California Institute
of Technology

Engineering & Science

**Volume LIX, Number 3
1996**

In this issue

*Smelling Yums
and Yucks*

Feynman Redux

*Turbulence
in Flux*

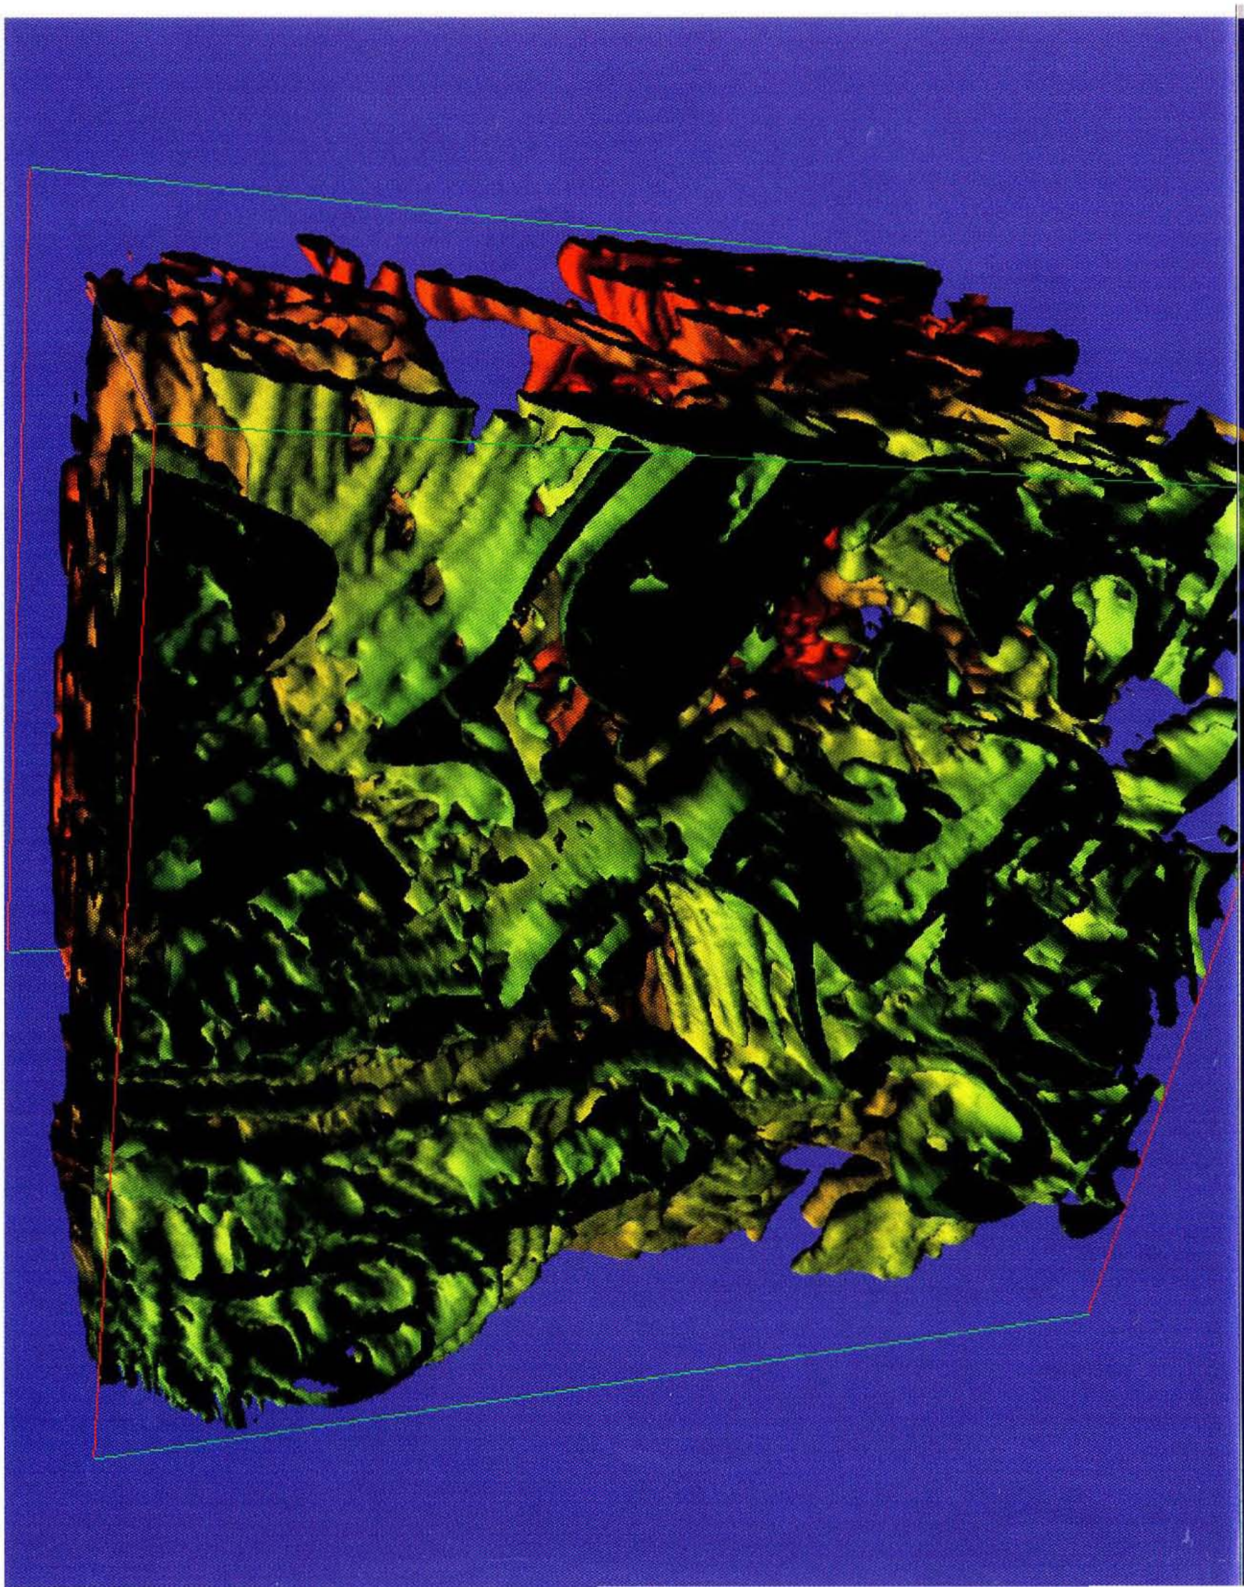

Richard Feynman entertains a group of students with his safecracking tales. This photo, the last of a set of four that first appeared in *E&S* in June 1964, was captioned: "He's sitting there all this time thumbing through a magazine, with this big fat smile on his face. So I let about 5 minutes go by and then I swing the thing open . . . He is *flabber*gasted!" Feynman also delivered a lecture in 1964 that was recently unearthed and made the subject of a book, an excerpt from which begins on page 14.

California Institute
of Technology

Engineering & Science

Volume LIX, Number 3
1996

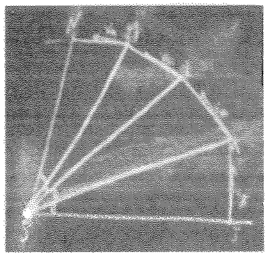

On the cover: This sponge is actually a portrait of the turbulence caused by squirting a dye jet into a tank of standing water. The cube's front face is a 5-x-5-centimeter square, and looking into the page corresponds to looking backwards in time. The sponge's cross section any depth consists of all the points in the square with a given dye concentration at that instant of time. For more on how scientists are getting a better look at the face of turbulence, see the story on page 22.

Departments

Engineering & Science (ISSN 0013-7812) is published quarterly at the California Institute of Technology, 1200 East California Boulevard, Pasadena, California 91125. Annual subscription $10.00 domestic; $20.00 foreign air mail; single copies $3.00. Third class postage paid at Pasadena, California. Published by the California Institute of Technology and the Alumni Association. Telephone: 818-395-3630. Postmaster: Send change of address to Caltech 1-71, Pasadena, CA 91125.

PICTURE CREDITS: Inside front cover — Kent McCaulley; 4, 6, 9, 10, 11 — Erik Severin; 5, 9 — Brett Doleman; 5, 26 — Doug Smith; 6, 10, 11, 21 — Bob Paz; 7, 8 — Mark Lonergan; 12 — Hillary Bhaskaran; 14, 17, 18, 20, 39 — Caltech Archives; 16 — Igor Bitman; 25, 27, 28 — Haris Catrakis; 34 — Paul Dimotakis; 36 — Jean-Paul Revel, Stefan Offermanns, Mel Simon; 41, 42 — James McClanahan; 43 — Herb Shoebridge; 44 — Bill Varie; inside back cover — Tom Bida/Judith Cohen

Edward M. Lambert
President of the Alumni Association
J. Ernest Nunnally
Vice President for Institute Relations
Robert L. O'Rourke
Associate Vice President for Institute Relations

STAFF: *Editor* — Jane Dietrich
Managing Editor — Douglas Smith
Copy Editors — Barbara DiPalma, Michael Farquhar, Danielle Gladding, Julie Hakewill
Business Manager — Debbie Bradbury
Circulation Manager — Susan Lee
Photographer — Robert Paz


CALIFORNIA INSTITUTE OF TECHNOLOGY
1891
CALTECH

The Caltech Electronic Nose Project

Can we teach a computer to "smell" in the same way that we can teach it to "see"?

by Nathan S. Lewis

Although not blessed with the keenest noses in the animal kingdom, humans (in this case, the author's three-year-old son Jeffrey) can smell the difference between yum and yuck almost from birth. Photo and subject courtesy of Dr. Carol Lewis, Jet Propulsion Laboratory.

Of our five senses—sight, smell, taste, hearing, and touch—we understand three well enough to build machines that mimic them. Touch is basically a pressure sensor. There are artificial cochleas—mechanical resonators that transmute sounds into signals that our brain, or a machine, can recognize. And we can build cameras that are essentially electronic eyes. But we know very, very little about the molecular basis of taste and smell, and even less about how to model them. So my lab is trying to build something that will give a value judgment—a number—to a smell, taking design lessons from biology without necessarily mimicking the exact way that a human nose works. We can assign a visual magnitude, a brightness, to a star; can we teach a computer to "smell" in the same way that we can teach it to "see"? This project began as a crazy idea in January of 1993, but there may be something to it.

Smell is a remarkably subtle sense, because most smells are not pure substances, but complex mixtures of different molecules. There are some 700 different chemical vapors in a glass of beer, yet somehow we can take a sniff and say it's beer. The human nose is generalized enough to sense almost all possible molecules, yet discriminating enough to tell the difference between strawberries and raspberries. How can we model that?

The way that most chemists have approached this problem is epitomized by what Arnold Beckman [PhD '28] did when he invented the pH meter. He built a chemical sensor that measures the concentration of one thing (protons in water) very selectively and very sensitively. People have since extended that idea to measure other molecules, such as glucose. In almost every case, the strategy is to design a molecule that has a hole in it—a lock—such that only the right key, i.e., glucose, will fit and generate a signal. (There are, of course, more generalized sensors that measure some physical property of the molecule, but they don't really "recognize" it—they merely tell you that they've detected a molecule with, say, the same charge-to-mass ratio as the molecule you're looking for.) Nature uses the lock-and-key approach very successfully—in enzymes, for example—but it takes evolution millions of years of work to make the molecules fit just right. You can see the daunting task that a chemist would face in trying to build 700 such locks to detect the 700 odor components in a glass of beer. And we'd have to build all 700, because we don't know which components are critical for identifying the smell of beer, and determining whether it smells good or stale. And what would happen when we encountered the 701st molecule in a different odor, like in another brand of beer? We'd have to build another sensor. And we'd have to make each lock specific enough that a very slightly different molecule wouldn't also fit, because even if the other molecule fits poorly we'd still get a signal. Designing such exact locks from scratch is a very, very complex problem at the frontiers of chemistry, and hundreds of groups around the world are working on it.

We abandoned this approach in favor of a pattern-recognition strategy. We decided that the biological olfactory system must employ a set

The electronic nose, high-school science-project style. Two electrodes are taped to a sponge. When the sponge is dry (top), the electrodes are in contact, completing the circuit and lighting up the bulb. As the moistened sponge swells (bottom), the electrodes move apart and break the circuit. The light goes out.

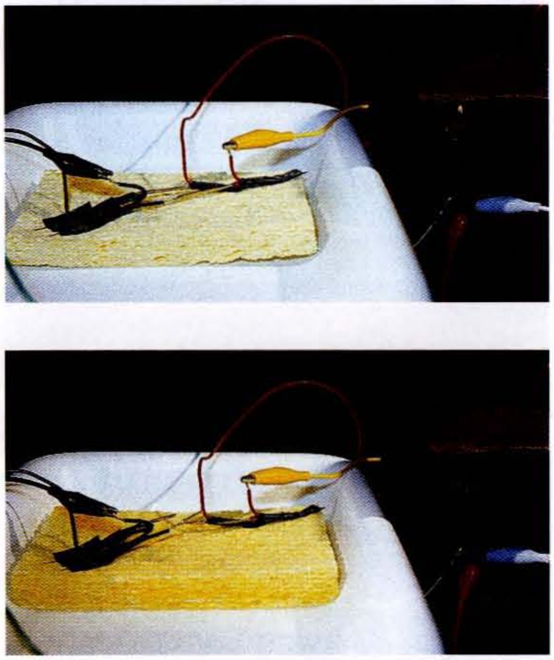

Dogs have no reason to sniff out cocaine in the wild, and yet they can be trained to do so in airports. The dogs must be learning to recognize a pattern, because one certainly couldn't train them to develop a new receptor overnight, or even in a few months.

of generalized sensors that respond to everything, but in different ways to different stimuli. Evolution might have developed specific receptors for fruits and wines, for example, but it's unlikely that dogs would have evolved receptors to smell drugs. Dogs have no reason to sniff out cocaine in the wild, and yet they can be trained to do so in airports. The dogs must be learning to recognize a pattern, because one certainly couldn't train them to develop a new receptor overnight, or even in a few months. So the task facing anyone trying to develop an artificial nose is to develop a generalized sensor whose output patterns will announce the difference between the vapors emitted by a rose and a dead fish. Then we train an electronic circuit to recognize those patterns, in the same way that signals fired to our brain get recognized as yum or yuck.

The sensor in our electronic nose must meet several basic requirements. We want it to give us an electrical signal that we can analyze on a chip. We want the signaling event to be reversible—that is, the sensor should return to its initial state when the sniff goes away, so we can use it over and over again. We want it easy to make. We want it to be stable in all sorts of environments, so we can just leave it sitting out in the air. And we want to be able to make it very small, so that we can put a million of them on a little chip.

Our solution is embarrassingly simple. In fact, I'm proud to say that a well-known physicist who wasn't familiar with this project came into my lab recently, looked at our nose and said, "This is a high-school experiment." And I said, "That's exactly right! That's what makes it so wonderful to study, because it works for anyone anywhere." Our sensor is a sponge made of insulating plastic, much like a bathtub sponge, but containing little conducting particles scattered here and there within it. When we pass a current through it, the electrons have to hop from one conductor to the next, so the sponge has a characteristic, measurable resistance. If we were to moisten the sponge, it would swell, and the conducting particles would move farther apart. It would get harder for the electrons to jump between the conductors, and the resistance would go up. Later, as the sponge dried, it would shrink, and the resistance would go back down. (If you soak a sponge, it won't shrink all the way back to its original size when it dries, but if you just add a few droplets, the swelling can be fully reversible.) The same thing happens with vapors—the sponge "sniffs" an odor by absorbing it and swelling up, causing a measurable resistance change, as you can see on the opposite page.

The linchpin of our design is to use an array of sponges with different chemical affinities. Each individual sponge will swell more (and exhibit a higher resistance) when it soaks up something it likes. For example, hydrophobic plastics don't like water at all. If you expose them to a water-like vapor, such as an alcohol, they'll repel it. The sponge won't swell much, and there's not much signal change. But hydrophobic materials *do* like oil, so an oily vapor—benzene, for example—will swell them a lot. So some of our sponges like oil better than water; some like charged molecules more than uncharged molecules, and so on. There's no lock-and-key design that says that Sponge A will only respond to the molecule "methanol," and Sponge B will only respond to the molecule "benzene." We don't have to worry about the details—instead, the molecule tells *us* what its important properties are by the signals it generates in the various sponges. We don't actually know if we have enough diversity amongst our sponges yet, so the nose will evolve as we pull out sponges that don't work very well, and put in ones that we hope will work better. We're still trying to figure out how best to choose them.

