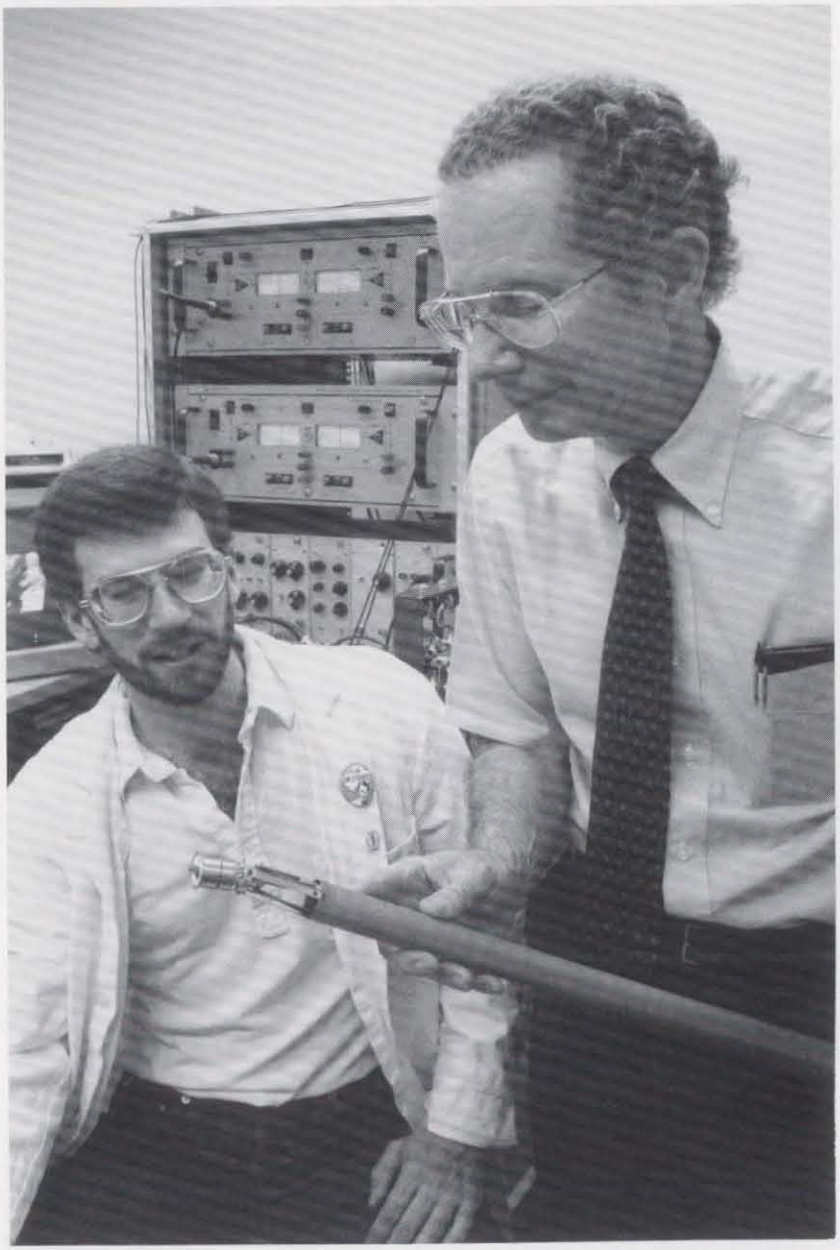

from the STM tip to a gold pad on the top plate (the middle plate serves as a spacer). Instead of positioning these plates with a piezoelectric drive (which has thermal drift and noise detrimental to a sensor) as in STM, these plates are brought within tunneling range of each other and held there by an electrostatic field; applying a voltage creates an attractive force that moves them closer together. Any slight vibration, say a seismic wave, will cause the spring-supported plate to wobble, changing the tunnel current. Inexpensive and easy to manufacture with current lithography techniques, the tunnel sensor could be used for miniature seismometers, pressure sensors, microphones, and a variety of applications in space and on the ground. The JPL group is also investigating using it as the transducer in a pneumatic infrared detector.

Kaiser's innovative adaptations of STM technology has also kindled a campus collaboration with Professor of Chemistry John Baldeschwieler, who is interested in using the technique to look at molecular structure. "We couldn't have done it without Kaiser, because the support from his group in terms of designing this kind of instrumentation was crucial. It would have taken us years to learn," says Baldeschwieler.

What Baldeschwieler needed was an STM that worked in temperatures close to absolute zero, and Rick LeDuc, of JPL's superconductivity group, and Kaiser had providentially built just such a device for studying superconductors. It operated perfectly well immersed in a dewar of liquid helium. Baldeschwieler needed low temperatures (which give tunneling electrons a narrower energy distribution) for an experiment in "inelastic electron tunneling spectroscopy," looking at how electrons interact with molecular vibrations. Molecules can be thought of as atoms held together by little springs that vibrate at certain frequencies. Chemists observe the radiation a sample emits (or absorbs) at those frequencies to determine what molecules are present in a complicated mixture. When an electron collides inelastically with a surface molecule, some of the electron's energy goes into exciting vibrations in the springs. This causes a change in the tunneling current. These changes in current provide the vibrational spectrum of the molecule with a field of view of a single atom. Chemists are interested in molecular surface structures because the reactive sites on catalyst surfaces are often an atomic step or dislocation that changes the reactivity of a molecule that binds to it.

Baldeschwieler and Kaiser are coprincipal investigators on the NSF-funded experiment,

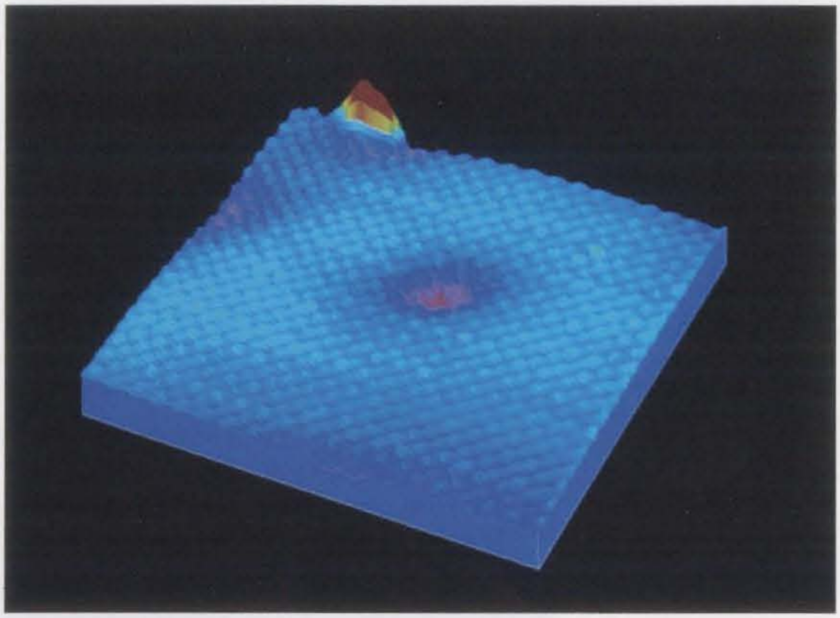

Top right: STM reveals the surface atomic lattice of gallium arsenide. The 125-angstrom-square image, made by grad student Robert Driscoll, resolves the gallium atoms along with an adsorbed electronegative defect (magenta depression in center). The height of the atomic corrugation is 0.03 angstroms.
Below right: Youngquist obtained this large-scale topographic image of graphite, 7,500 angstroms square, on the first test run at room temperature of his low-temperature STM. The vertical scale is exaggerated to enhance surface features. Smaller-scale images (not shown) resolved individual carbon atoms. Youngquist also built a functioning STM out of Lego blocks.

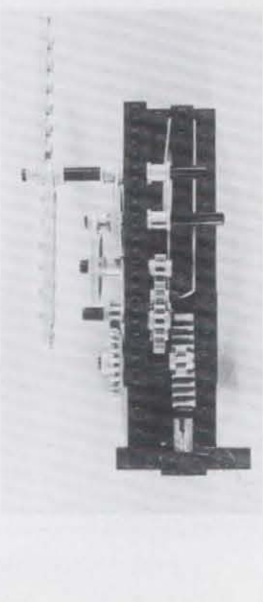

and have collaborated both on the experiment's design and its theory of operation. "This is a project where there's a significant sharing of time, energy, and talent," according to Baldeschwieler. Initially Mike Youngquist, a grad student in Baldeschwieler's lab, spent several months up at JPL studying the technology with LeDuc. Youngquist has since built several prototypes of his own (including the functional but facetious Lego-block instrument shown in the photo at left), and the Caltech group has significantly advanced low-temperature STM instrumentation.