We were originally going to vary the chemical affinities by modifying the conducting particles. My colleague Bob Grubbs, the Atkins Professor of Chemistry, has discovered ways of making electrically conductive plastics that you can paint on anything. [See *E&S*, Summer 1988.] But then-postdoc Mike Freund, who began this whole project (and is now an assistant professor at Lehigh University), realized that we didn't need to alter the conductor. All we really needed to do was to make one paintable conductor, and then use assorted commercial plastics with various properties, available from any supply house, as the insulators. So that, being simpler, is what we do. We dissolve the insulator, add the ingredients needed to make the conductor, and then apply the resulting solution while the reactions that make the conductor are going on. The solvent eventually evaporates, leaving us with our sniffer sponge.

In hindsight, it turns out that we didn't have to go to all the trouble of making conductive plastics. Any electrical conductor will work, as long as we can find a way to disperse it into the sponge. For example, last summer, SURF [Summer Undergraduate Research Fellowship] student Sara Beaber started using little particles of silver. And Pinocchio, our newest, most improved, nose uses carbon-black particles—the same stuff you find in asphalt and pencil lead. Postdoc Mark Lonergan did most of the work on this, aided by grad student Erik Severin, and Bob Grubbs, as usual, had the idea. Carbon black is a very stable compound, unlike the temperamental conducting polymers, and it's really easy to come by. If you break apart a Radio Shack resistor, you'll find little balls of carbon black inside.

And the way we attach wires to our sponge is incredibly inexpensive—we break apart a 10-cent capacitor. Capacitors store electric charge on thin sheets of palladium–silver foil, separated by a good insulator—a sand-like material called mica—so they don't short out. We use a belt sander to grind the top half of the capacitor down until we expose the foils, and then dip it in our

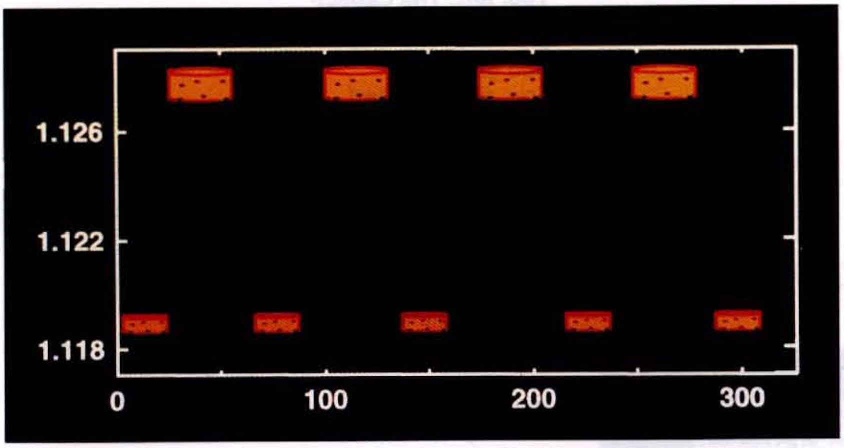

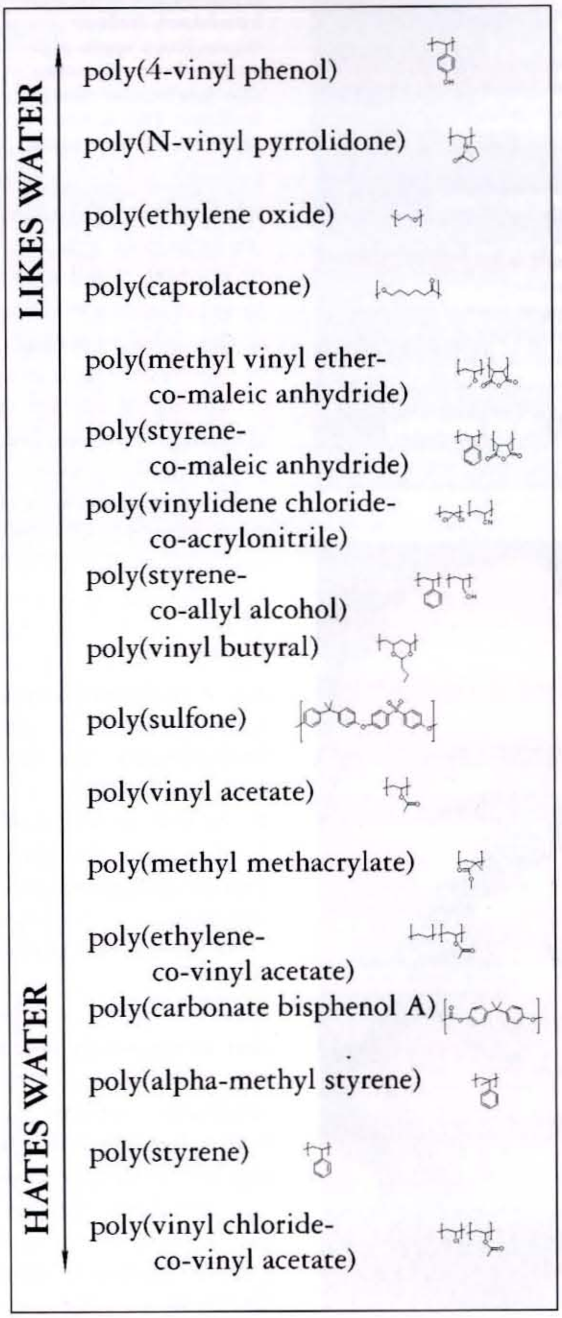

Above: When a sponge [here, poly (ethylene-co-vinyl acetate) containing carbon-black particles] sniffs something it likes (air containing 0.1 percent benzene), it swells and its resistance jumps. When the odor vanishes, the sponge shrinks, so it can be used over and over. The vertical axis is resistance in millions of ohms, and the horizontal axis is time in seconds. Right: The insulating plastics that make up the sponges can be ranked by their properties on many scales, including affinity for water. The chemical structures in brackets are the repeat units that make up the corresponding plastics.

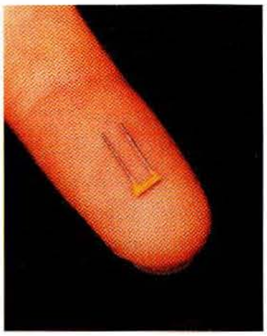

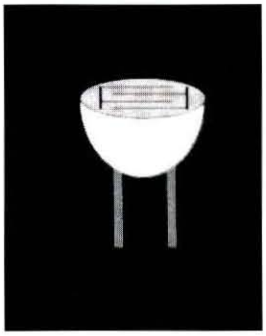

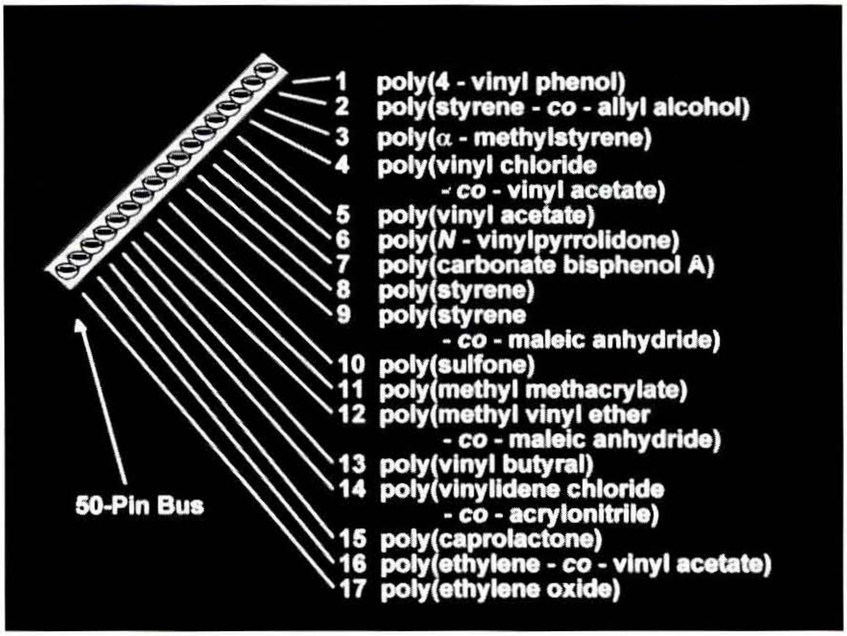

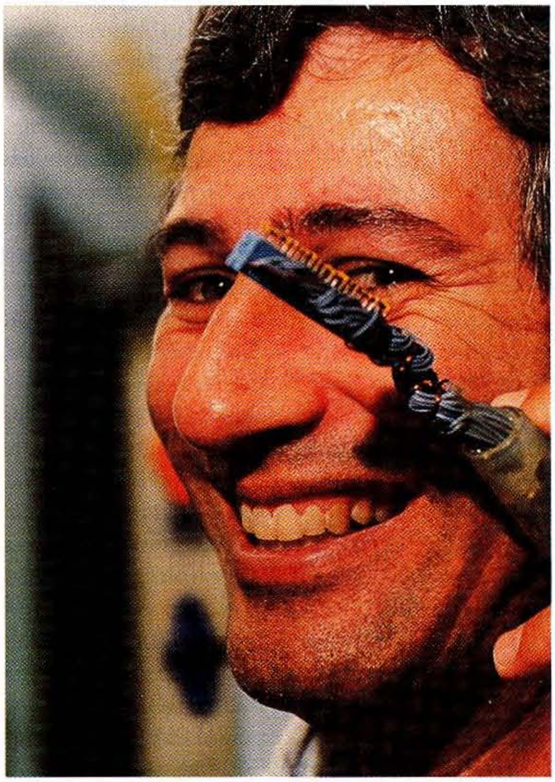

Top, left: The capacitors used in the nose are about the size of rice grains.
Top, right: Sealed in epoxy within each capacitor are two sets of interleaved parallel plates, separated by an insulator. Sanding the top off the capacitor exposes a cross section through the plates. Applying a sponge coating to the exposed surface completes the circuit.
Middle: Solutions of the 17 plastics listed on the previous page were doped with carbon black before capacitors were dipped in them to make this particular nose.
Bottom: The author and two of his noses.

solution, bridging the insulator. Then we plug the capacitors right back into where they came from. A so-called bus chip a few centimeters long can hold a whole array of capacitors, each with a different sponge. The output signals then feed directly into the computer.

The signal's height and shape depend on both the thing being smelled, and the thing that's doing the swelling to sense the smelling. As you saw before, the resistance rises as the sponge swells, plateaus at some value characteristic of the vapor for as long as the vapor remains, and then falls off as the sponge shrinks once the vapor disappears. The swelling and shrinking rates depend on how the sniffer and the sniffed interact. A hydrophobic sponge, for example, will slurp up benzene because it's greasy, and won't let it go easily. But the same sponge won't soak up as much chloroform, and will release it faster. Right now we only look at the maximum signal-height change, but the curve's shape should provide additional, and maybe more valuable, information in the long run.

When we look at the overall pattern of all the signals from all the sponges, we get a fingerprint that—we hope!—will be different for everything that we expose the nose to. (So far, that's been true.) One sponge by itself does not identify a compound—another compound that didn't swell it as well might give a signal that's half as high, but if there were twice as much of that second compound, we might get a very similar signal. But the signals from the entire array provide a pattern that will be diagnostic of a given odor. On the facing page is an example from an array

When we look at the overall pattern of all the signals from all the sponges, we get a fingerprint that —we hope!— will be different for everything that we expose the nose to.

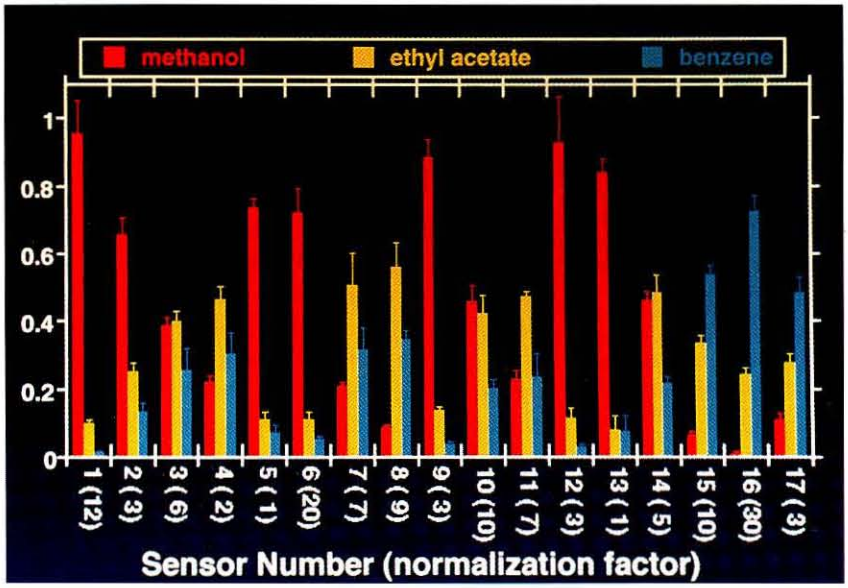

The 17-sensor carbon-black nose gave three different response patterns for three different vapors. The numbers on the vertical axis indicate the relative resistance change in each sensor. Because the response range of each indivdual sensor is different, the values were "normalized" to make them fit on a common scale by dividing them by the number shown in parenthesis below each sensor number.

of 17 different sponges. The yellow bars show the pattern that we get for ethyl acetate, a solvent commonly found in paint thinner. The blue bars are the pattern that we get for benzene, and the red bars are methanol. You don't have to have a trained eye to see that they are different, so we can certainly distinguish them electronically.

But it's hard to quantify *how* different they are. You can't tell me if they're 10 percent different, or 20 percent. How can we teach a machine to discriminate between patterns whose differences we can't easily describe ourselves? How do we know how much leeway we can allow between two patterns and still call them a match, for example? We use a statistical method called principal component analysis (which we did not invent) to analyze the data. The method takes all the signals from the individual sensors and plots them as points in what we call odor space, in which it's easier to see the patterns. Unlike ordinary three-dimensional space, however, we have one dimension per sponge. Therefore, even though it's easier to see the patterns, the analytical process can still get quite elaborate.

Last year we did an experiment where we exposed 17 sensors to nine pure vapors—methanol, ethanol, isopropanol, acetone, ethyl acetate, chloroform, hexane, benzene, and toluene. We gave the nose sniffs of the various vapors, repeated in random order, over a period of five days. We didn't control the temperature of the room, and we didn't control the humidity in the air, so this experiment was essentially a worst-case scenario to see how well we could do. The shapes enclosing the data for each compound would have been

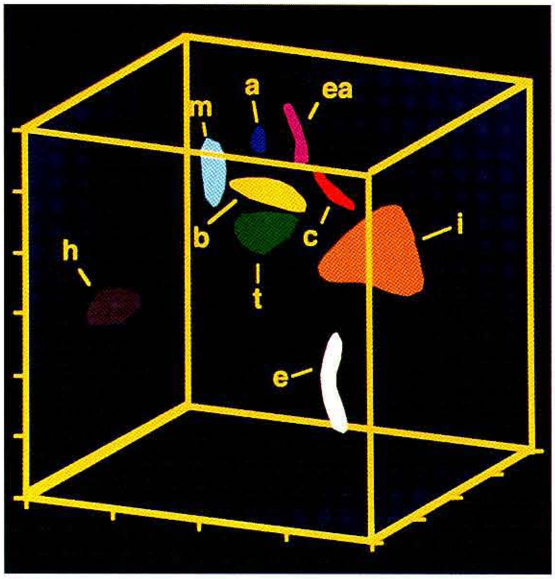

When the nose was exposed to nine different vapors, each one turned up in its own little corner of odor space once the right set of dimensions was plotted. The vapors are labeled as follows: a = acetone, b = benzene, c = chloroform, e = ethanol, ea = ethyl acetate, h = hexane, i = isopropanol, m = methanol, and t = toluene.

smaller in a controlled climate.

I can't plot a 17-dimensional space, so the plot above shows the three dimensions that contain the most differences between those nine patterns. The three alcohols (methanol, ethanol, and isopropanol—methanol has one carbon atom, ethanol two, and isopropanol three) separated very cleanly. Benzene and toluene, which are chemically only very slightly different—toluene has an extra methyl group and so is just a little bit bigger than benzene—are close together, but distinguishable. By contrast, hexane—a molecule about the same size as benzene and toluene, but with a different shape and very different properties—appears quite a distance away. And ethyl acetate and acetone (the solvent in nail-polish remover) are also chemical cousins, but they aren't as closely related to each other as benzene and toluene are, so they show up farther apart than benzene and toluene do. Chloroform, which isn't related to any of these guys, also registers separately.

We can also tell how much of something there is, because the responses grow larger with increasing vapor concentration. All the sponges continue to swell in approximately the same relative way as the odor gets stronger, and we retain the fingerprint.

Each coordinate axis represents some unknown property—it might be how big the molecule is, how it is shaped, how much it likes water, or, usually, some combination of properties. We've already seen how hydrophobicity works, and polarity works much the same way—we can make our sponges hospitable to positive, negative, or neutral charges. We can also discriminate between molecules of different sizes, because the plastics' pores differ in size and shape. Molecules that are too big for the pores don't fit very easily, so the sponges don't swell as much. Molecules that are smaller than the pores do fit, but not very well, and so again the sponges don't swell as much. Discovering what the coordinate axes actually correspond to is a very interesting problem. We're working very hard to try to associate the chemical and physical characteristics of the sniffed molecule with our sniffer data.

Since we have 17 dimensions to choose from, we can select the three that best discriminate between whatever specific compounds we're interested in. If I wanted, for instance, to separate chloroform and toluene, I could plot three other dimensions that would separate chloroform from toluene much better, but wouldn't separate methanol from ethanol as strongly.

A computer, of course, isn't limited to "seeing" things in three dimensions, as we are, but can look at all 17 at once.

A computer, of course, isn't limited to "seeing" things in three dimensions, as we are, but can look at all 17 at once. We had to learn how to analyze such data, so we're collaborating with Rod Goodman, professor of electrical engineering and director of the Center for Neuromorphic Systems Engineering, which is devoted to developing machines that mimic, on some level, the way biological brains—what's known in the trade as "wetware"—work. Rod and grad students Jeff Dickson and Alyssa Apsel are developing a model to handle our data flow based on how our brains might analyze the firing of neurons as we recognize an odor. And last summer, a SURF student of Rod's named Wei Qin set up for us a

Right: A two-dimensional plot of the nose's response to methanol and ethanol mixtures. The line of red squares indicates the response to air–methanol mixtures, and the line of green circles is for air–ethanol mixtures; in each case the deeper the color, the higher the vapor's concentration. The nose was then given whiffs of five mixtures of methanol and ethanol (methanol–ethanol ratios of 11:1, 4:1, 2:1, 1:1, and 1:2) at two different flow rates to give two sets of concentration values. The data fell neatly onto the broken lines. The arrow marked χ_M shows the direction of increasing methanol content, so where a mixture appears on the graph is directly related to its composition. Below: If all the sensors respond linearly to individual vapors—that is, if the response increases in proportion to the vapor's concentration—the array's response to a mixture of vapors will be the sum of the responses to each vapor as if it were by itself. Here, for example, sensor 1 registers a 5 for vapor A and a 1 for vapor B, so a half-and-half mixture of the two registers as 2.5 + 0.5, or 3.

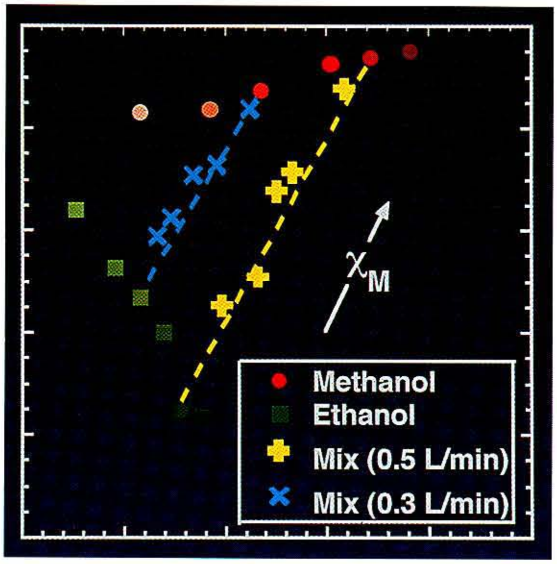

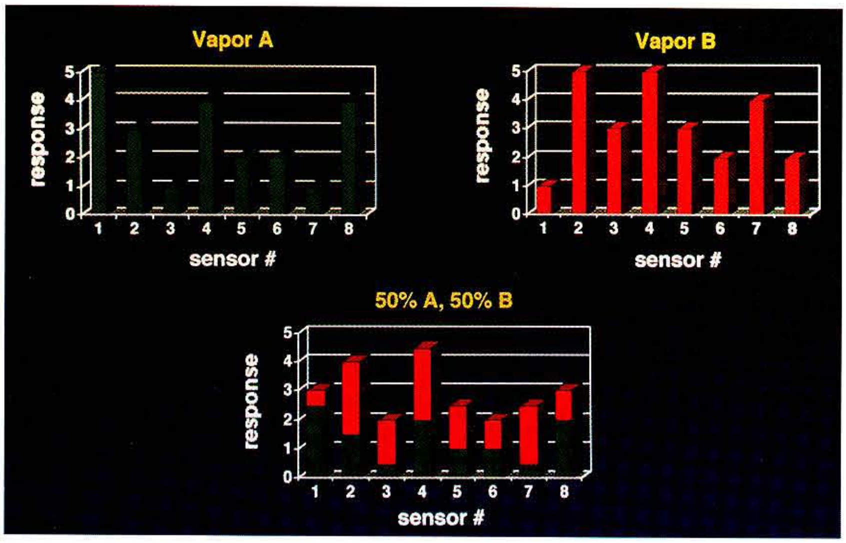

data-processing algorithm called a neural network, which basically mimics a whole bunch of interconnected nerve cells all firing messages back and forth at one another, and which can learn to recognize patterns. [See *E&S*, Summer 1990.] The network took our patterns, processed them, and identified each of our solvents by number. Such a neural network could easily be trained to recognize anything the nose can smell, as long as the sniff gives a reproducible pattern. Wei wrote the algorithm as a piece of software, but neural nets can also be built directly into chips as hardware, and Rod's working on that right now. Brett Doleman, a grad student in my group, is working with Rod's group to figure out how best to classify the different odorants.

Discriminating between pure vapors is a start, but what about mixtures? If we give the nose a mixture that's half methanol and half ethanol, the new pattern should be at the midpoint of the line segment connecting the two pure smells in odor space. Will the nose break this pattern down into the two known ones, or think it's a brand-new smell? It turns out that as long as the responses are linear, the mixture simply registers as the linear combination of the individual smells. If the responses are nonlinear, then we have to train the nose on the mixture as if it were a new compound, which is obviously a lot less useful.

Conversely, can we fool the nose by giving it a new compound? If we don't tell the nose that this is a new thing, will the nose tell us that it's smelling a linear combination of known smells? Or will the nose know that there's something new in the air? We took the data we got from

We want our nose to be able to tell the difference between a rose and a dead fish.

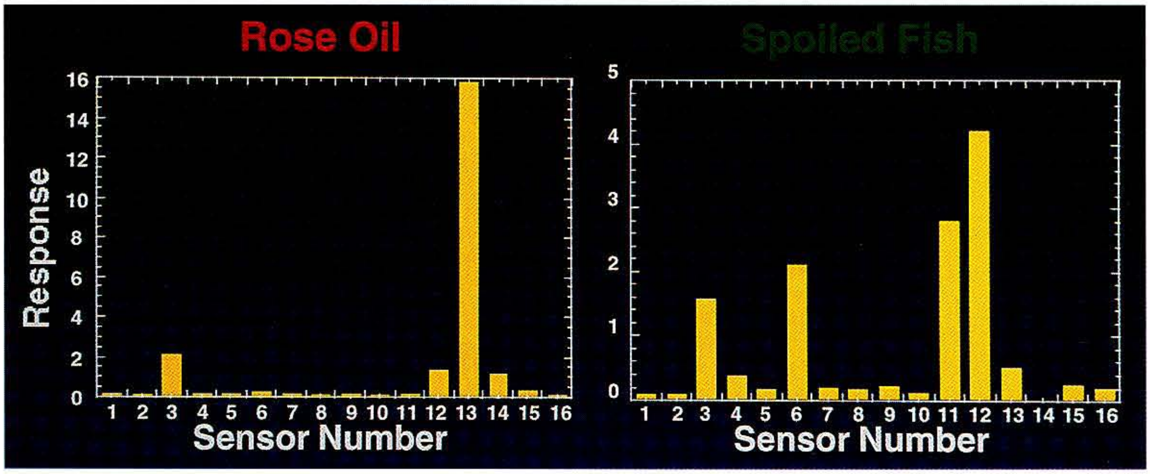

Brett Doleman (right) and Erik Severin (left) give a fish and a rose the once-over.

seven smells (methanol, ethanol, isopropanol, acetone, chloroform, hexane, and benzene), and tried to see if some combination of them would reproduce the pattern we got from ethyl acetate. The only stipulation was that all the components had to be positive—we didn't want a recipe that included, say, −15 percent ethanol. And with just that one constraint, we could not make the new smell out of any combination of the other seven smells. Of course, the more you know about the sample, the easier this is; the more different smells you're allowed to use, the harder it gets. There will be a happy medium somewhere, and we don't know what the trade-offs will be; but we do know that in certain instances we can't fool the nose. This is a very powerful test of the electronic nose's information content.

I said at the beginning that we want our nose to be able to tell the difference between a rose and a dead fish. So grad student Erik Severin went to the store and bought one generic fish, and put it in a flask. The human nose has evolved to smell raw meat, so the fish stank to our noses earlier than it did to the electronic nose. People were complaining by noon, but it took the nose all day to pick up the scent. Nevertheless, above right is the pattern Erik got for spoiled fish. (For unspoiled fish, the pattern is just water vapor, which we null out, so there is no pattern. We think this is what the human nose does, because people can't smell water vapor, either. It must be that our nasal sensor cells are in a constant-humidity environment, so they zero out water.) Erik also bought some rose oil, and its pattern (above, left) is quite unlike the fish's.

We can't yet tell red wine from white. We can, however, tell beer from wine from hard liquor by the alcohol content. We actually tried to tell wines apart initially—the Athenaeum is interested in sponsoring this project. In retrospect, perhaps we should have tried nulling out the water vapor with the wines, as we did with the fish.

Our electronic nose can't do what a mass spectrometer does, and say that there is one part per trillion of molecule X in the complex mixture we call "strawberries." But we don't always care about molecule X—sometimes we just want to know that it's strawberries and not raspberries. Sometimes we just want to know, does the cheese smell the same as it did yesterday, or has it rotted? The pattern-recognition approach to smelling does this very, very well.

You can imagine the quality-control applications for such a device. For example, cheese manufacturers pay people to sit on the production line and smell the cheese as it goes by. But they can only smell for two hours at a time, because their noses get saturated. And a quality-control lab can't analyze every single cheese with a gas chromatograph or mass spectrometer. You don't even know what you're looking for, necessarily—sometimes the cheese just smells bad! But a little electronic nose could just sit on the line all the time and say, "The cheese is the same as it was yesterday. The cheese is the same as it was yesterday. It's OK." The nose would beep whenever the cheese smelled different, and then you'd stop the line, and smell the cheese yourself to find out whether it really was OK or not.

Pinocchio, the new supernose, accommodates up to 20 sensors and lives in a stainless steel case on legs (right). The array of glassware in the background holds the pure liquids—the acetone, benzene, and so on—through which air is bubbled to generate the vapors that are then piped to the nose. The three black boxes are computer-controlled flow regulators. Pinocchio's case is so big because the sponges now bridge metal contacts plated onto glass slides, as seen below—an even simpler (and more reproducible) process than buffing down capacitors.

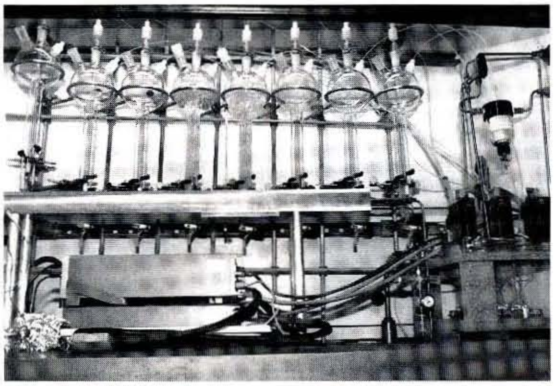

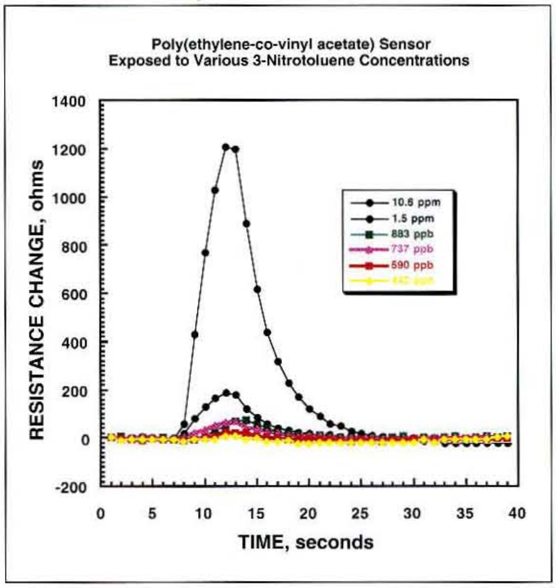

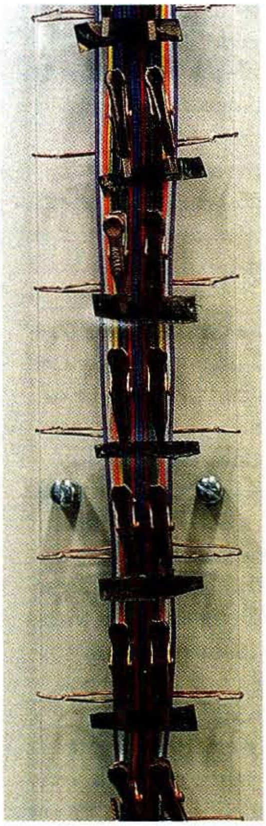

Similarly, you could program the nose to beep when a room smelled differently than normal. In a potentially hazardous situation, you might not even need to know what that difference was. You'd just leave the room (or not enter it, as the case may be), and wait until a more specific sensor had registered hydrogen sulfide from a gas leak, perhaps, and then you'd take appropriate action. NASA is interested in this for the space station, so they're helping sponsor our work. When humans will be up in confined atmospheres for years, in some cases, NASA doesn't necessarily know how to anticipate what might get into the air, and whether or not it will be safe to breathe. This way, they don't have to worry about designing a specific sensor for a substance they don't even know might be up there. The nose would just beep if something new appeared in the environment and the astronauts would reach for their oxygen masks. To this end, we are setting up a gas-handling system so that Pinocchio can try to measure toxic gases. We'd like to find out if Pinocchio can respond to gases that are odorless to us, such as carbon monoxide, but we don't know yet. There may be whole classes of gases that the nose can't smell.

I should also point out that we have no idea about the longevity of these noses. We've only been working on this project intensively for 18 months, so even our first nose isn't *that* old. It's too early to tell if this is really a durable device.

Right now, the nose's sensitivity is limited by our very primitive electronics. We use a simple voltmeter, just like the one you might have in your garage, and we can read what's known as 16-bit resolution. We can detect methanol in air down to 70 parts per million, which is roughly as good as a human nose can do. But we can detect 3-nitrotoluene, which is much less volatile, down to about 600 parts per billion, as shown above. (The less volatile a vapor is, the easier it is to detect at relatively low concentrations, because it prefers to stay liquid and is thus better held by the sponge.) We calculate that the ultimate detection limits will be about 10 parts per billion. Each sensor also needs what's called a Wheatstone bridge, which is adjusted to null out the sensor's baseline resistance. That way, we're measuring a small resistance on top of a zero. Right now, we're measuring a change of a few ohms on top of a 40,000-ohm baseline. We do care about the signal-to-noise ratio, because if there are things in very, very small concentrations that are critical for, say, distinguishing wines, we don't want to lose that information.