They are now testing the instruments by looking at well-known molecules. They have achieved atomic-resolution images of graphite with their system, and plan to look at the stretching of the bonds holding hydrogen atoms to silicon (in a thin layer of silicon hydride on the surface), an experiment that promises a good chance of success. "This is the most straightforward experiment we can design at the moment, and that's because the hydrogen will cover the whole surface; we won't have any trouble finding it. We'll be looking for the appearance of the inelastic tunneling transition, and then we'll verify that this is real by substituting deuterium for hydrogen, so the vibrational frequency should shift. This whole experiment is very difficult. In the beginning it's a challenge to prove that it's even working. If it does work, we'll be able to understand what the sensitivity of the technique really is, and then, of course, optimize it for molecules that are of chemical interest."
□ *—JD*

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.

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.

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.

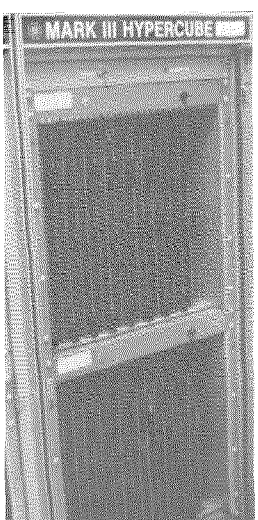

"We're connected to everything," says Messina, who spends much of his time on computer networks. Caltech has a high-speed line to the San Diego Supercomputer Center's very large Cray, funded by the National Science Foundation. Caltech has recently leased a Cray for JPL's Supercomputing Project; JPL takes 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

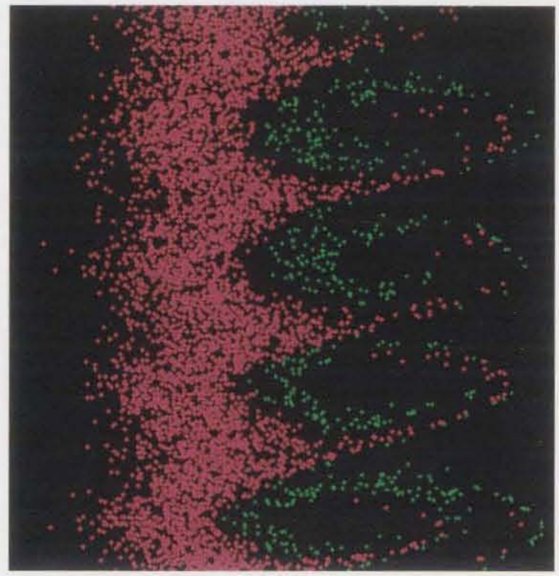

In modeling the behavior of a plasma when a beam of electrons is shot through it, the Mark IIIfp tracked the position (y-axis) and velocity (x-axis) of the electrons. The magenta dots represent the electrons of the background plasma, which tries to eliminate the electron beam by trapping its electrons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

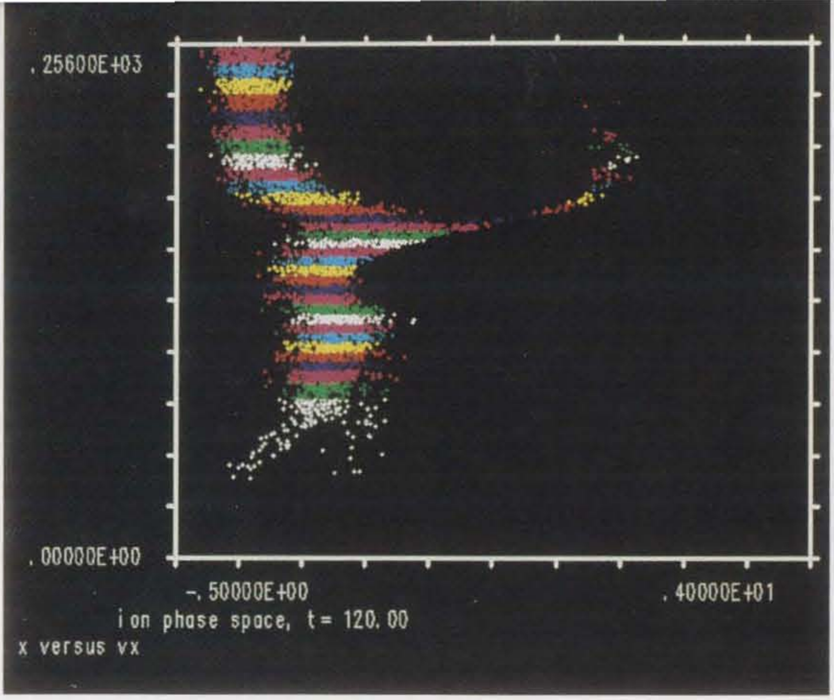

A simulation of the earth's bowshock shows the solar-wind ions streaming in from the top and slowing down as they hit the shock about a third of the way down. The vertical axis represents ion position, and the horizontal axis, velocity. Each color stands for ions in the region of space tracked by one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

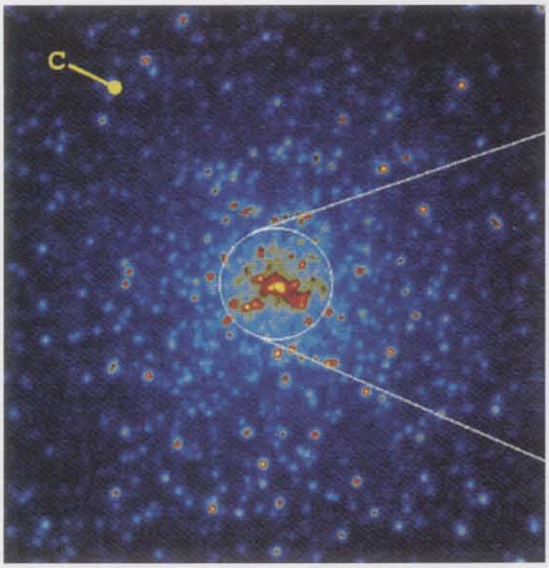

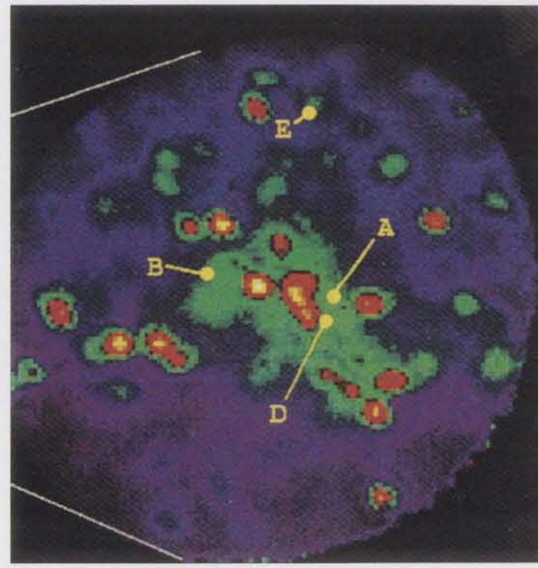

So far 19 such pulsars have been discovered worldwide, and 8 of them have been found using the Caltech supercomputing facility.

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-

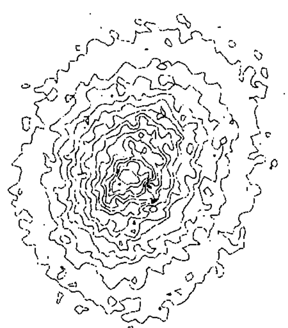

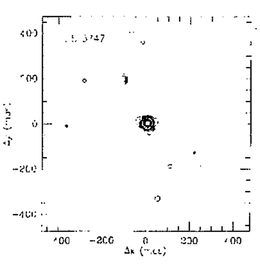

Left: The Caltech group has found five faint pulsars in globular cluster M15 using a concurrent supercomputer to perform the intensive calculations of Fourier transforms in order to distinguish their frequency signatures. C is a binary pulsar somewhat farther out than the others. Below: The ragged, fuzzy 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*