We can also adjust the sensor's threshold sensitivity by changing the ratio of conductor to insulator in the sponge. When the conductors are close enough to touch one another, the electrons essentially percolate from conductor to conductor through the points of contact. The electrons travel quite rapidly through the sponge (low resistance), even if they have to go through a tortuous path. On the other hand, if we swell the sponge to a little bit above that percolation threshold, they're going to have to hop across the intervening insulating regions. The resistance will jump dramatically with just that little bit of swelling. It's an on–off signal. We have

A meeting of the noses. Back row, from left: Sara Beaber, Bob Sanner, Nate Lewis, Erik Severin. Front row: Brett Doleman, Mark Lonergan.

Mike Freund makes an olfactory value judgment.

actually shown that this works, and you can see how this could be a very sensitive alarm. The alarm wouldn't tell us what's out there, because there's no pattern of linear responses that would allow the chip to recognize what the alarming substance is. But we can set the alarm's sensitivity by adjusting the percolation threshold. The more conductor we have, the more the sponge has to swell before the last percolation pathway is broken and the resistance skyrockets.

What are we going to do next? We'd like to do what we think human and dog noses do—use a large number of incrementally different elements. We'd like to make a million sensors on a chip. We think we know how to do it; it's just a fabrication issue. The electronics aren't the problem—building and adjusting a million different Wheatstone bridges; reading out a million different resistances; making a two-dimensional grid of 1,000 by 1,000 wires; and addressing each of a million individual intersections, even if the sensors are only 10 microns big, is not stretching current chipmaking technology. Such a chip would be a modest size—one centimeter by one centimeter—much smaller than *my* nose! Overlaying the grid of wires would be a matching matrix of little wells—also easily made—to hold the sponges. The issue is, how do we make a million different plastic sponges? Bob Grubbs and I had an idea, which Bob Sanner [PhD '78], now a visiting faculty member from Lawrence Livermore National Laboratory, is trying to implement.

We start with one monomer—one component of the plastic that makes up the sponge—that might like water, say, and another one that might like oil, and then spray the water-loving monomer left to right and the oil-loving one up and down while smoothly increasing the dilution of each. The wells will fill with an array of sponges gradated by water-loving-ness on one axis and oil-loving-ness on the other. It wouldn't even matter if the gradation varies slightly from chip to chip, because each chip would learn its own response. As long as the response is consistent every time the chip smells that smell, it doesn't matter what the details are.

We don't know yet how much benefit there will be in making a million different intermediate materials instead of just the two extreme cases. We do know that there's no point in doing so if all the responses are linear. Then the intermediate sponges' responses are just linear combinations of the two extremes, and there's no new information. But if the intermediate sponges behave differently, then they give us new signals to the extent that they have different swellabilities. The algorithms for this system are much like those for antenna design, it turns out, although as chemists we don't know enough about our "antennas" to decide just how many we need. That's one question we want to answer: what minimum number of elements is sufficient to distinguish very subtle differences in smells? A million channels is an awful lot of signals—can we get away with fewer? So Brett Doleman is working with Rod's group to figure out how many sensors we actually need.

And if we're building a chip with a million sensors, we could make a composite array in which some sensors beep when something appears in the environment at very low levels, and others wait a little bit and then tell you what that thing is. Or maybe you'd just get out of the room, depending on how many of the low-level sensors beep.

We'd also like to see if we can train this nose to make "human" value judgments—to say that this is a good perfume, or a bad perfume, or to set the price of a bottle of wine. Or, to restate the question more scientifically, can we assign a number to a fragrance based on these patterns? Can we assign a number to a bottle of wine or a cigar that somehow quantitatively reflects a human value judgment? This is a very interesting intellectual problem. We're working with the neural-network people to find how best to approach it. We're sniffing a fine wine versus a jug wine to see if there are any differences.

We're also interested in stereo smell. We can make a sponge so thin that it responds very quickly. It then becomes possible to use the time

One can envision a little robot equipped with stereo smell crawling along a fume-filled ventilation duct, coming to a junction and telling us that the smell is coming from the left, say, and following it back to its source.

difference between when a stimulus arrives at two separate arrays to determine where a smell is coming from. Except for cockroaches, there is no creature that has stereo smell—that can locate smells based on concentration gradients between the left and right parts of its nose. Even other insects, although they have two separate antennae, turn their heads to find out, much like we do. But one can envision a little robot equipped with stereo smell crawling along a fume-filled ventilation duct, coming to a junction and telling us that the smell is coming from the left, say, and following it back to its source. Building such a robot is, at this point, an engineering task. We know that the response is fast enough, in some systems, to allow us to build one, and we know that we can make the sensors small, but I don't know if we can make them *that* small yet. We might also want to align them along a rod, perhaps, instead of in a plane, to make insect-like antennae. It's very interesting to think about bringing the sense of smell into the same electronic regime that the sense of sight has been brought to by small TV cameras, and to use smell to guide robotic systems.

We didn't invent the idea of using conductive arrays to detect odors. The British thought of it first, although we didn't know about their work when we first started ours. In 1982, K. Persaud and G. H. Dodd built a nose that used bulk conducting polymers as the swellers, as we initially did, but there aren't really that many chemical differences between the various conducting polymers. Then, a few years later, several Japanese research groups started experimenting with tin oxide, an inorganic resistor from which you can make broadly responsive films. But in order to make the films different chemically, you have to sprinkle catalysts on the tin-oxide layer, and no one really knows how to control what those catalysts do. People have also experimented with quartz crystals, similar to what's in your watch. You launch a 100-megahertz wave, much like an ocean wave, across the surface of the crystal and look at the response. If odor molecules adsorb onto the surface, the wave's frequency will change measurably. But the electronics to launch 100-megahertz waves and then read tiny changes in their frequency are quite complex, and it's difficult to envision making an array of a million such sensors on a small chip. The crystal can also be coated with swellable plastic films, as in our work, but the signal transduction is much more difficult. The beauty of our approach is that we get all of our chemical differences from the insulating sponge, whose properties we can vary broadly and systematically in a very precisely controlled way. We only rely on the conducting phase to transduce the signal into electronic form.

In conclusion, I'd like to note that I got my first taste of research when I was an undergrad working for [Beckman Professor of Chemistry] Harry Gray. When I was ready to leave Caltech, I asked him, "What should I do? How will I know what a good project is?" Harry is a very wise person, and I remember to this day what his answer was. He said I should just follow my nose. And I did. □

Nathan S. Lewis earned his BS and MS in chemistry from Caltech in 1977, and his PhD in inorganic chemistry from MIT in 1981. He went on to Stanford and tenure before Caltech lured him back in 1988. He became a full professor in 1991.

Lewis, an electrochemist, first gained national attention in the year of cold fusion as a co-leader of the Caltech team whose meticulous experiments concluded that the phenomenon couldn't hold water, much less heat it (see E&S, *Summer 1989). But his "real" research has been in the development of liquid-based solar cells that produce electricity, chemical fuels, or both when struck by sunlight.*

Lewis, who has taught freshman chemistry for the past eight years, is also the electromotive force behind the Chemistry Animation Project (CAP) videos (see E&S, *Fall 1994).*

This article was adapted from a recent Watson lecture.

Dynamics

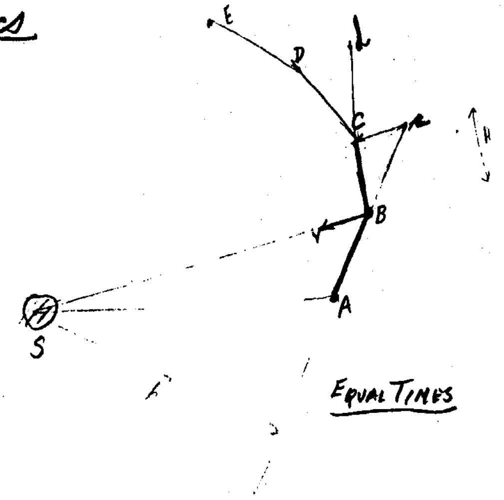

EQUAL TIMES

a straight line, no force
C. Force reduced to impulse at C means motion is to C rather than c.

~~CB is in plane ABS~~

area ABS = BcS
BcS = BCS ∴ ABS=BCS

or
Equal areas in equal times.

~~Vel inversely~~ Vel × h = const.

Equal angles in times prop to square of distance.

Also orbit in plane.

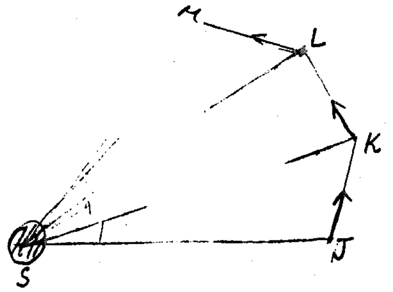

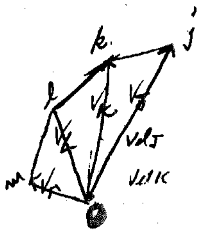

jk is ll to KS
lk is ll to LS
lm is ll to MS etc.
lk = jk = lm

∴ Ends of Vel on regular polygon → circle.

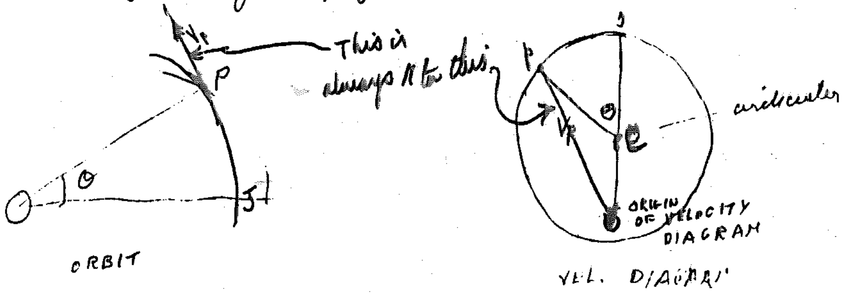

Feynman's Lost Lecture
The Motion of Planets Around the Sun

This lecture is an opportunity for anyone who has mastered plane geometry to see the great Feynman at work!

by David L. Goodstein
and Judith R. Goodstein

All that remained of Feynman's "lost" lecture were the audiotape and a few pages of notes that Feynman jotted down for himself. Most of the lecture derives from the page at left: the figure at upper left is copied from Newton's *Principia*.

For the lecture itself Feynman fans will have to read the book of the same title. Published last month, it may be found in local bookstores or ordered directly from the publisher, using the coupon on page 21. But for a little foretaste, we excerpt here Judith Goodstein's Preface, the Introduction, and a part of David Goodstein's reminiscences. (Copyright © 1996 by the California Institute of Technology. Reprinted with permission of the publisher, W. W. Norton & Company, Inc. Proceeds from the book will be used to support scientific and scholarly research at Caltech.)

Preface

This is the story of how Feynman's lost lecture came to be lost, and how it came to be found again. In April 1992, as Caltech's archivist, I was asked by Gerry Neugebauer, the chairman of the Division of Physics, Mathematics and Astronomy, to go through the files in Robert Leighton's office. Leighton was ill and had not used his office for several years. Marge Leighton, his wife, had told Neugebauer that it was all right to clean out the office—she'd already collected her husband's books and personal effects. I could take what I wanted for the archives, and the division would dispose of the rest.

Besides heading the Division of Physics, Mathematics and Astronomy from 1970 to 1975, Leighton, together with Matthew Sands, had overseen the editing and publication of Richard Feynman's two-year course of lectures in introductory physics, delivered to Caltech freshmen and sophomores. The lectures, published in the early 1960s in three volumes by Addison-Wesley, dealt with virtually every subject in physics, with a point of view that remains fresh and original to this day. I was hoping to find some tangible evidence of the Leighton-Feynman collaboration.

It took me a couple of weeks to sift through the stacks of paper, which were stashed everywhere, but Leighton didn't disappoint me. I unearthed two folders, one marked "Feynman Freshman Lectures, unfinished," another labeled "Addison-Wesley," wedged between budget sheets and purchase orders from earlier decades and reams of yellowing computer paper covered with endless columns of numbers, all thrown together in a storage closet just outside his office. Leighton's correspondence with the publisher contained details about the format, the color of the cover, comments by outside readers, adoptions at other schools, and estimates of how well the volumes would sell. That folder I put in the "Save" pile. The other folder, the one containing the unedited Feynman physics lectures, I carried back to the archives myself.

In his June 1963 preface to *The Feynman Lectures on Physics,* Feynman commented on some of the lectures not included there. He'd given three optional lectures in the first year on how to solve problems. And, indeed, three of the items in Leighton's folder turned out to be the raw transcripts for Reviews A, B, and C, offered by Feynman in December 1961. A lecture on inertial guidance, which Feynman gave the following month, didn't make the cut either—an unfortu-

***Signora e Signore Goodstein*, portrayed by artist (and Goodstein friend) Igor Bitman as contemporaries of Galileo. The signora, however, holds Feynman's notes for the 1964 lecture, and the more familiar Goodsteins can be seen as modern Roman tourists in the painting above the book (published 1996).**

But in the end, we decided that the only lecture that still had the vitality, originality, and verve we associated with Feynman's presence in the classroom was the 1964 lecture on planetary motion —the one lecture that demanded a full complement of blackboard photographs. And we didn't have them.

nate decision, according to Feynman—and I found a partial transcript of this lecture in Leighton's folder. The folder also contained the unedited partial transcript of a later lecture, dated March 13, 1964, along with a sheaf of notes in Feynman's handwriting. Entitled "The Motion of Planets Around the Sun," it was an unorthodox approach to Isaac Newton's geometric demonstration of the law of ellipses in the *Principia Mathematica*.

In September 1993, I had occasion to draw up a list of the original audiotapes of the Feynman lectures, which had also been contributed to the archives. They included five lectures that were not to be found in the Addison-Wesley books. Then I remembered the five unpublished lectures in Leighton's file; sure enough, the unedited transcripts matched the tapes. The archives also had photographs of the blackboard diagrams and equations for four of these lectures—the four mentioned by Feynman in his preface—but I could find none for the March 1964 lecture on planetary motion. (In the course of selecting illustrations for this book, I did stumble upon one photograph of Feynman taken during this special lecture. It is reproduced here [on page 18].) Although Feynman had given Leighton his notes on the 1964 lecture, which included sketches of his blackboard drawings, Leighton apparently decided not to include it in the last (1965) volume of *The Feynman Lectures on Physics*, which dealt primarily with quantum mechanics. In time, this lecture was forgotten. For all practical purposes, it was lost.

The idea of rescuing all five unpublished lectures from oblivion appealed to David and me. So the following December, when we went, as we often do, to the Italian hill town of Frascati, we took along copies of the tapes, the transcripts, the blackboard photographs, and Feynman's notes. In the course of the next two weeks, we listened to the tapes, took notes, laughed at the jokes, strained to hear the students' questions and Feynman's answers after each lecture was over, took more notes. But in the end, we decided that the only lecture that still had the vitality, originality, and verve we associated with Feynman's presence in the classroom was the 1964 lecture on planetary motion—the one lecture that demanded a full complement of blackboard photographs. And we didn't have them. Reluctantly, we abandoned the project.

Or so I thought. As it turned out, bits and pieces of the lecture haunted David, especially when he came to teach the same material in freshman physics the following year. He had the tape. But could he reconstruct the blackboard demonstrations from the few tantalizing sketches in Feynman's notes and the few words Feynman had jotted down more for himself than for the students? "Let's try again," he announced, early in December 1994, as we were packing for a trip through the Panama Canal. This time, we would take along only the transcript of the 1964 lecture, the lecture notes, and selected pages from Kepler's *The New Astronomy* and Newton's *Principia* for good measure.

It took the S.S. *Rotterdam* 11 days to sail from Acapulco to Fort Lauderdale. For two or three hours each day, David would hole up in our cabin and work on deciphering Feynman's lost lecture. He began, as Feynman had, with Newton's geometrical proofs. The initial break came when he was able to match up Feynman's first sketch [here on page 14] with one of Newton's diagrams, on page 40 of the Cajori edition of the *Principia*. We'd been at sea for three, maybe four days, Costa Rica's shoreline plainly visible, when David announced that he, too, could follow Newton's line of reasoning up to a point. By the time we'd exchanged the Pacific Ocean for the Atlantic, he was completely absorbed in Feynman's sparse, neatly labeled pencil drawings of curves and angles and intersecting lines. He stayed in the cabin, ignoring the scenery in favor of geometric figures—Newton's, Feynman's, and his own—longer and longer each morning and in the evening as well. When we arrived in Fort Lauderdale, on December 21, he knew and understood Feynman's entire argument. On the plane home, the book took shape. . . .