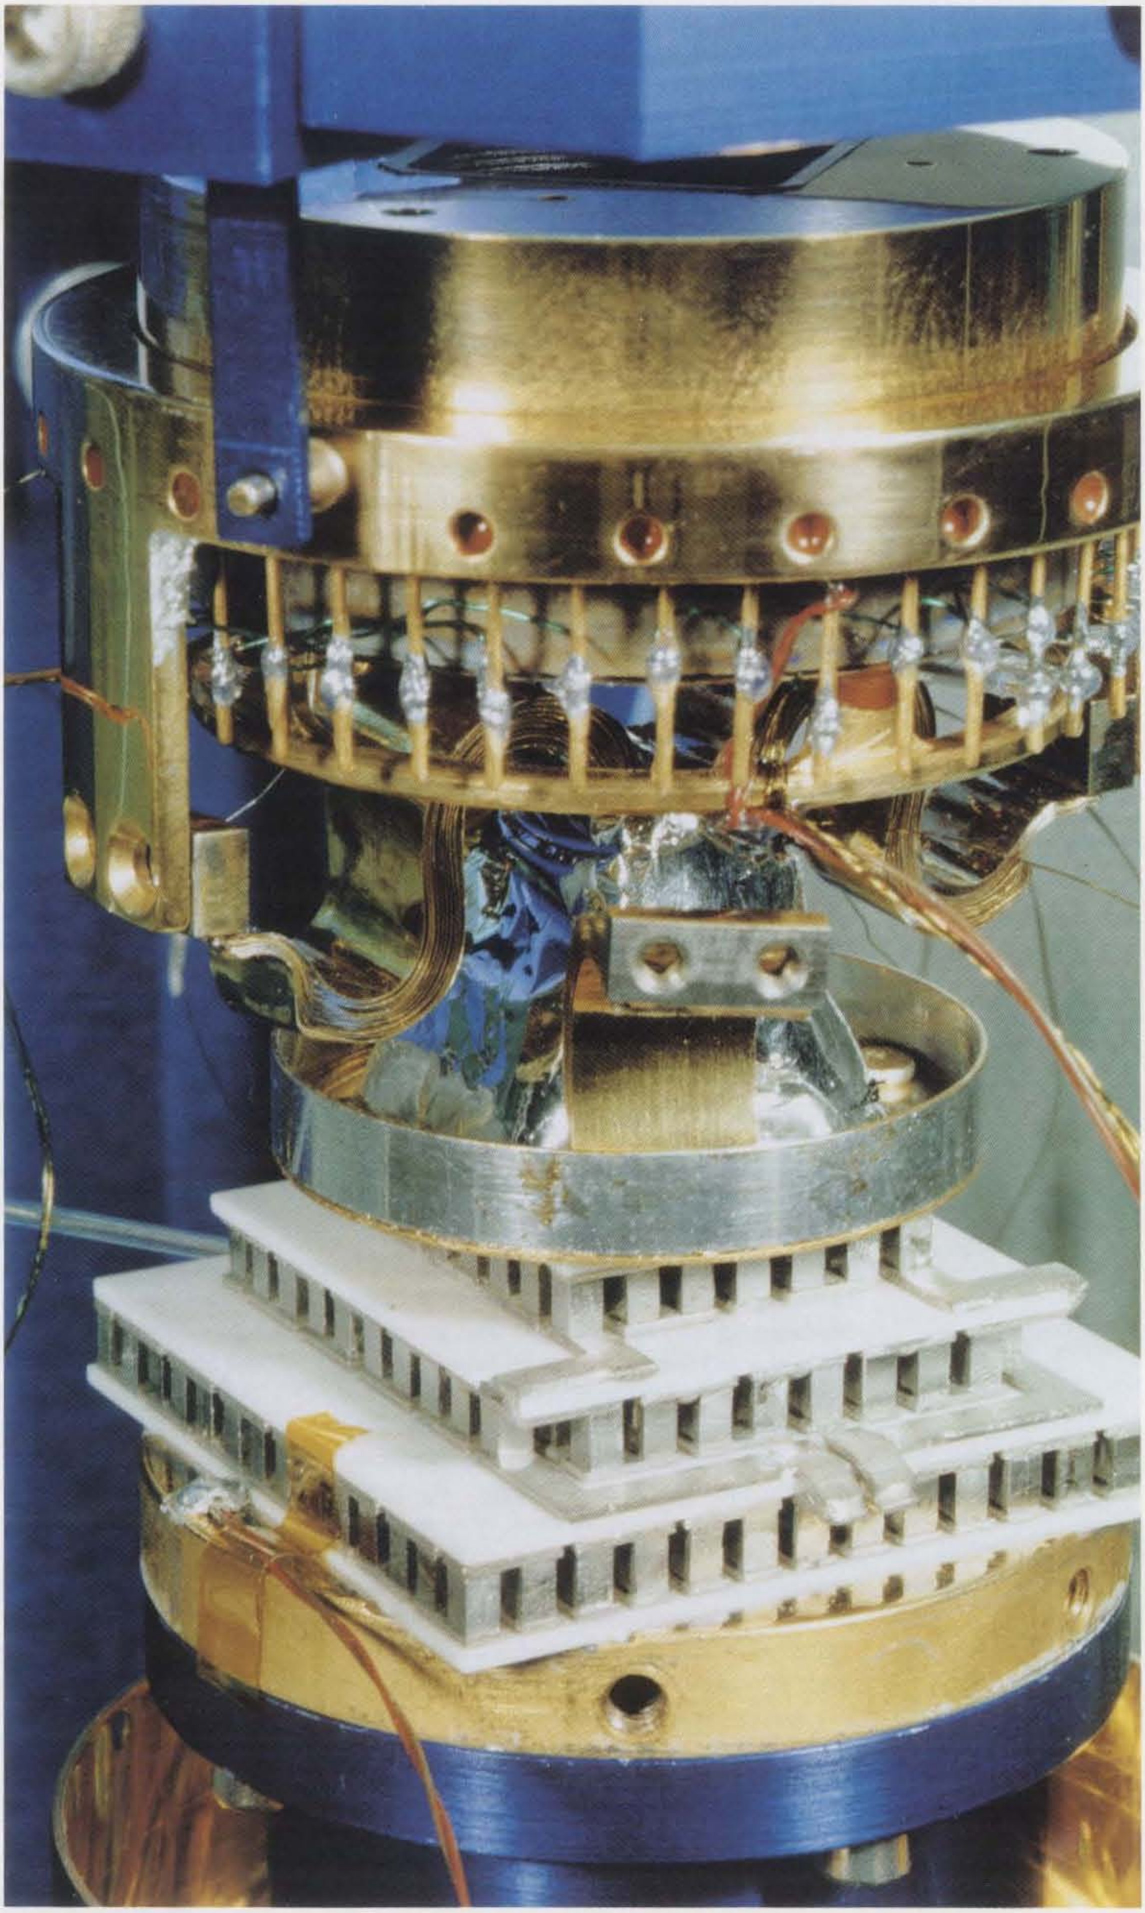

Camera Ready (Telescope Not)

"When the two cultures try to work together, lots of learning has to go on on both sides."

"And I said, 'You're out of your mind. Neither one of us works in that world; we don't want to spend our time up there dealing with that bureaucracy and counting beans and making viewgraph presentations and not being allowed to make marks on a blackboard and all that sort of stuff.'"

Initially, Professor of Planetary Science Jim Westphal was not exactly enthusiastic about the prospect of working with NASA. That was in 1977. On April 25, 1990, however, when the Hubble Space Telescope was finally launched into orbit 381 miles above the earth's distorting atmosphere, it carried the product of a remarkably successful collaboration between Caltech and JPL that won over even such a stubbornly free spirit as Westphal. That product—the Wide-Field/Planetary Camera (WF/PC, pronounced "Wiffpick")—was capable of imaging stars about 10 times fainter than the 200-inch Hale Telescope can resolve, and, in its other, higher-resolution mode, of observing Jupiter with as much detail as could Voyager five days before encounter.

One of the Wide-Field/Planetary Camera's sensors is located behind the pins in this CCD package (see photo on page 36 for scale). The white pyramid at the bottom is a thermoelectric cooler, which keeps the CCD at −100°C.

But heaped onto all the project's other delays came the discouraging discovery in June that a "spherical aberration" blurs the focus of the Hubble's mirrors, which will make the telescope's spatial resolution no better than that of ground-based telescopes—about 1 arc second. (If the entire error is in the primary mirror, it corresponds to a mirror curvature that is too shallow by about two microns from center to edge.) The show will go on, although drastically cut back, for the instruments operating in the ultraviolet (which is important because UV astronomy can't be done from the ground at all), but Wiffpick, which does its significant work in the visible part of the spectrum, was rendered virtually useless. It is, however, acting as optometrist to the stars by diagnosing the flaw.

Meanwhile, back at JPL, Wiffpick II is already being built, originally intended to relieve the first instrument after three years. With considerably more urgency and some compensating changes to the shape of the optics that will cancel out the mirror's aberration, Wiffpick II may be able to fly to the rescue. The whole show may indeed go on, but three years late; NASA officials call it "deferred science."