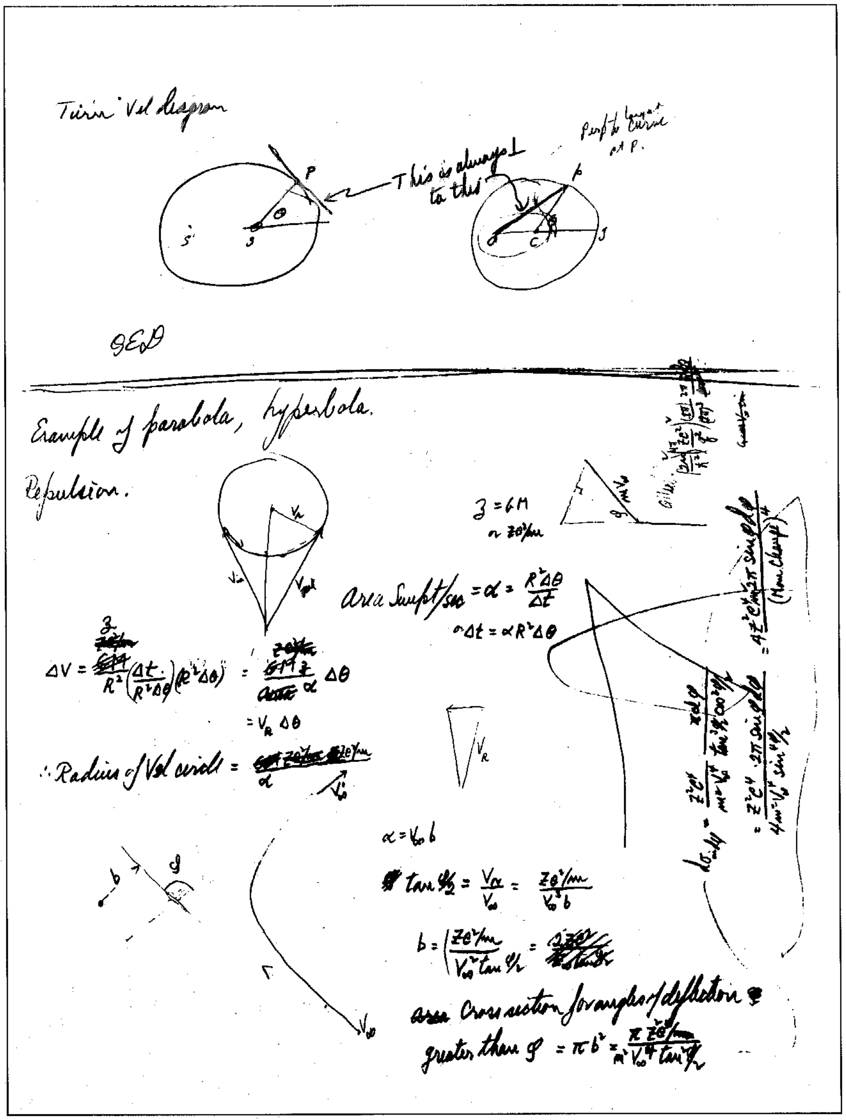

Another page of Feynman's lecture notes shows the final steps of his proof of the law of ellipses (above the line) and Rutherford's law of scattering (below the line).

Introduction

I would rather discover a single fact, even a small one, than debate the great issues at length without discovering anything at all.
—Galileo Galilei

This book is about a single fact, although certainly not a small one. When a planet, or a comet, or any other body arcs through space under the influence of gravity, it traces out one of a very special set of mathematical curves—either a circle or an ellipse or a parabola or a hyperbola. These curves are known collectively as the conic sections. Why in the world does nature choose to trace out in the sky those, and only those, elegant geometrical constructions? The problem turns out to be not only of profound scientific and philosophical significance but of immense historical importance as well.

In August of 1684, Edmund Halley (after whom the comet would be named) journeyed to Cambridge to speak to the celebrated but somewhat strange mathematician Isaac Newton about celestial mechanics. The idea was abroad in scientific circles that the motions of the planets might be a consequence of a force from the Sun that diminished as the inverse square of the distance between the Sun and the planets, but no one had yet been able to produce a satisfactory demonstration. Yes, Newton let on, he had been able to demonstrate that such a force would give rise to elliptical orbits—exactly what Johannes

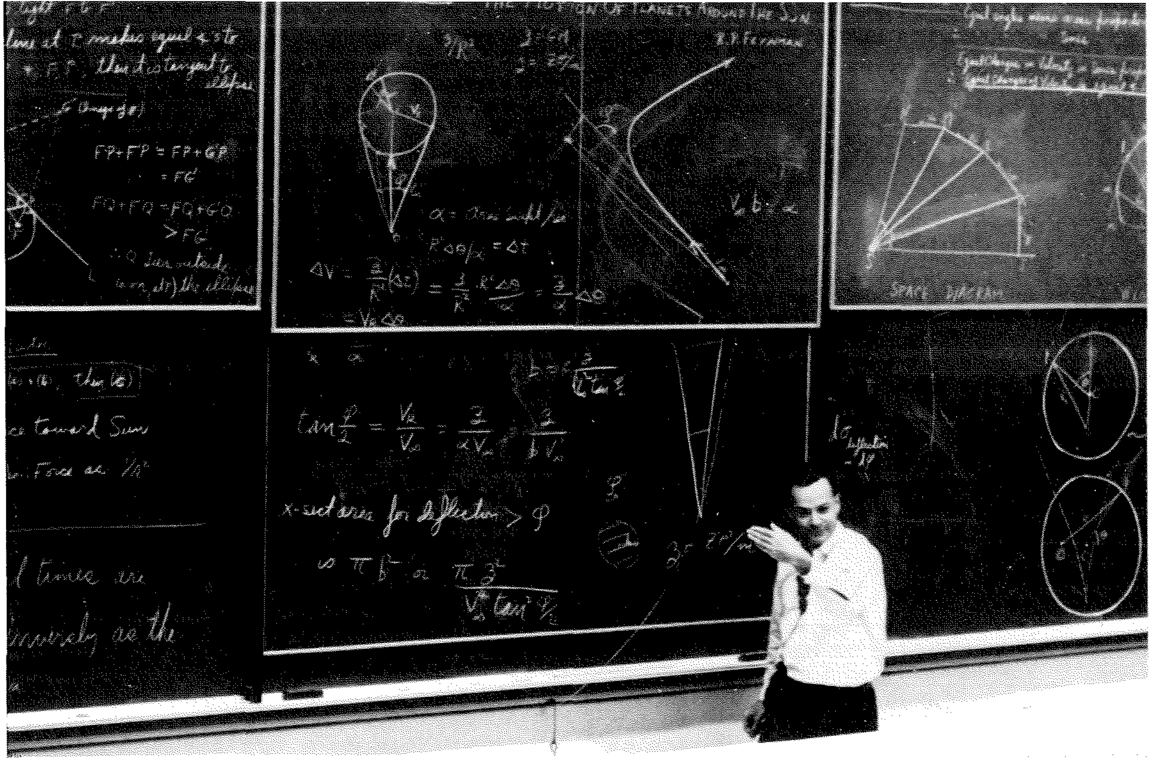

Although blackboard photographs of all the rest of Feynman's lectures survive, none were ever found for his 1964 lecture on the motion of the planets around the sun. Only this one shot surfaced of Feynman actually giving the lecture—with part of the blackboard behind him.

Feynman tried to follow Newton's proof, but he couldn't get past a certain point, because Newton made use of arcane properties of conic sections. . . that Feynman didn't know. So, as he says in his lecture, Feynman cooked up a proof of his own.

Kepler had deduced some 70 years earlier from observations of the heavens. Halley urged Newton to let him see the demonstration. Newton apparently begged off, saying he had misplaced it, but promised to work it out again and send it to Halley. In fact, a few months later, in November 1684, Newton did send Halley a nine-page treatise in which he demonstrated that an inverse-square law of gravity, together with some basic principles of dynamics, would account for not only elliptical orbits but Kepler's other laws of planetary motion as well, and more besides. Halley knew that he held in his hands nothing less than the key to understanding the universe as it was then conceived.

He urged Newton to let him arrange for its publication. But Newton was not entirely satisfied with this work and delayed, wanting to make revisions. The delay lasted almost three years, during which Newton, now thoroughly hooked on the problem, seems to have done nothing else but work on it. What emerged at the end, in 1687, was *Philosophiae Naturalis Principia Mathematica*, Newton's masterpiece and the book that created modern science.

Nearly 300 years later, the physicist Richard Feynman, apparently for his own amusement, undertook to prove Kepler's law of ellipses himself, using no mathematics more advanced than elementary plane geometry. When he was asked to give a guest lecture to the Caltech freshman class in March 1964, he decided to base it on that geometric proof. . . .

The discovery of Feynman's lost lecture notes affords us an extraordinary opportunity. For most people, Feynman's fame rests on the picaresque exploits, recounted in two anecdotal books (*"Surely You're Joking, Mr. Feynman!"* and *"What Do You Care What Other People Think?"*) which he produced late in life in collaboration with Leighton's son, Ralph. The stories in these books are amusing enough, but they take on a special resonance because the protagonist was also a theoretical physicist of historic proportions. Yet for the nonscientist reader there is no way to peer into Feynman's mind and see that other side of him—the powerful intellect that left an indelible imprint on scientific thought. In this lecture, however, Feynman uses all his ingenuity, insight, and intuition, and his argument is not obscured by the layers of mathematical sophistication that made most of his accomplishments in physics impenetrable to the uninitiated. This lecture is an opportunity for anyone who has mastered plane geometry to see the great Feynman at work!

Why did Feynman undertake to prove Kepler's law of ellipses using only plane geometry? The job is more easily done using the powerful techniques of more advanced mathematics. Feynman was evidently intrigued by the fact that Isaac Newton, who had invented some of those more advanced techniques himself, nevertheless presented his own proof of Kepler's law in the *Principia* using only plane geometry. Feynman tried to follow Newton's proof, but he couldn't get past a certain point, because Newton made use of arcane properties of conic sections (a hot topic in Newton's time) that Feynman didn't know. So, as he says in his lec-

"I couldn't do it. I couldn't reduce it to the freshman level. That means we don't really understand it."

ture, Feynman cooked up a proof of his own.

Moreover, this is not just an interesting intellectual puzzle that Feynman has doodled with. Newton's demonstration of the law of ellipses is a watershed that separates the ancient world from the modern world—the culmination of the Scientific Revolution. It is one of the crowning achievements of the human mind, comparable to Beethoven's symphonies, or Shakespeare's plays, or Michelangelo's Sistine Chapel. Aside from its immense importance in the history of physics, it is a conclusive demonstration of the astonishing fact that has mystified and intrigued all deep thinkers since Newton's time: nature obeys mathematics. . . .

Feynman: A Reminiscence . . .

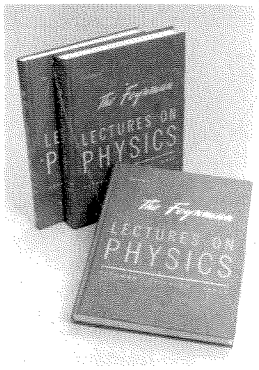

In 1961, Feynman undertook a project that would have far-reaching impact on the entire scientific community. He agreed to teach the two-year sequence of introductory physics courses that were required of all incoming Caltech students. His lectures were recorded and transcribed, and all the blackboards he filled with equations and sketches were photographed. From this material, his colleagues Robert Leighton and Matthew Sands, with help from Rochus Vogt, Gerry Neugebauer, and others, produced a series of books called *The Feynman Lectures on Physics*, which have become genuine, enduring classics of the scientific literature.

Feynman was a truly great teacher. He prided himself on being able to devise ways to explain even the most profound ideas to beginning students. Once, I said to him, "Dick, explain to me, so that I can understand it, why spin one-half particles obey Fermi-Dirac statistics." Sizing up his audience perfectly, Feynman said, "I'll prepare a freshman lecture on it." But he came back a few days later to say, "I couldn't do it. I couldn't reduce it to the freshman level. That means we don't really understand it."

Feynman delivered the *Feynman Lectures* to the Caltech freshman class in the academic year 1961–62 and to the same students as sophomores in 1962–63. His taste in physics topics was perfectly eclectic; he devoted just as much creative energy to describing the flow of water as to discussing curved spacetime. Of all the subjects he covered in that introductory course, perhaps his most impressive accomplishment is the presentation of quantum mechanics (Volume III of the series); in only slightly disguised form, it is the new view of quantum mechanics that he himself had developed.

While Feynman was a riveting, dramatic performer in the classroom, the period 1961–62 was to be the only time he ever taught formal undergraduate courses. For the rest of his professional life, before and after, he taught only courses designed for graduate students. The lecture that is the subject of this book was not part of the original course but rather a "guest lecture" to the freshman class at the end of the winter quarter in 1964. Rochus Vogt had taken over the teaching of introductory physics by then, and he invited Feynman to give the talk as a treat for the students. The *Feynman Lectures* were never successful as introductory textbooks—not even at Caltech, where they originated. They would instead make their lasting contribution as a source of insight and inspiration for accomplished scientists who had learned their physics by more conventional means.

In the immediate aftermath of his Nobel Prize in 1965, Feynman suffered a brief period of dejection, during which he doubted his ability to continue to make useful, original contributions at the forefront of theoretical physics. It was during this time that I joined the Caltech faculty. The Feynman physics course was now being taught by Gerry Neugebauer. When Feynman himself had been giving the lectures, Gerry, as a young assistant professor, had had the difficult job of making up homework assignments from them for the 200 or so students—difficult in large part because no one, maybe not even Feynman himself, knew in advance exactly what he was going to say. Just as he did for the lost lecture in Chap-

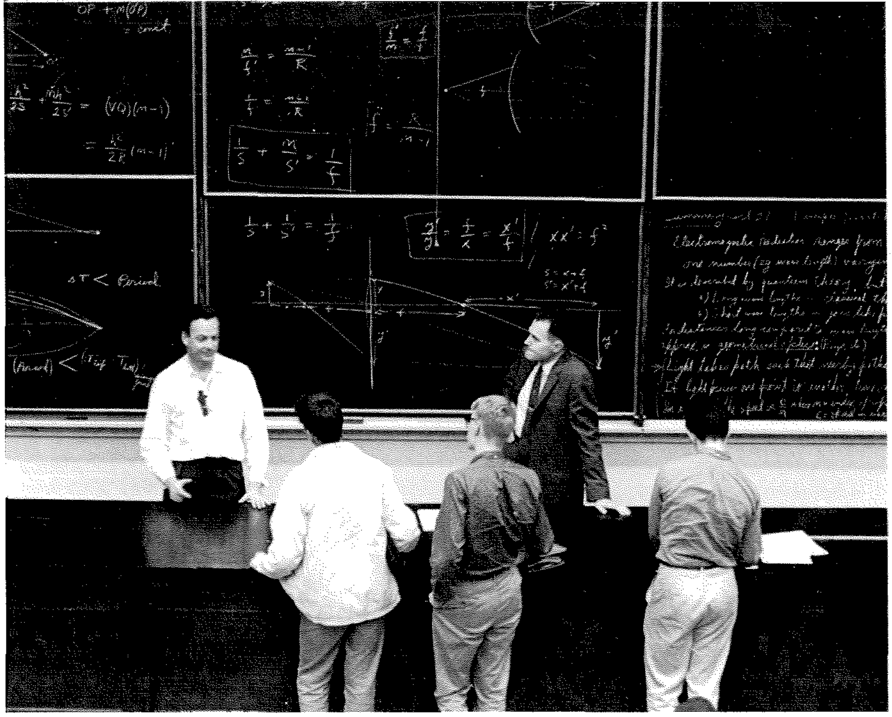

“You have to worry about your own work and ignore what everyone else is doing.”

1962: Feynman and Leighton in 201 East Bridge.

ter 4 in this book, Feynman would come to class with no more preparation than one or two pages of scribbled notes. Neugebauer, to make his own task somewhat easier, would join Feynman, Leighton, and Sands for lunch after each lecture, in the Caltech cafeteria, known to generations of students as “the Greasy”; Caltech’s elegant faculty club, the Athenaeum, was not Feynman’s style. During these lunches, the lecture would be rehashed, with Leighton and Sands competing to score points with Feynman, while Neugebauer desperately tried to figure out the essence of the lecture.