In the 13 years he spent on a project he didn't want to be involved with in the first place, Westphal got used to deferrals. It was CCDs (charge-coupled devices), a new solid-state detector, which JPL had and Westphal didn't, that dragged Westphal kicking and screaming into collaboration. The solid-state revolution had not yet quite reached astronomy in 1977. Telescopes and spacecraft were still using photographic film and vacuum tubes, although Westphal had begun experimenting with silicon sensors at Palomar. But in the early seventies, the JPL image group under Fred Landauer was looking for a better sensor for the Jupiter Orbiter Probe (renamed Galileo and launched last fall) and was investigating the CCD, a solid-state device that had been invented at Bell Labs. Intrigued by its potential, JPL began work with Texas Instruments (under NASA funding) to develop it for spacecraft imaging, and by the

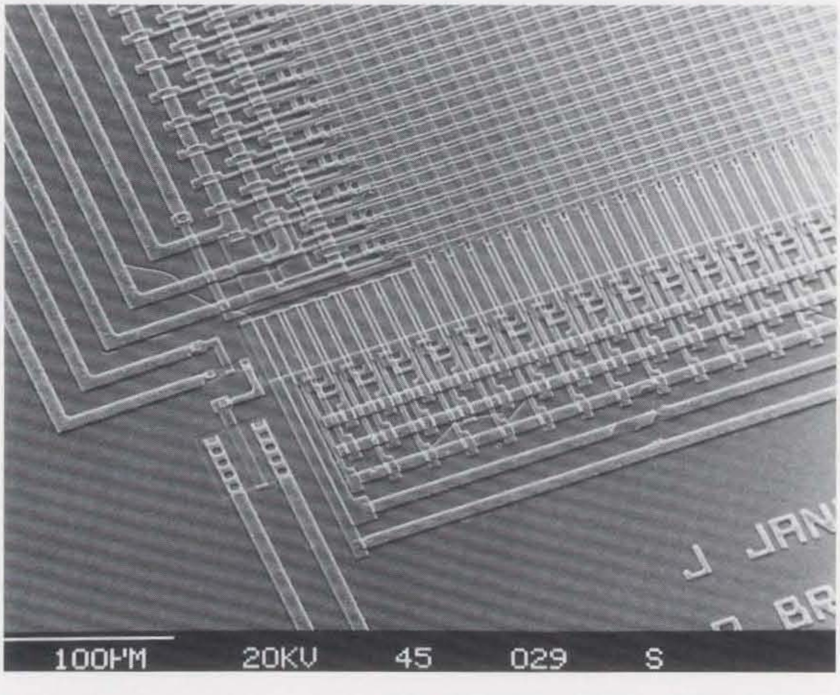

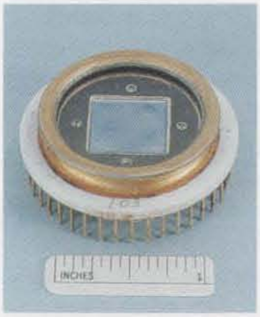

A highly magnified corner of a CCD shows part of the grid of pixels (800 × 800, each 12 microns in size). The charge created by a photon striking a pixel is transferred through the grid and finally dumped into the horizontal register along the bottom. It then goes through a capacitor (the backward L left of center) and out through a control gate (leading off the bottom). Below is the whole CCD—the blue square in the center.

time Westphal entered the scene already had some samples measuring 100 × 160 pixels. Engineer Jim Janesick was a key player in the development of CCDs at JPL. A member of the imaging group and an amateur astronomer as well, he took one of the devices home and built what was probably the first CCD camera for his own little telescope. When he saw what it could do, he "got really, really excited." But Janesick saw a wider field of application than just Galileo. "Once we got some money to start developing the CCDs, the next thing was to convince the astronomical world and the scientific world that the CCD was the thing of the future."

It was indeed, and Westphal knew it as soon as he went up to JPL at the behest of a suspicious Committee on Lunar and Planetary Exploration (of the Space Science Board of the National Academy of Sciences), to figure out what they were doing at JPL with this new sensor. Westphal reported what he had seen to Jim Gunn (then professor of astronomy at Caltech and now at Princeton), who after a few seconds of calculating declared that CCDs "are going to wipe out every other detector astronomers use." And the two desperately wanted "to get our hands on those things and get them on the 200-inch." A short time later Janesick found himself down on campus at Westphal's lab helping Gunn and Westphal build a camera.

A CCD consists of a grid of perpendicular channels on a tiny slab of silicon. When a photon hits one of the individual pixels confined by the channel barriers, it interacts with the silicon to create an electron-hole pair. The charge from the electrons that collect in each pixel's potential well goes up proportionally to the number of photons that hit it, and an image is formed in units of electrons. By manipulating voltages to shift potentials, the charge in each pixel can be transferred out across the barrier phases beneath conductive gates. Each pixel's charge moves sequentially, pixel by pixel, into an amplifier and is eventually reconstituted as video. The beauty of the things is their extraordinary sensitivity to a wavelength range from 1 to 11,000 angstroms (visible light is about 4,000 to 7,000 angstroms). Wiffpick's actual wavelength range extends from 1,150 to 11,000 angstroms—from the ultraviolet to the near infrared. "With a one-electron read noise, [the current CCDs are] the world's most perfect detectors," says Janesick. "It's amazing what the device can do."

Even today Westphal says of the solid-state physics of CCDs that "it's not a science; it's a black art. But it's a wondrous black art. It's still not really something that I have an automatic, clear, warm feeling that we can do what we clearly can do with these devices. We're dealing with two or three electrons at a time—moving them around and doing all kinds of stuff with them. And somehow that just seems like it could hardly be true."

Back in the late seventies almost everyone thought it could hardly be true, with the exception of the Galileo camera team and a handful of astronomers proselytized by Jim Janesick, Jim Westphal, and Jim Gunn (collectively known as the J^3 team). As the Space Telescope struggled into existence, NASA invested heavily in developing a wide-field camera using another kind of detector that wasn't turning out well. ("They never would have made it work anyway," says Westphal.) But during a meeting of a NASA science working group on the Caltech campus, the decision was suddenly made, after a couple of presentations on CCDs, to open up the camera project for competitive bids from principal investigators—scientists who would actually use it. Westphal just happened to be an innocent bystander at this event ("I didn't want anything to do with the Space Telescope; it wasn't my kind of thing; I didn't want to be involved in any way"), but he, Gunn, and six colleagues ended up submitting the successful proposal for the wide-field camera. Its sensors were CCDs.

Westphal's proposal necessarily included JPL, and not just because they had possession of the CCDs. "Of course we weren't competent to do the design," says Westphal. "Designing things that go in spacecraft is a very special art; special talent is needed and a lot of experience. And we

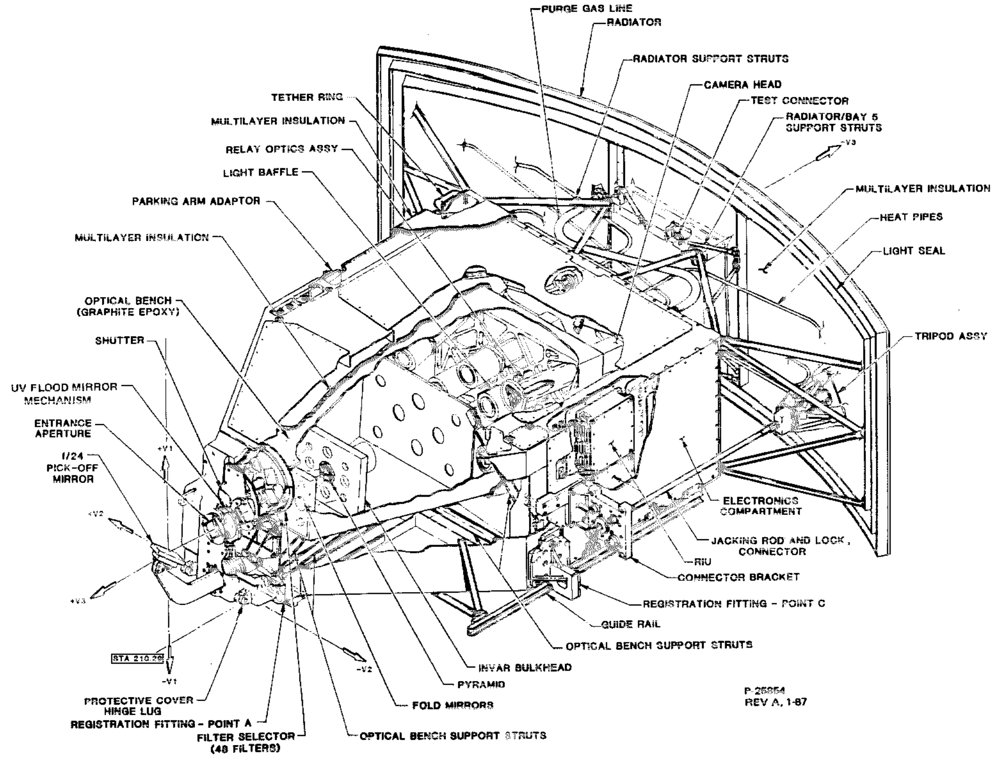

had none of those things." Ed Danielson from JPL helped write the proposal and has remained a part of Westphal's team. And despite Westphal's initial reluctance, the collaboration with the Lab worked—not without problems, but it worked. The project was, according to Westphal, a classic study in Caltech–JPL interaction. "It illustrates not only all the wonders that can be done this way, but also the pain and suffering it takes to make it happen. When the two cultures try to work together, lots of learning has to go on on both sides."