Now, in 1966, Neugebauer was giving the lectures, and I was pressed into service as a T.A. (teaching assistant), in charge of one of the small recitation sections that supplemented the main course of lectures. The by now traditional lunches at the Greasy continued, with Feynman still in attendance. It was here that I first really got to know him, mostly exchanging ideas with him on how to teach physics. That fall, he got an invitation to give a public lecture at the University of Chicago the following February. At first he was inclined to refuse (invitations to speak arrived almost daily), but then he decided to accept and to talk about our ideas on teaching, if I would agree to come with him. He said that he would pay for my travel expenses out of the absurdly large ($1,000) honorarium they were offering. I thought the matter over carefully for a microsecond or so, and agreed to go. When he told the University of Chicago that I would be joining him, they were no doubt mystified about who I was and why I was needed, but they invited me with good grace and paid my way in the bargain.

At Chicago, Feynman and I shared a suite in the Quadrangle Club, the university’s faculty club. On the evening after his talk, we had dinner at the home of friends, Val and Lia Telegdi. The next morning, I wandered down to the faculty club dining room for breakfast a bit late. Feynman was already there, eating with someone I didn’t know. I joined them, introductions were mumbled but not heard, and I sleepily drank my morning coffee. As I listened to the conversation, it dawned on me that this person was James Watson, discoverer with Francis Crick of the double-helical structure of DNA. He had with him a typed manuscript entitled *Honest Jim* (the title would later be changed by the publisher to *The Double Helix*), which he wanted Feynman to read, in the hope that Feynman might contribute something to the dust jacket. Feynman agreed to look at the manuscript.

That evening there was a cocktail party and dinner in Feynman’s honor at the Quadrangle Club. At the cocktail party, the worried host asked me why Feynman wasn’t there. I went up to the suite and found him immersed in Watson’s manuscript. I insisted that since he was the honoree, he had to come down to the party. Reluctantly, he did, but he fled after dinner at the earliest moment permitted by civility. When the party broke up, I went back up to the suite. Feynman was waiting for me in the living room. “You’ve gotta read this book,” he said.

“Sure,” I said, “I’ll look forward to it.”

“No,” he shot back, “I mean right now.” And so, sitting in the living room of our suite, from one to five in the morning, with Feynman waiting impatiently for me to finish, I read the manuscript that would become *The Double Helix*. At a certain point, I looked up and said, "Dick, this guy must be either very smart or very lucky. He constantly claims he knew less about what was going on than anyone else in the field, but he still made the crucial discovery." Feynman virtually dove across the room to show me the notepad on which he'd been anxiously doodling while I read. There he had written one word, which he had proceeded to illuminate with drawings, as if he were working on some elaborate medieval manuscript. The word was "Disregard!"

"That's what I'd forgotten!" he shouted (in the middle of the night). "You have to worry about your own work and ignore what everyone else is doing." At first light, he called his wife, Gweneth, and said, "I think I've figured it out. Now I'll be able to work again!". . .

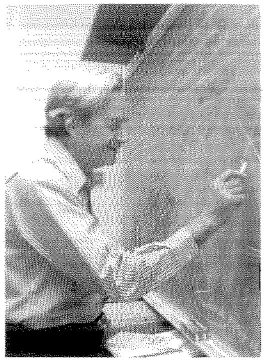
Feynman in 1985—still rarely far from a blackboard.

Feynman's lost lecture on planetary motion was by no means the only ad-hoc lecture he ever gave for the benefit of the Caltech undergraduates. Over the years, he was often asked to make a guest appearance, and he nearly always complied. The last of these guest lectures took place on Friday morning, December 4, 1987. I was now teaching the freshman introductory physics course, and he agreed to my request to give the final lecture of the fall quarter.

The subject of Feynman's lecture on this occasion was to be curved spacetime (Einstein's theory of general relativity). Before starting, however, he had a few words to say on a subject that excited him greatly. That year a supernova had occurred at the edge of our galaxy. "Tycho Brahe had his supernova," Feynman told the class, "and Kepler had his. Then there weren't any for 400 years. Now I have *mine*!"

This remark was greeted with a stunned silence by the freshmen, who had reason enough to be in awe of Feynman even before he opened his mouth. Dick grinned with obvious pleasure at the effect he had created, and defused it in the next breath. "You know," he mused, "there are about a hundred billion stars in a galaxy—10 to the 11th power. That used to be considered a huge number. We used to call numbers like that 'astronomical numbers.' Today it's less than the national debt. We ought to call them 'economical numbers.'" The class dissolved in laughter, and Feynman went on with his lecture.

Richard Feynman died two months later, on February 15, 1988.

Caltech Registrar and Archivist Judith Goodstein and her husband, David (professor of physics and applied physics, the Frank J. Gilloon Distinguished Teaching and Service Professor, and vice provost) have been, separately, loyal contributors to Engineering & Science *(see page 40). Their last joint appearance in these pages was in October 1980, when they viewed the scientific method from different, but concurring, perspectives.*

FEYNMAN'S LOST LECTURE, by David and Judith Goodstein, can be ordered directly from the publisher. It comes packaged with a compact disk recording of the entire original lecture in an attractive boxed set. Please send your order, along with a check or money order for $35.00 per set to:

Department FM
W. W. Norton & Company, Inc.
500 Fifth Avenue
New York, NY 10010

Please send me ____ copies of FEYNMAN'S LOST LECTURE at $35.00 per set.

I enclose ____check
____ money order for $________ (New York and California residents please add sales tax)

Name________________________________

Address________________________________

City, State, Zip________________________________

Turbulence, Fractals, and CCDs

by Paul E. Dimotakis

While turbulence has captured people's imagination for millennia, the beginnings of our current understanding date from the 1930s and 1940s.

This was one of the first images ever to capture a high level of detail within a turbulent flow—detail enough to convince Benoit Mandelbrot (MS '48, Eng '49) that turbulence was an example of the class of mathematical creatures he called fractals. The picture was made by injecting a jet of water carrying a fluorescent dye into a tank of standing water. A laser then sliced through the flow, lighting up only the dye molecules within that slice. From Dimotakis, Miake-Lye, and Papantoniou, *Physics of Fluids*, 1983.

Two and a half thousand years ago, the philosopher Heracleitus sat on the banks of a small river near the ancient Greek town of Ephesus, in Asia Minor, tossing little sticks into the water. As he watched them float irregularly downstream on the turbulent river, he remarked, "Twice into the same river you could not enter." Despite reasonably steady initial conditions (the spring is the same) and boundary conditions (the banks are the same), the turbulent flow in the river is never the same twice. Heracleitus put in a nutshell the problem that bedevils researchers in turbulence to this day—how can you analyze something that changes randomly and uncontrollably from moment to moment?

Now, turbulence isn't always a curse. It is often a blessing—without turbulence, we wouldn't have much animal life on this planet. When we exhale, for example, our breath comes out as a little jet of gas that mixes with the surrounding air. Then, when we inhale, only a very small part of the exhaled carbon dioxide comes back in. Without turbulence, we would reinhale most of it, although, as my 11-year-old son Manolis noted, not for long. And turbulent vortex rings shed from our heart valves are crucial in helping them close. It doesn't take a large change in the flow through the valve to alter its dynamics and cause life-threatening difficulties, as the work of Professor of Aeronautics Mory Gharib (PhD '83) and others is helping us appreciate. Any creatures that didn't master the dynamics of turbulence in their breathing and internal circulation, as well as in other important turbulent-flow phenomena (such as swimming and flying) would have rapidly gone extinct.

More broadly, we rely on turbulent mixing to sustain and drive all kinds of things, including many flow and combustion devices in which chemical reactions occur. Consider a jet engine, for example. Our ability to fly at high speeds is limited, in part, by our ability to mix fuel and air quickly and efficiently at flow speeds that are high compared to the speed of sound, i.e., at high Mach numbers. The inherent unsteadiness that leads to and sustains turbulence tends to diminish as the Mach number increases. Flows that would be strongly turbulent at low Mach numbers often aren't at high Mach numbers, and less mixing results. But at the same time that we're trying to maximize mixing within the engine, we need to minimize mixing (and thus heat transfer) in the flow along the engine's interior surfaces, so that they don't melt. Partly as a consequence of such considerations (and many others—flight, especially commercial flight, is a complex interplay between economic as well as aerodynamic forces), we've been flying at the same speed for the last 30 years or so—except for the Concorde, which is not economically viable because of its high fuel consumption for its size. That's a remarkable statistic, considering commercial aviation's enormous progress in so many other ways. So, if you ask whether it will always take this long to fly across the Pacific, or to Eastern Europe, the answer partly depends on learning how to both promote and limit turbulent mixing.

While turbulence has captured people's imagination for millennia, the beginnings

This shadowgraph, and others like it, provided the first evidence of large-scale order in turbulent flows. Here, a stream of nitrogen at four atmospheres pressure (left) traveling at 1,000 centimeters per second blows by a helium-argon mixture (right) with the same density and pressure but traveling only 380 centimeters per second. The zone where they mix is made visible by their different refractive indices, in exactly the same way that you see heat shimmers when looking across a blacktop parking lot in August. From John H. Konrad's PhD thesis, 1976.

of our current understanding date from the 1930s and 1940s, when Ludwig Prandtl in Germany, Theodore von Kármán at Caltech, G. I. Taylor in England, A. N. Kolmogorov in the then Soviet Union, and others elsewhere proposed that descriptions based on local averages and other statistical tools such as spectral analysis could provide useful information about the nature of turbulence, and that it was possible to describe, and even predict, some aspects of turbulent flows. Because turbulence is chaotic, irregular, and nondeterministic, statistical treatments appeared to be the only possible way to describe it. These methods often work well, in fact, but it's difficult to extract much information from them about many properties of turbulence—such as drag, entrainment, and mixing—that are important to engineers.

Then, in the late '60s to early '70s, largely as a result of experiments initiated at Caltech by Garry Brown, then a research fellow and later a professor of aeronautics (and now head of Mechanical Engineering at Princeton), and Anatol Roshko (MS '47, PhD '52), von Kármán Professor of Aeronautics, Emeritus, the picture changed. Brown and Roshko found that, despite its obvious disorder, turbulent flow is organized to a fair extent, primarily at its largest scales, as shown in this photo (left) by one of Roshko's students. The dynamical properties that engineers were struggling to understand depended on the behavior of these large-scale structures, which were present even in intensely turbulent flows. These discoveries provided hope that it would be possible to describe the dynamics of turbulence in relatively simpler terms than previously thought necessary. Heracleitus' epigram didn't sound so daunting anymore. You still couldn't step twice in the same river, but at least you could describe the river better.

From the days of von Kármán, much of the progress in turbulence has rested on the use of flow-visualization techniques. It's difficult to see a complicated, nonperiodic geometrical pattern in a sequence of numbers, when such a pattern may be obvious from a casual glance at a photograph. Our brain has an uncanny ability to decipher complexity and discover order in visual data. A two-year-old can look at a drawing and tell you whether it's a cat or a dog; that distinction cannot be made easily using the largest conventional computers.

Unfortunately, such visual data tended to be "soft" back then, because extracting quantitative information from pictures was difficult. The data were recorded on photographic film, a few measurements were made from the pictures, and a limited statistical analysis was laboriously done by hand. The "hard" mathematical treatments that most researchers were interested in mostly relied on point measurements. You'd put an instrument, or an array of instruments, in the flow and get a series of readings as the flow moved past the array. Only a few numbers—mean values, or, at most, spectra from wave analyzers—were recovered. The rich continuum of spatial and temporal properties of turbulence could not easily be discerned in such data.

In retrospect, the evidence of large-scale order in turbulence can be seen in the old point-mea-

The laser-induced fluorescence apparatus (left) is essentially a high-tech aquarium on a stand. The jet, tagged with a fluorescent dye, shoots straight down from the plenum, whose lower surface is immersed in the reservoir water. A rotating mirror of adjustable height sweeps the laser rapidly through the flow perpendicularly to its direction of travel, illuminating a cross section of the flow (right). The CCD camera then records the frozen slice of turbulence through the tank's glass bottom. With slightly different optics, the system can also take slices along the flow's axis, as in the picture on page 30.

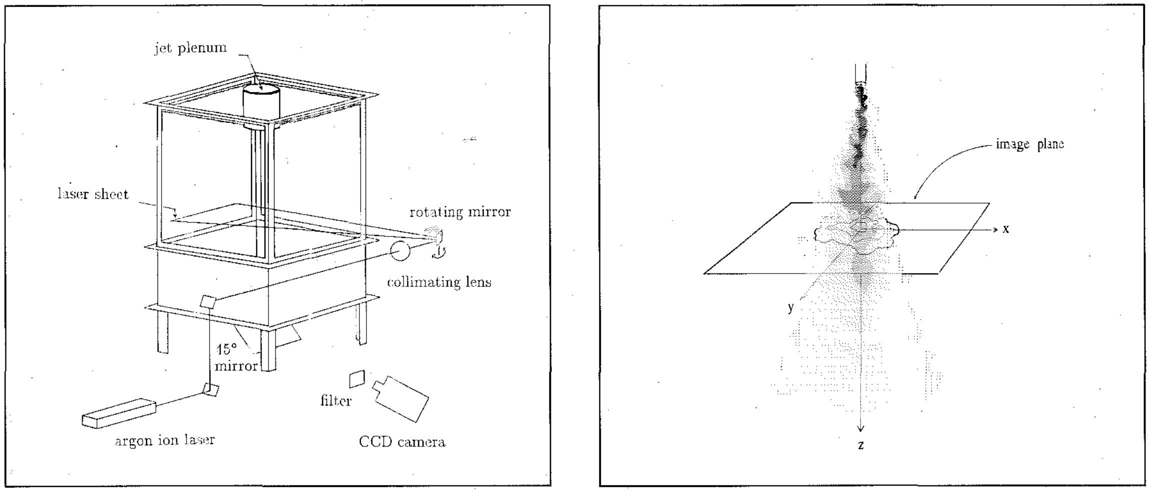

surement data, but it was so contrary to expectations that it was overlooked. In the late '40s, for example, Hans W. Liepmann, now the von Kármán Professor of Aeronautics, Emeritus, but then a young Caltech professor, was analyzing point-velocity data from a hot-wire array and found strong evidence that the points near the edge of a turbulent flow were only turbulent intermittently. Brown and Roshko's experiments some 20 years later showed why: the probe was periodically being engulfed by the largest vortices—the ones you can see in the photo on the opposite page—in the same way that a piling just above the tide line gets immersed in the swash from each breaking wave. Liepmann also noticed that these vortices tended to pair up. However, von Kármán pooh-poohed the results, and there the matter stood for two decades.

Today, with the advent of CCD (charge-coupled device) cameras, and the image-compression and data-handling technology developed by Caltech's Jet Propulsion Laboratory (JPL) and elsewhere to send us breathtaking images from planets we will not be able to visit ourselves in the foreseeable future, we can record two-dimensional information at a million or more points simultaneously, with an accuracy that matches yesterday's best point-measuring instruments. A 1,000-x-1,000-pixel CCD array is equivalent to placing one million measuring instruments in the flow, all recording at the same time without disturbing the flow or getting in each other's way.

Back in the mid-1970s, our lab at Caltech was the first to develop laser-induced fluorescence techniques for fluid mechanics. Coupled with digital CCD imaging, these methods have provided quantitative, multidimensional (field, as opposed to point) flow measurements. (We used the first [linear] CCD arrays at about the same time.) Much of our work has focused on turbulence generated by shooting a jet of water, tagged with a fluorescent dye, into a reservoir of quiescent, untagged water. A laser selectively excites the dye, which fluoresces with an intensity proportional to its concentration. A CCD camera then records the fluorescence, which shows how the jet fluid mixes with the entrained reservoir fluid. By rapidly sweeping a laser across the jet (or pulsing a sheet of laser light) we can, in effect, freeze any slice of the flow at an instant of time.