"The science team defined the science objectives; then we converted those into engineering terms."

"The science team defined the science objectives; then we converted those into engineering terms," says Dave Swenson, now program engineer in JPL's Office of Space Science Instruments. Swenson, who began working on the camera part time in 1977 and full time in June 1978, was at various times over the course of the project instrument manager, system engineer, deputy project manager, and gofer, he says. "Defining the instrument—how it fits into the telescope—was a long process. What we originally proposed in many ways is not the instrument we ended up building—mostly for engineering reasons."

In terms of its science objectives the Wide-Field/Planetary Camera was built pretty much as proposed. It's the size of a telephone booth, weighs about 600 pounds, and consists of two cameras in one, each with a different focal length, sort of like a camera with interchangeable lenses—a wide angle and a telephoto. The incoming light beam is "folded" from the telescope's 2.4-meter primary mirror to the secondary mirror and then to a pickoff mirror, which deflects it into Wiffpick's aperture. After passing through one or a combination of several of 48 filters, perched on a spindle like a stack of records (for polarization and spectroscopy and for picking out particular wavelengths), the beam is focused on a pyramid "light switch."

In the f/12.9 wide-field mode the light reflects off the pyramid's four faces through the outer four holes (see diagram), through a set of optics that sets the focal length, and then onto four camera heads, each containing an 800-×-800-pixel CCD; in the wide-field camera each of these pixels is 0.1 arc seconds across. Although the image is optically split apart into four (with an overlap of a few pixels), it can be put back together in a computer to make a mosaic containing the information of all 2,560,000 pixels. The field of this camera is not really wide; it's limited by the size of the CCDs (each CCD covers a quarter of a 2.6-arc-minute square), but it's wide in comparison to the field of its companion camera and can cover a substantial piece of sky.

Rotating the pyramid 45° reflects the light beam into the inner four holes to the planetary camera's four heads, which with a focal length of f/30 can cover only one-fifth as much of the sky. The pixels of the CCDs in these camera heads are, however, 0.043 arc seconds across, providing 2½ times greater resolution.

Putting all this into engineering terms and integrating it with the telescope itself was not a piece of cake. Take the filters, for example. "We were really having trouble with the filter mechanism, figuring out how to do it," says

Wiffpick's eight camera heads assembled and in place (right). The photo below shows one of the camera heads before assembly. At top right is the heat pipe saddle, which removes the heat from the thermoelectric cooler; next to it are the electronics. The thermoelectric cooler and the CCD assembly (wires attached) are in the center.

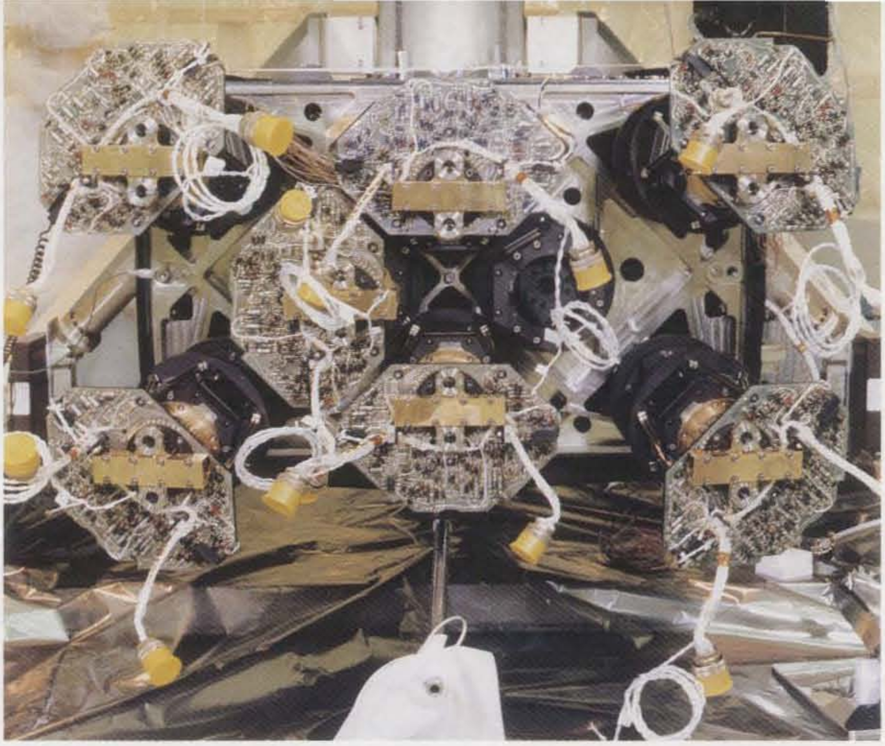

Swenson. He describes one design approach called the "Wurlitzer" for its jukebox-like arm that snatched up one filter from the stack and dropped it into place. "But this thing had 121 springs in it, and if any one of those springs had busted, it would have jammed the whole thing," says Swenson. Jim Gunn, an astronomer, not an engineer, eventually came up with the concept for the final mechanism, which applies magnetic fields to the filter wheels to move them.

Heat was also a problem, and the external radiator, which is about the size of a door and forms part of the telescope's outer skin, was added after the original design. The CCDs, which Swenson describes as "the bread and butter of the whole thing," must be cooled to $-100°$ C with a thermoelectric cooler, but the heat from the power to accomplish this must be removed. Heat pipes filled with ammonia run from the thermoelectric cooler to the external radiator, where the temperature varies from $-20°$ C to $-60°$ C, depending on where it's pointing. The ammonia at the warmer end of the heat pipe vaporizes and collects in the cooler end, where it returns to the liquid phase and is then wicked back up to the radiator. No moving parts—perfect for flight.

Solving such problems put Wiffpick behind schedule (the eventual delays caused by the telescope itself and by the Challenger disaster could not, of course, be foreseen) and threatened cost overruns. Bob Lockhart joined the project in July 1980, when Wiffpick got "projectized"—NASA jargon for focusing more attention on it. Lockhart, who is now project manager of the

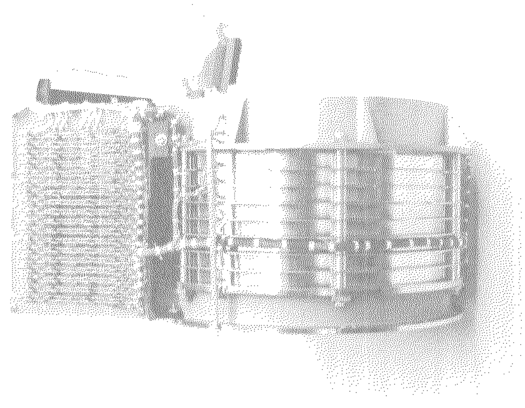

In SOFA, the Selectable Optic Filter Assembly, the filters are arranged on 12 wheels. Of the five slots on each wheel, four are filters and one is clear.

Visible Infrared Mapping Spectrometer for the Mars and CRAF/Cassini missions, took over responsibility for development of the electronics as well as the CCDs and camera heads. The design was basically complete when Lockhart signed on, but the CCDs were experiencing a major problem with noise. The signal-to-noise ratio of a telescope trying to peer to the edge of the universe is obviously rather important; dim objects at such vast distances can easily be drowned out by a few electrons. Each noise electron was worth half a billion light years.

The noise performance requirement for Wiffpick's camera heads was 15 electrons, but only 25 or 30 had been achieved by 1980, after the system had already been put together. Lockhart and his engineers managed to reduce the noise to 12 electrons (each CCD well holds something like 30,000 electrons). Getting rid of noise is like peeling an onion; about a dozen noise sources (peels) had to be eliminated before the 12-electron level was achieved. This took months. As does Westphal, Lockhart invokes black magic when describing CCDs, but he also mentions hands-on engineering. "We could figure what theory tells you minimum noise ought to be, but when you really implement it, moving the wires a few centimeters can make all the difference in real performance of the system."

"This was pretty much an engineering thing. It was the instrument builders at Caltech working with the instrument builders at JPL."

Since there's no way of knowing how to move wires around until the whole system is built, an engineering model for working out all the bugs is necessary for all flight projects. Wiffpick's engineering model was scrubbed to save time and money. But in actual fact, according to Lockhart, it was not. When you start eliminating models, he explains, you are really eliminating them from the last model backwards: first you eliminate the spare, then the flight model, then the prototype. And what you have left is the engineering model. "They call it a 'protoflight' unit," he says. "What that means is that you just build one and you fly that one."