With this dense, quantitative turbulent-flow data, we can begin to ask questions about the complex geometry that turbulence generates. Geometry, to most people, brings to mind triangles and circles, spheres and cubes—the simple, regular shapes that fascinated the ancient Greeks. Well, turbulence isn't so kind. It generates irregular shapes that aren't amenable to the analyses that Thales of Miletus (near Ephesus); Pythagoras, a short swim away from Ephesus on the island of Samos; and many others developed, and that were so eloquently documented by Euclid in Alexandria in the third century B.C. How do we describe the geometric characteristics of the interface between the jet fluid and the entrained reservoir fluid in the photo on page 22, for example? How can we measure that interface's surface-to-volume ratio (a way of quantifying mixing), and determine whether it increases

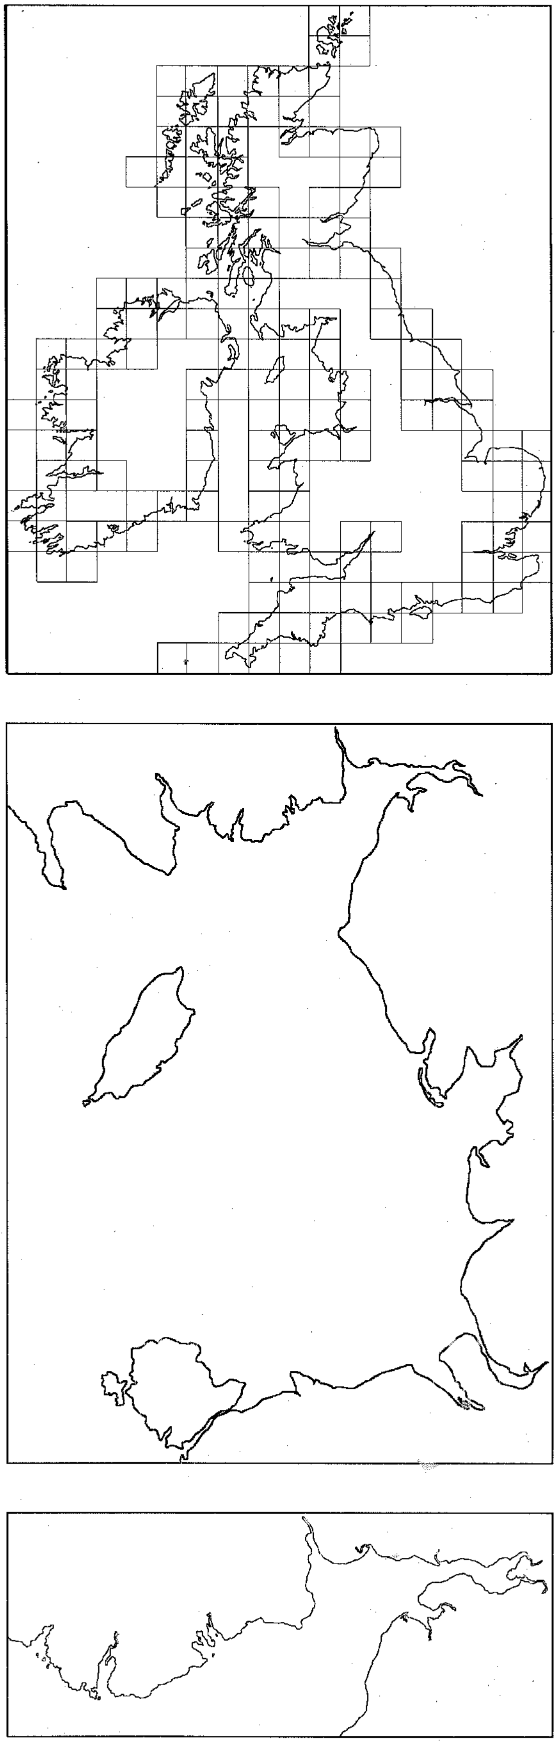

The coastline of Britain remains crinkly, whether you're looking at the entire thing (top), or just the part along the Irish Sea (middle), or just Solway Firth (bottom). If you tiled the entire map like a bathroom floor and then counted how many of those tiles covered some piece of the shoreline, you'd have a measure of how long it was. And if you made the tiles smaller and smaller, the measured length of the coastline would get longer and longer.

or decreases as the flow velocity increases?

A very exciting development took place about 20 years ago, when Benoit Mandelbrot (MS '48, Eng '49) coined the term "fractal" to describe the geometry of irregular objects, and suggested that special tools previously limited to a relatively arcane branch of mathematics could be applied to study such a geometry. His idea was that a fractal object looks equally complex, no matter at which scale you choose to examine it. You can zoom in on a small piece, or pull back and look at the whole thing, and it will look similarly complex. A coastline, for example, looks convoluted, whether you're looking at the entire coast of Britain, or just the part along the Irish Sea, or just Solway Firth, or just the harbor at Kirkcudbright, or just a piece of the rocky strand at low tide.

The ideas behind fractal mathematics had been put forth in bits and pieces by many people, but were first applied to natural phenomena by Lewis Fry Richardson, who, in a paper published posthumously in the early '60s, actually did a fractal analysis (he didn't call it that, of course) of Britain's coastline. Imagine that you have a map of Britain, and you're trying to describe how crinkly the coastline is. There are several possible approaches, but a good one is to draw a so-called bounding rectangle that just barely contains the coastline. You then fill the rectangle with contiguous, nonoverlapping tiles and count how many of them cover some segment, however small, of the coast, including all the nearby islands. As we make the tiles smaller and smaller, we need more and more of them to cover the

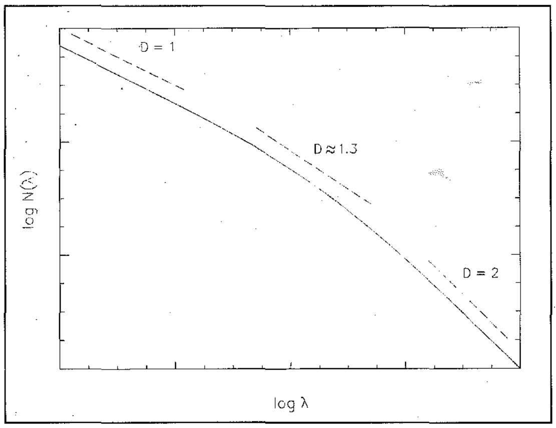

Plotting the logarithm of the number of tiles needed to cover the coastline, N(λ), versus the logarithm of the tile size, λ, should give a straight line whose slope, D, is the fractal dimension. But the line isn't exactly straight, so D goes from 2 when the tiles are large down to 1 when the tiles are tiny.

same coastline. We can plot the logarithm of the number of coastline-covering tiles, $N(\lambda)$, for a given tile size, λ, versus the logarithm of λ. According to Richardson and Mandelbrot, we should get a straight line with a negative slope, i.e.,

$$\log N(\lambda) = -D \log \lambda + \text{constant}$$

In this expression, the negative slope, D, is a constant, which means that the number of coastline-covering tiles is:

$$N(\lambda) \propto \lambda^{-D}$$

This is a power-law relation, because $N(\lambda)$ depends on a variable (λ), raised to a constant power ($-D$). The exponent, D, is called the fractal dimension. If the coastline is straight, then $D = 1$, corresponding to a one-dimensional object, i.e., a line. If the coastline is all scrunched up and visits nearly every point in the interior of the bounding rectangle that contains the tiles, then D is closer to 2, and the coastline approaches the solidity of a two-dimensional object. For the west coast of Britain, D is about 1.3, which means that the British coastline is not as baroque as, say, the fjords of Norway, for which D is about 1.5.

But, if you look at the log-log plot of $N(\lambda)$ versus λ more closely, you see that the line isn't exactly straight. At first, each time we cut the tile size in half, the number of tiles that contain a part of the coastline (however small) is squared. D still equals 2, in other words. This is because if you're cutting very large tiles, each smaller piece will still cover some stretch of shore. The subdivided covering tiles still fill the entire bounding rectangle, as for a two-dimensional object. This is called the embedding dimension, because our fractal island (of dimension 1.3) lives in a two-dimensional space. At the other end of the curve, for very small tiles, the plot's slope approaches 1. This is because you now need as many tiles as the "arc length" divided by the tile size, λ. The arc length is simply a number—the length of the coastline as it's drawn on the map. This geometric figure is represented down to a particular resolution, and thus has a corresponding arc length, even though the actual object—the coastline itself—effectively does not. And since the real object is approximated on the map by a line, which is a one-dimensional entity, we call this the topological dimension.

The fractalists say we understand the slope = 2 region and the slope = 1 region. So we'll ignore those extremes and study the region in between, where we hope D has some fixed intermediate value. For the most part, fractals discussed to date have been of this power-law variety and describe objects whose complexity is the same regardless of scale. Mandelbrot, in fact, adopted this attribute as the defining property of fractals.

Our first laser-induced fluorescence photos were taken in a small fish tank in 1976, as part of a research project with Rick Miake-Lye (BS '78), at about the same time that Mandelbrot was formulating his proposals. Mandelbrot visited Caltech in the late '70s, and I showed him our pictures. He was very excited and asked for a slide of the picture on page 22. He presented it at the physics colloquium he gave later that day as a clear demonstration of fractal behavior in turbulence. He even referred to it as such in the subsequent edition of his book on fractals. But we had already spent some time trying to do a fractal analysis of that picture, and tried again after he left, and could not get a power law.

I hesitated to publish this counter-result, however, because it was just one picture. Perhaps no power-law behavior emerged because our statistical sample simply wasn't big enough. Every picture is different—Heracleitus's insight: the river is not the same twice—so several pictures would have had to be analyzed, and average tile counts computed for each value of λ. However, our analysis methods back then were primitive and very time-consuming. I projected the slide on a lecture-room wall, and tried to measure $N(\lambda)$ by counting the number of times a λ-length string was needed to get from one end of a contour line to the other—the technique Mandelbrot had recommended.

At this point, I should explain what, exactly, we were measuring. What is the "coastline" of a turbulent jet? There are actually many lines one

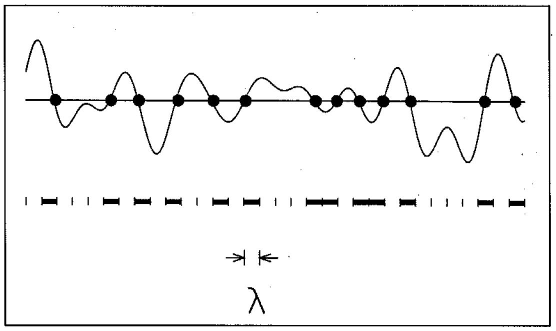

A quick lesson in chaos management: You start with some raw data (top)—a succession of random peaks and valleys—and you pick an elevation, such as the horizontal line. The set of all points at that elevation (the heavy dots) is called a "level set" and preserves the random qualities of the original data. You can now use line segments as tiles of size λ (bottom), and count how many segments it takes to cover all the data points in the level set.

can measure. Since the dye fluoresces in proportion to its concentration, given a laser sheet of uniform intensity all points of equal brightness in the image will correspond to the same concentration of jet fluid. Connecting all points of equal concentration on the image gives a set of contour lines—analogous to the elevation contours on a topographic map—called isoconcentration lines, or isocontours. Isocontours are also called "level sets," because every point in the set is at the same level—the same elevation in a topographic map, or, in this case, the same concentration. We approximated these isocontours photographically by varying the exposure while making a set of high-contrast prints (or slides) of each image.

So collecting enough data by hand to provide decent statistics would have been unthinkable, but getting the images into a computer for analysis was also difficult. We graduated from measuring isocontours off the wall to using a film scanner at JPL to get the image in digital form. However, even with the scanner, it was still so laborious to analyze a single picture that doing many of them was not in the cards. Also, despite the novelty of Mandelbrot's fractal ideas, it wasn't clear whether this approach was leading anywhere, and I didn't dare ask students to spend much time on it.

But even after we finally learned how to do a computerized analysis, we still weren't home free. There were problems with the way we were estimating the fractal dimension. We called our method the "stretched-string" algorithm because that was how we had done it on the wall. Just as a rock climber negotiates a tricky face by securing a safety rope to pitons at closely spaced intervals, the computer belayed an imaginary string from a point on the isocontour that had been reached by stretching the string from the previous point reached in the same way, and continued to do so until a complete circuit of the contour had been made. Unfortunately, the coverage counts thus derived were not unique, because there were often many choices of where to place the other end of the string, which led to different counts. And the stretched-string algorithm could cheerfully yield fractal dimensions that exceeded 2 for two-dimensional data, which we regarded as nonsensical.

We really had to wait until we could use digital imaging, initially in the form of a linear CCD array oriented perpendicularly to the jet's axis, to acquire good data in bulk. In those days, you couldn't buy a digital camera. You beat the bushes until you found a manufacturer who'd sell or give you a noisy CCD chip. Then you built all the electronics around it, converted the voltages to numbers with expensive and difficult-to-use analog-to-digital converters, and recorded the numbers any way you could. We had to develop, from scratch, the electronics to acquire and store digital data for subsequent computer processing; a technology we've been refining ever since. Our setup was first used for fractals in 1985, as part of an Aeronautics 104 class project by grad students Sheldon Green (MS '85, PhD '88) and Giancarlo Losi (MS '85, PhD '90). Able now to record digital records in the computer from the start, we could gather enough data to obtain reliable statistics. Our results continued to be inconclusive, however, and I still did not wish to venture a publication. The stretched-string algorithm remained troublesome, among other issues.

By the late '80s, we had done enough thinking and doodling that we decided we could design an experiment to settle at least some of these issues once and for all. Grad student Paul Miller (MS '87, PhD '91) and I began using laser-induced fluorescence to make long, digital records of the jet fluid's concentration, as a function of time, at a fixed point on the jet's axis. These plots looked like a slice through a very jagged mountain range—peaks and valleys in succession. We then selected all the points in time where the jet-fluid concentration crossed a fixed threshold—the one-dimensional analog of an isocontour—and "tiled" them with line segments, again counting the number of (one-dimensional) tiles required to cover the threshold-crossings as a function of tile size. And, having abandoned the stretched-string algorithm in favor of contiguous, nonoverlapping tiles, we also revisited the linear-array data from

Below: This plot of fractal dimension, D, versus the logarithm of the tile size, λ, was calculated from the one-dimensional temporal data. If the data had a power-law fractal region, it would appear as a horizontal plateau (or at least a kink) in the curve. After Dimotakis, *Nonlinear Science Today*, 1991.

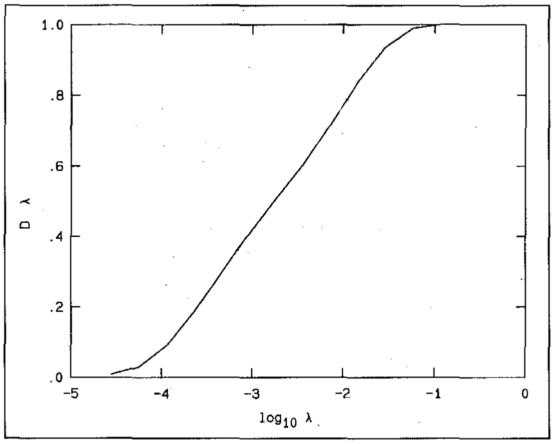

Below: Noisy data can give your level set a lot of spurious members. In this example, the real signal (heavy line) crosses the chosen elevation (dashed line) twice and twice only. But when this data is cloaked in noise (light line), scores of new points appear. To make matters worse, the new points don't just surround the real level crossings, but can also appear where the original signal approaches, but does not cross, the chosen elevation. Fortunately, during World War II, Norbert Wiener proposed a very clever scheme (classified at the time, declassified later, but not well understood until later still, when it was described by someone else) for recovering a signal from noisy data and computing its level crossings quite reliably. After Miller and Dimotakis, *Physics of Fluids A*, 1991.

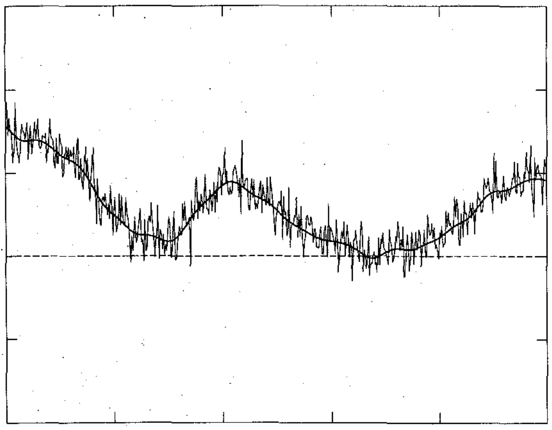

Sheldon and Giancarlo's Ae 104 experiment—still luxuriously spinning on a disk in our computer network—and calculated one- and two-dimensional tiling statistics for them too.