During the protoflight unit's final assembly the engineers discovered that the blue-light sensitivity of CCDs diminishes with time and had even disappeared on a few of the sensors during testing. Lockhart and his team (which included Janesick) traced the problem to trapping sites on the back of the CCD, where incoming blue photons get caught. The solution (which was awarded a patent) was to add a light pipe to the side of the instrument to expose the CCDs to sunlight (ultraviolet). Flooding the CCD with UV charges the back side of the CCDs, which repels signal electrons to the front side where they can be collected. This problem was later solved more efficiently starting from scratch on Wiffpick II: bias gates on the backs of the CCDs and a chemical converter enhance the instrument's response to blue light without the necessity of recharging the surface.

Design of Wiffpick I took about three years and actual fabrication another two (it was finished in September 1983), at a total cost of $65 million. During that time Westphal and his staff spent much of their time at JPL. There was a core group of about 40 to 50 JPLers who worked on it most of the way through, and at the peak probably 90. But what they really could have used, according to Swenson, was "a two-inch gnome with a soldering iron to get in and put it together."

"The working relationship [with Caltech] was excellent," says Swenson, who still remembers Westphal's phone number. "Both Gunn and Westphal were good scientists, but they were good engineers also—very good engineers, in fact." Westphal, too, credits the working relationship to a congruence of interests: "This was pretty much an engineering thing. It was the instrument builders at Caltech working with the instrument builders at JPL. Because Gunn and I are hardware people and build things with our hands, we got on with those guys like crazy."

"There were times, though," adds Westphal, "when we thought we understood it a lot better than the folks at JPL did, and there were various clashes as time went on. But we found our common ground, and things were made to work without, I think, any huge pain and suffering. There were days when you wished you had never

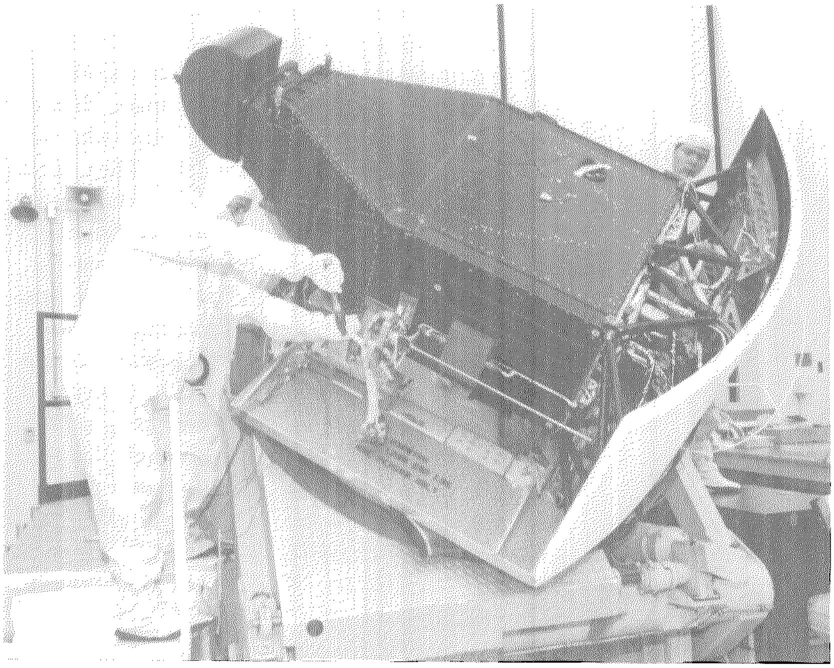

Technicians put the final touches on the Wide-Field/Planetary Camera.

"And I'm sure there were many days when the people up there wondered if there were any way in the world that they could send all of the scientists to Chile and leave us there."

started this thing, or you wished you could just kinda walk in and take it away from all these people that wanted to do all this dumb stuff when you wanted to do something different. And I'm sure there were many days when the people up there wondered if there were any way in the world that they could send all of the scientists to Chile and leave us there."

When the shuttle carrying the Hubble Telescope into orbit finally blasted off in April, there must have been an enormous sigh of relief as the scientists and engineers assumed they could now go their separate ways. But now it's back to the drawing boards after all. The relay optics (between the pyramid light switch and the camera heads) of Wiffpick II have already been built but new ones can be made with a different "prescription." That's the easy part. "To the extent that we can determine what's wrong, we can cancel it out exactly," says Westphal.

Determining what's wrong—which mirror (if not both) and the precise nature of the deformation—is the hard part. The best solution will be to find the error in the paperwork documenting the mirror's manufacture, but while that investigation is going on the Wiffpick science team is using the camera to try to diagnose the problem—taking pictures in various positions of focus and then doing computer simulation of the images. At JPL work on Wiffpick II is accelerating. "They're looking to find people for a second shift to speed things up," according to Westphal. "But they need the right people." JPL engineers are optimistic about meeting their scheduled delivery date of late 1992, even with redesigned and rebuilt optics.

At any rate, it will be a while longer before astronomers can start looking for the beginning of time. The 31 hours of observing time that Westphal earned for his 13 years of devotion to the camera will have to wait too. But at least Westphal and Gunn were successful in their ulterior motive of getting their hands on some CCDs; this produced a revolution in ground-based astronomy while the Space Telescope sat waiting for launch. Some 70 telescopes around the country, including the Hale 200-inch at Palomar, will not be put out of business by Hubble after all because they have received a piece of its technology (and they do still have some other advantages over the Space Telescope). "I've even heard people go to such an extreme as to say that if the Space Telescope never worked it was still a success because it brought CCDs to astronomy," Westphal said—words that may have been more prophetic than he intended.

Janesick at JPL is still working on CCDs and doesn't buy the "black art" stuff. Even after 17 years he claims, "About every two months we get a major discovery." He recently solved the long-standing problem of "blooming" or smearing of a bright image. He's also working with some Caltech biologists to furnish microscopes with CCDs to track very fast fluorescent images. And science is not the only beneficiary of CCD research: CCDs, a suspect curiosity 14 years ago, are now a part of all currently manufactured television sets and video cameras. And soon they'll be standard in ordinary 35-mm cameras. □ —*JD*

Steady As She Goes

Earthbound telescopes are bolted to mountaintops, but the orbiting ones slated for the next century will have to ride lightweight, flexible trusses.

Fanson and the Lab's latest active structure. The 11-inch by 16-inch trusswork weighs 50 pounds. The test mass at the arm's end weighs 52 pounds. The whole structure is built on a 3500-pound base to eliminate outside vibrations.

A telescope's optics must maintain position to within a small fraction of a wavelength of the light being focused, in order to produce a sharp image. Interferometers, telescopes that combine light gathered by two mirrors a fixed distance apart and measure the resulting interference patterns, are particularly sensitive to changes in the baseline distance between the two mirrors. Earthbound instruments are bolted to mountaintops, but the orbiting ones slated for the next century will have to ride lightweight, flexible trusses. Whenever on-board equipment with moving parts kicks in or cuts off, the vibration shakes the truss, setting the instruments dancing like the dangling butterflies on a baby's cribside mobile. With no air resistance to dampen the motion, it persists much longer than it would on the ground. Some items—such as the "reaction wheel," a free-spinning flywheel from which torque can be drawn to pivot the spacecraft—run constantly, giving the structure a persistent throb. How can something so inherently shaky be made rigid?

Thomas Caughey (PhD '54), professor of applied mechanics and mechanical engineering, got interested in this question in the late 1970s, as did other people. The received wisdom was that vibrations should be sensed as velocity changes and damped by attitude-control thrusters—the small jets that keep the spacecraft oriented. Unfortunately, thrusters are a prime source of bad vibes. So Caughey and then-grad student Chuen Goh (PhD '83) studied ways to use displacement information to generate damping forces within the structural members instead. Caughey and Goh showed that displacement control could dampen low-frequency flexing without starting higher-frequency vibrations that could make the structure unstable.

But how to generate the internal forces? Piezoelectric substances expand and contract very predictably when a voltage is applied to them. Robert Forward, an engineer at Hughes Research Laboratory and a noted science-fiction writer, first proposed using them in structures. (Piezoelectric quartz crystals' precise vibrations, like electronic tuning forks, are highly accurate frequency controllers essential to radios and timepieces. Other materials are almost as ubiquitous.) Piezoelectrics have no moving parts that might set up additional vibrations, and they could run off a spacecraft's solar panels.

When the time came for James Fanson (MS '82, PhD '87) to pick a thesis topic, Caughey sent him up to JPL to talk to Jay-Chung Chen (MS '64, ENG '67, PhD '72) about various control issues. Chen had heard Forward speak, and was keen to try piezoelectrics. Fanson and Caughey concurred, and so JPL's first "active structure" to use internal forces was built with money from the director's discretionary fund. The idea was apparently one whose time had come—a handful of researchers elsewhere were taking the plunge, too. Soon afterward, NASA organized a Control-Structure Interaction (CSI) program to deal with the vibration problem and other issues of spacecraft control. This program, which encompasses several NASA centers, became the eventual home for the half-dozen or so JPLers now working on active structures.

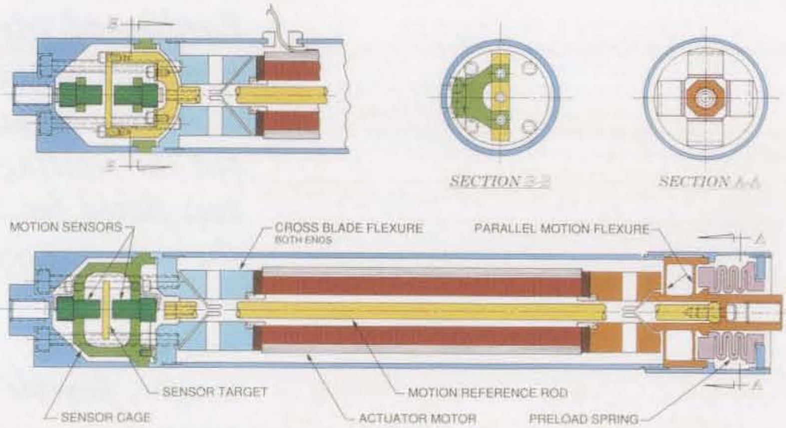

The piston's innards (above). Like a chainsaw-cut tree that's ready to topple, the "cross blade flexure" (upper right) is almost sliced through, tapering down to two thin, flexible regions—the "blades"—at right angles to each other. The blades absorb any sideways load on the piston. The "preload spring" (lower right) is really four springs in one. The machining was done in the Lab's own high-tech machine shop.

The group is on its fourth structure, a nine-foot-tall trusswork tower topped by two four-foot arms at right angles. Built of stock aluminum tubes attached to 1½-inch-diameter ball joints, it looks like it came straight from Tinkertoy heaven. This elaborate web includes eight piezoelectric pistons the size of spring-loaded toilet-paper rollers.

"We're trying to improve structural performance using humans as a model. Our floppy, fleshy bodies can do very precise things because of continuous feedback from our eyes, middle ears, and other sensors. We're essentially trying to put nerves and muscles onto steel skeletons," says Fanson.

The piston's muscle is a stack of lead zirconate titanate (PZT) rings, each about half the size of a Lifesaver and a millimeter thick. PZT is a standard high-performance piezoelectric ceramic—it's the buzzer in smoke detectors and digital watches, among other things. The piston has a total stroke, or expansion range, of 0.003 inches (0.076 millimeters), and its overall length can be controlled to an accuracy of a few nanometers (billionths of a meter)—about ten atomic widths. A rod from one end of the piston runs through the hole in the stack of Lifesavers and connects to a sensor that measures the motion between one end of the piston and the other, and a force sensor on the piston's fixed end tells how much load the ceramic is taking.

The piston itself is a fancy bit of machining. The parts must move freely without touching each other, because on this atomic scale of motion, the slightest bit of friction, the tiniest dead spot, the least tendency of contacts to stick and then suddenly break free—"stiction"—would be ruinous; lubrication is out of the question in the hard vacuum of space. Furthermore, ceramics are brittle, cracking when flexed, so the piston can't take forces in any direction except along its length. So each piston end has an ingenious fitting, sculpted from a titanium block, that absorbs any sideways load and yet allows axial motion. Inside the piston, soft aluminum "crush wafers" nestle the ceramic, distributing the remaining force evenly around the ring, and a four-spring unit, machined from a single block of steel, puts enough pressure on the assembly to keep everything together.

While the muscles are piezoelectric pistons, the nerves are wires connecting a central controller to sensors scattered throughout the structure. Just knowing what's going on at the piston isn't sufficient—you have to know what the whole structure is doing, just as your brain has to know what your legs are up to as well as your arms when you catch a ball. The piston might be near a boom's base for efficiency's sake, because a small motion there translates to a larger movement at the free end, but it's the free end that you want to keep stable.

Associate Professor of Electrical Engineering John Doyle's group has been working with the CSI group to design the algorithms that will constitute the controller. The basic issue is: how do you close the feedback loop between a piston and a sensor some distance away when there's a lot of flex in the framework between them? If the feedback starts going the wrong way, even

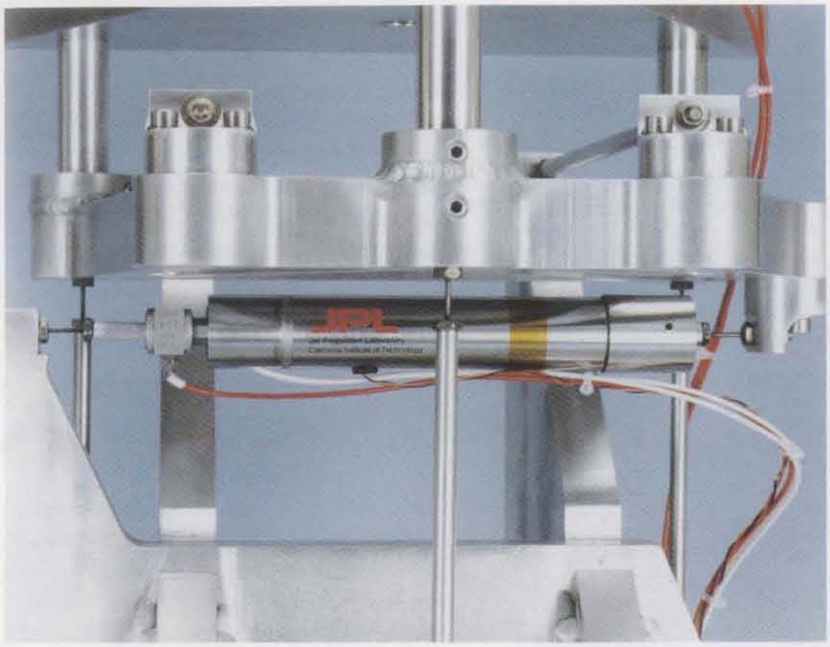

"We're essentially trying to put nerves and muscles onto steel skeletons."

Above: The isolation plate and its active member.
Below: Even the ball joints get wired with position sensors.

by just a little bit, it can rapidly amplify the vibrations, and the truss can shake itself apart. But any controller is based on some mathematical representation, or model, of the system—masses, lengths, stiffnesses, and so forth—that will necessarily be incomplete. You need to develop a "robust" controller, unfazed by small behavioral differences between the model and the real thing. The trick, according to Doyle, is to put these uncertainties into the model as explicitly bounded differences. For example, if the model pictures a long, thin boom as a one-dimensional line, then some part of the system has to know that the boom *does* have a thickness, and under what circumstances—high-frequency vibrations, in this case—that thickness matters. "Engineers have developed intuition about the uncertainties in their models," Doyle says. "But as something becomes more complex, the effect of a particular action becomes less obvious, as do the consequences of modeling uncertainty. We're trying to develop a set of mathematical tools that will help engineers expand their intuition. The tools apply to any feedback-based control system, from flying an airplane to chemical processing."

The controller must be able to handle a wide frequency range. Every structure has a set of frequencies at which it vibrates naturally. Overtones of the basic resonances proliferate at higher frequencies. And as structures get larger, they begin to resonate at lower and lower frequencies. When the resonant frequencies creep down into the region where the attitude-control thrusters operate, the problem becomes very serious indeed. Mariner 10, launched to Venus and Mercury in 1973, was nearly lost when its solar panels and low-gain-antenna boom began to flex in resonance with the thrusters, according to William Layman, JPL's CSI task manager. The controller interpreted the motion as an attitude change and kept firing the thrusters to compensate, putting the spacecraft into a stable oscillation. Half of the attitude-control propellant was blown before Mission Control could find and fix the problem.

The CSI program is using the proposed Orbiting Stellar Interferometer (OSI) as a test case to identify the problems that need to be solved, and then to develop and apply the relevant technology. The tower is part of that study, incorporating a laser interferometer and three layers of active control. The first is the active members in the truss, of course. The second is a mounting to which a vibrationally noisy component is bolted, and which is attached to the truss with an active isolation system. And the third is an "optical delay line"—an active mounting for the interferometer optics that nudges them a fraction of a wavelength to keep the interferometers precisely separated—based on a design that OSI's principal investigator, JPL's Michael Shao, developed for ground-based use. (Keeping the pieces of a segmented mirror, such as the Keck telescope's, in position falls into the same category. The Keck has mechanical pistons for the job.) The group will start experiments with the active components soon. Right now they're shaking the truss over a wide range of frequencies to find its natural resonances.□—*DS*

Random Walk

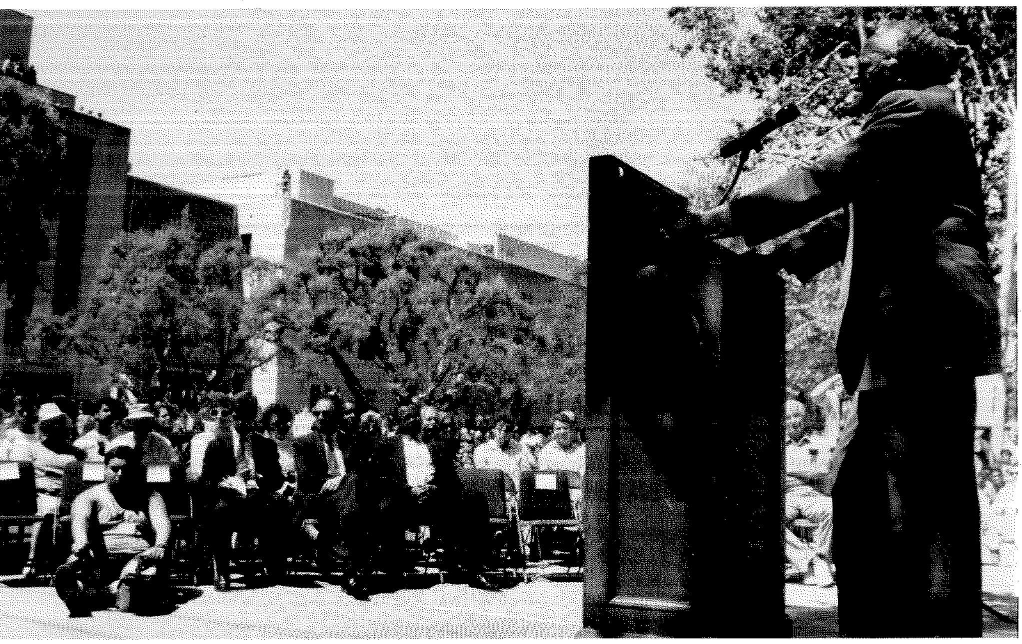

Archbishop Desmond Tutu of South Africa addressed the Caltech community on May 22. His speech, given on the Olive Walk, drew a capacity crowd, some of whom found the view best on nearby rooftops.

Under(graduate) Achiever—and Proud of It

Mike Chou was one of 20 college and university juniors nationwide selected by *Time* magazine for its 1990 Achievement Awards. Chou, who is pursuing an option in physics, won $3,000, a free trip to New York City, and a profile in an upcoming issue. He was cited for his 1989 SURF (Summer Undergraduate Research Fellowship) project on particle emissions during solar flares, sponsored by Senior Research Associate Richard Mewaldt.

Honors and Awards

Three faculty members have been elected to the American Philosophical Society, the nation's oldest learned society. They are Don Anderson, McMillan Professor of Geophysics; Edward Lewis, Morgan Professor of Biology, Emeritus, who was also co-recipient of Brandeis University's 1990 Rosenstiel Award, bestowed annually to outstanding life scientists; and Rudolph Marcus, Noyes Professor of Chemistry.

Michael Aschbacher, professor of mathematics, John Bercaw, Shell Distinguished Professor and Professor of Chemistry, and Barclay Kamb, Rawn Professor of Geology and Geophysics, have been elected to the National Academy of Sciences, one of the highest honors that can be accorded an American scientist or engineer.

Jay Bailey, Chevron Professor of Chemical Engineering, is the 1990 recipient of the Marvin J. Johnson Award, presented annually by the Biochemical Technology Division of the American Chemical Society.

Seymour Benzer, Boswell Professor of Neuroscience, has been given the 1990 W. H. Helmerich III Award for Outstanding Achievement in Retina Research.

Four faculty members have been named Presidential Young Investigators for 1990 by the National Science Foundation. They are Assistant Professor of Biology William Dunphy, Assistant Professor of Mechanical Engineering David Goodwin, Assistant Professor of Chemical Engineering Julia Kornfield, and Assistant Professor of Physics Jonas Zmuidzinas.

Caltech President Thomas E. Everhart, Koepfli Professor of the Humanities Daniel Kevles, and Trustees Stephen Bechtel, Jr., and Gordon Moore have been elected Fellows of the American Academy of Arts and Sciences, one of the oldest honor societies in North America. Everhart has also been appointed chairman of the newly created Secretary of Energy Advisory Board.

Leroy Hood, Bowles Professor of Biology, has been honored by the American College of Physicians for his "outstanding work in science as related to medicine."

Assistant Professor of Biology and Computation and Neural Systems Gilles Laurent has been chosen a Searle Scholar for 1990, and with collaborators in England, Germany, and Japan, has been awarded one of the 29 "Human Frontiers of Science Program" grants given worldwide by NATO this year.

Institute Professor of Chemistry, Emeritus, John Roberts is a co-recipient of The Welch Foundation's 1990 Welch Award in Chemistry.

P. P. Vaidyanathan, associate professor of electrical engineering, was given the S. K. Mitra Memorial Award by the Institution of Electronics and Telecommunications Engineers.

Associate Professor of Electrical Engineering and Computer Science Yaser Abu-Mostafa and Assistant Professor of Applied Physics Kerry Vahala have been chosen recipients of the newly created Richard P. Feynman–Hughes Fellowship.

Kai Zinn, assistant professor of biology, has been selected a 1990 Pew Scholar in the Biomedical Sciences by the Pew Charitable Trusts.

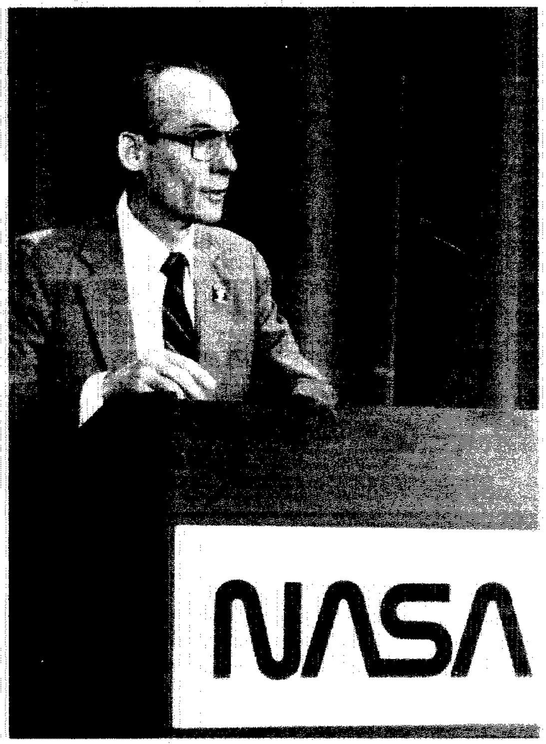

Edward Stone Named JPL Director

Edward C. Stone, Jr., a man long identified with the most spectacular of JPL's successes, will become director of the Jet Propulsion Laboratory in December, succeeding Lew Allen. Since 1972 Stone has been project scientist of the Voyager mission to the outer planets, coordinating analysis of the scientific data and communicating to the inhabitants of one of the inner planets (and in a special way to the readers of *E&S*) what the spacecraft saw.

Stone isn't planning to abandon Voyager, but his involvement will continue at a reduced level. He will also continue as chairman of the board of directors of the California Association for Research in Astronomy (CARA), which oversees all facets of the development of the W. M. Keck Observatory and Telescope in Hawaii. The 10-meter telescope is scheduled for first light in the fall. Stone will remain a vice president of Caltech.

"This is an exciting time for JPL," said Stone, "with all the missions on their way, the missions under development, and the plans for new ones. JPL continues to have a primary role in NASA's space science program. I'm looking forward to providing scientific leadership."

Stone received his AA degree from Burlington (Iowa) Junior College and his SM and PhD (1964) in physics from the University of Chicago. He began his Caltech career as a research fellow in that same year and became assistant professor in 1967, associate professor in 1971, and professor of physics in 1976. From 1983 to 1988 he served as chairman of the Division of Physics, Mathematics and Astronomy, and has been vice president for astronomical facilities since 1988. In addition to his position as project scientist for Voyager, Stone has been a principal investigator on nine NASA spacecraft and a co-investigator on five others. His primary field of interest has been the isotopic and elemental composition of galactic cosmic rays, which arrive from the interstellar medium, a region that Voyager may finally intercept in about 10 years.

As someone who has had one foot in each camp for a couple of decades, Stone is well placed to appreciate collaboration between campus and JPL. "The Caltech connection is one of the unique aspects of JPL as a NASA center," said Stone. "The opportunity to undertake joint efforts has grown under Lew Allen, and I would like to enhance it further."