Our new, one-dimensional, temporal data produced statistically reliable tile-coverage counts with a logarithmic slope that smoothly increased from 0 to 1. Similarly, the space-time data from the Ae 104 experiment yielded a slope that smoothly increased from 1 to 2. In neither case did the fractal dimension (the slope, D) pause at any particular value. No power-law region appeared. I should note that, by then, many investigators had reported having found a constant fractal dimension in all kinds of flows, and the same fractal dimension to boot. So ours was a very controversial result. We were comfortable, however, with the care we had expended in our analyses to eliminate the influence of noise on the data and to understand the subtleties of the various algorithms from determining N(λ). And as our earlier data had prepared us for a curve, we did not attempt to fit a straight line to the log N(λ) versus log λ plot, which would of course have given us a constant dimension. It took 14 months to get the paper, which was eventually published in 1991, through the reviewing cycles: answering all the reviewers' queries and objections, documenting that we'd addressed and eliminated all sources of error, and, incidentally, doubling the paper's length in the process.

There could no longer be any question, at least in our minds, that turbulence generated level sets of smoothly varying fractal dimension—i.e., a continuous dependence of D on tile size—whose values were bounded by the topological dimension from below and the embedding dimension from above. There was no choice but to extend Mandelbrot's inspired proposals—insistence on uniform geometric complexity, regardless of scale, had to be abandoned if fractals were going to be useful in describing turbulence. Mandelbrot's original (power-law) fractals had to be regarded as an important special case of a broader mathematical framework, but a special case nevertheless.

Our paper caused a lot of confusion. Was Pasadena turbulence again different, as had been alleged in the early '70s, when large-scale behavior was discovered? Was it because we had primarily relied on temporal data, even though the scant spatial data we had analyzed were also in accord? So Haris Catrakis (BS '91, MS '91, PhD '96) and I forged ahead to see if our conclusions would survive the test of time and the results of improved experiments.

By then, CCD technology had progressed to

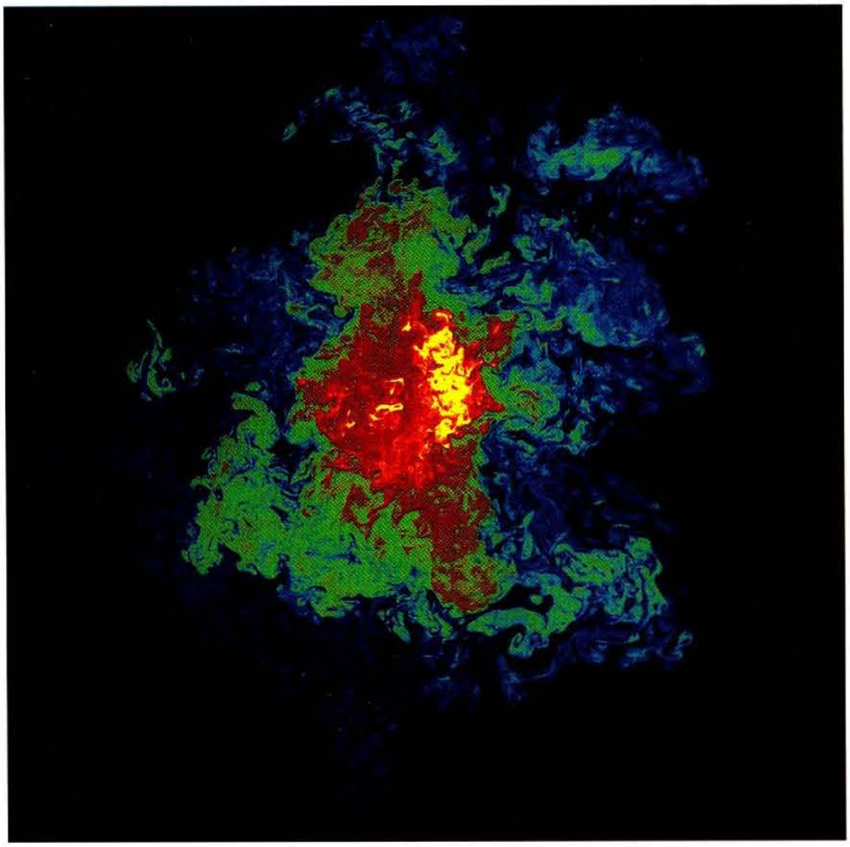

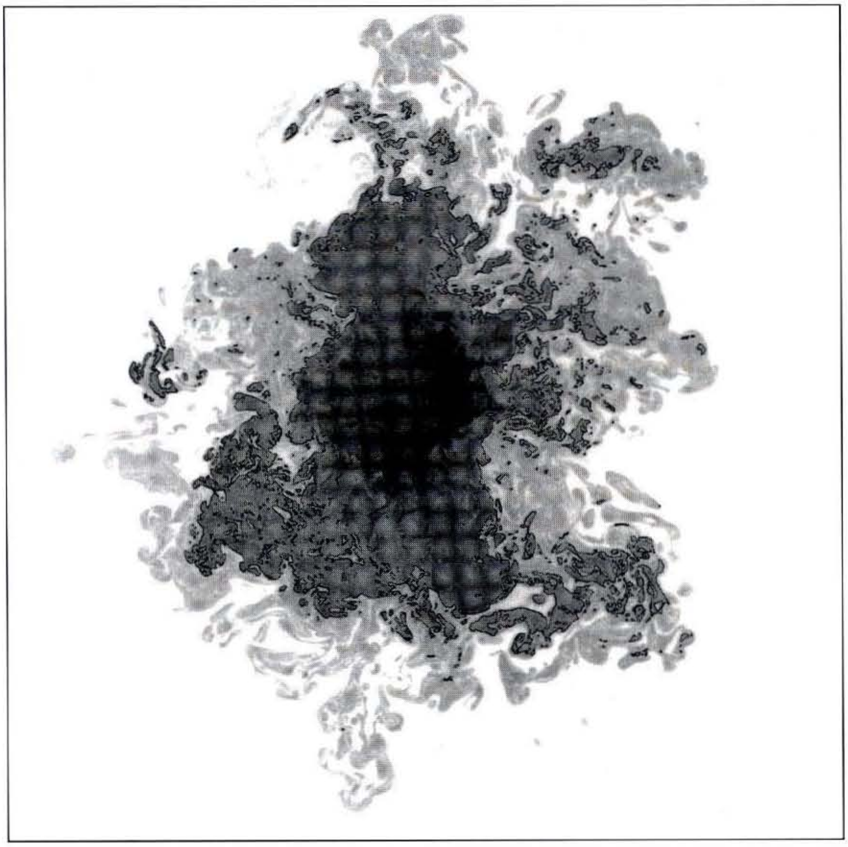

Top: A two-dimensional slice of a turbulent jet, taken with the apparatus described on page 25. The colors represent the jet fluid's concentration. Blue is the most dilute, while green, orange, red, and yellow are progressively more concentrated. Bottom: The heavy line is an isocontour calculated from the same image. Several "islands"—secondary isocontours outside the main one—and "lakes"—secondary isocontours inside the main one—can be seen. From Catrakis and Dimotakis, *Journal of Fluid Mechanics*, 1996.

where we could record two-dimensional spatial data (i.e., images) that were almost as good as previous point measurements. (A typical image appears at left.) These two-dimensional, laser-induced fluorescence slices, oriented perpendicular to the jet axis, allowed us to give the irregular geometry of two-dimensional isocontours the same rigorous statistical treatment that was only possible in one dimension before. And more powerful data-acquisition, storage, and processing systems allowed us to analyze several isocontours for each Reynolds number from images recorded over a range of Reynolds numbers.

The Reynolds number measures the relative importance of viscous diffusion in a flow. If the Reynolds number is low, viscous effects are important, and viscous damping prevents flow fluctuations and turbulence. For example, flowing honey has a very low Reynolds number and is hard to make turbulent. But the Reynolds number increases with flow speed, so water, for example, can easily be at a high enough Reynolds number to be turbulent. If you fill your bathroom sink very slowly, the water necks down as it leaves the faucet and you get a nice, smooth, laminar flow. If you turn the water up, the flow suddenly becomes unsteady. The Reynolds number has crossed a critical value above which the small, inevitable fluctuations in the flow in the pipe supplying the faucet are amplified and sustained by the flow's kinetic energy. Viscous damping is no longer sufficient to keep the flow calm. This doesn't mean that flows above some critical Reynolds number are always turbulent, only that turbulence requires a minimum

Reynolds number to be sustained.

Haris has made many important contributions in the course of his PhD research, but to make a long story short, the new images yielded the same fractal behavior as our previous data—a continuously varying D that spanned its possible range of values (in this case, from 1 to 2). This held true throughout the Reynolds-number and isocontour ranges we investigated. In summary, our experiments suggest that turbulence generates structures that are not equally complex at all scales. Instead, turbulence is more complex at larger scales (larger D), and less complex at smaller scales (smaller D), with a continuously variable D(λ) required to describe it. Haris has also found reports of such geometric behavior from other fields, for example in the analysis of alveolar tissue from rabbit lungs (the alveoli are the little sacs where gases exchange into and out of the bloodstream, so a high surface-to-volume ratio is obviously desirable), and in cloud-shape distributions. We've decided to call this variable D(λ) the "scale-dependent fractal dimension," and geometrical figures that display it "scale-dependent fractals," to distinguish them from the original "power-law" fractals.

Turbulence takes a large eddy, strains it, splits it into smaller eddies, and then again into smaller ones yet. It also merges eddies to make larger ones, and merges those again to make bigger ones yet, producing a very rich distribution of shapes and sizes.

So what have we learned from all this? We've recently realized that if you know the D(λ) curve, you can work backward and compute the distribution of spatial scales in the flow. In particular, you can say what the distribution of nearest distances to an isocontour is. Paul Miller had discovered this, in an inverse way, in the one-dimensional temporal data—a result we included in the 1991 paper. He used a random-number generator to sprinkle points on a line with a statistical distribution of his choosing to see what distribution gave a D(λ) that looked like the turbulent-jet data. He found that if the point spacings were log-normally distributed (that is, if the logarithms of the distances between successive pairs of points had a Gaussian distribution—the classic bell-shaped curve) then a fractal analysis of those spacings gave a D(λ) that very closely matched our one-dimensional data.

Haris and I have extended that to higher dimensions. Of course, what you mean by "spacing" in two dimensions must be defined, because it can be measured in many ways. In this context, we measure it as the (distribution of) sizes of the largest tiles, randomly placed, that do not touch the level set—the isocontour—at any point. Similarly, in three dimensions, one would be placing a box so it doesn't touch an isosurface, the level set in that case. We've also studied the size distribution of isocontour islands and lakes in our two-dimensional data. In this case, since we're talking about distinct objects, we can define "sizes" more naturally as the square roots of the individual areas—if the island or lake were square, its size would be the length of its edge. We found that this size distribution is also log-normal. Nature often generates log-normal distributions whenever it merges or subdivides things, and turbulence, in a way, makes islands and lakes by the fusion and fission of eddies. Turbulence takes a large eddy, strains it, splits it into smaller eddies, and then again into smaller ones yet. It also merges eddies to make larger ones, and merges those again to make bigger ones yet, producing a very rich distribution of shapes and sizes. The largest eddies are bounded by the full spatial extent of the turbulent region, while the smallest ones have sizes dictated by viscosity and, in the case of concentration data, diffusion.

What does this mean in the real world? These geometrical properties are important in describing the non-premixed combustion of hydrocarbons, for example. If we ignite aviation fuel in a jet engine, the burning rate is typically not limited by the rate of the chemical reaction of fuel and oxygen in air. The limiting factor is the rate at which the fuel mixes with the air and finds the oxygen it needs to burn, which is almost entirely determined by the characteristics of the turbulence that brings the two reactants together. The burning is confined to the unsteady, three-dimensional surface on which the mixture of fuel and oxidizer is at the stoichiometric ratio—the ratio at which the two will completely consume each other, with no leftover fuel or oxygen. This constant-concentration surface is also a level set, like the ones we've been studying in our water jets. Knowing the statistics of the distance distribution from a point to that isosurface tells us how far the fuel has to diffuse to meet the oxygen, or vice versa. Premixed combustion, as in an internal-combustion engine—in which fuel and air are mixed ahead of time and ignited later on—occurs on an equally complex combustion surface (of more-or-less constant temperature), which can also be treated as a level set.

One has to be careful when extrapolating our water-tank results to air, however, because of the differing diffusion properties of the two fluids. From a molecular viewpoint, a gas is mostly empty space. If you mix two gases together, the molecules go zipping past one another and carry their momentum some distance between collisions—there's a long mean free path. The diffusivity of mass (the molecules) and the diffusivity of their momentum is very nearly the same. But in a liquid, the molecules are, effectively, in contact with one another. There's no such thing as free flight. You can transport momentum without transporting the molecules themselves, as happens in those toys with a row of hanging steel balls—you hit the ball at one end and the ball at the other end takes off. Momentum has been transported with almost no transport of mass. To transport a molecule some distance in a liquid, many other molecules have to get out of the way. This doesn't happen easily, so the diffusion of molecules in a liquid is about a thousand times slower (slower still for a large molecule) than the diffusion of momentum. For a chemical reaction to proceed, it is individual molecules that must mix, not their momentum—we have to get the acid to meet the base, or the fuel to meet the oxygen.

In the midst of all this fuss about such images, we should not forget that turbulence is not two-dimensional. It's a three-dimensional process that evolves as a function of time, so it's really a four-dimensional phenomenon. To look at this, we would need to capture three-dimensional data—ideally, as a function of time. One way to do this is to slice the flow very thinly, very quickly—before it changes—and assemble the slices into a three-dimensional image. Kelley Scott (BS '84) tried to do this as a SURF (Summer Undergraduate Research Fellowship) project over ten years ago, but the technology was not there. Werner Dahm (PhD '85) finally did so at the University of Michigan after leaving here, but he used low spatial resolution and low flow speeds (low Reynolds numbers); it's easy to appreciate that the data rates required for decent spatial and temporal resolution quickly push this approach beyond the reach of present-day technology.

But in the last few months, we've managed an early peek at what such data might look like. Dan Lang (BS '76, MS '77, PhD '85), the group's expert in electronics and many other things, has integrated a new CCD into a digital camera and high-speed data-acquisition system that records at a resolution of 1,024-x-1,024 pixels with an excellent signal-to-noise ratio. Designed by Jim Janesick, Andy Collins, and other digital-imaging wizards at JPL, this CCD is a spare late-technology chip whose sibling will be on the upcoming Cassini mission to Saturn. Haris, Dan, and I have used the system to assemble some of the first three-dimensional, high-resolution data sets ever taken of turbulence. We recorded successive image slices in a flow for which a rate of one frame per second was almost fast enough. (By contrast, TV cameras, which have much lower resolution and signal-to-noise ratio, generate 30 frames per second.) We got our first peek at the data by computing some isosurfaces from a small portion of that space-time data, using about 60,000 polygonal facets to render each isosurface. Now David Laidlaw (MS '92, PhD '95), a postdoc in computer science, has rendered a much larger portion of the image data using up to 3,000,000 polygonal facets per surface. One of his images appears on the cover of this magazine.

And we continue to push forward—we can now achieve a framing rate 10 times higher than that first effort, yet with the same spatial resolution and signal-to-noise ratio. Or, by "binning," i.e., summing 2-x-2-pixel regions on the CCD before reading them out, we can read 512-x-512 images at 20 frames per second. To study fully developed turbulent flow, however, we estimate we would need a resolution of no less than 1,000-x-1,000 pixels per frame, read out at something like 1,000 frames per second, at Reynolds numbers large enough for bona fide turbulence. This translates into a minimum of a billion pixels per second, or, if we digitize the data stream at 12 bits per pixel, a data rate of 1.5 gigabytes per second. Dan is currently developing this kilo-frame-per-second system.

We're also improving related laser-imaging technology so that we can study gas-phase turbulent mixing. To date, we've used water as our fluid because it's a thousand times denser than air, so we can get much more fluorescence signal per unit volume. We hope that, in a few years, we'll be able to compensate for the three-order-of-magnitude signal loss we'd encounter if we probed gas-phase flows. In fact, we're almost there now, after recent advances in gas-phase imaging made by Senior Research Fellow Dominique Fourguette. Since the same flow equations describe gases and liquids, comparing liquid- and gas-phase flows directly will be tremendously valuable. The only difference would be in the diffusivity of mass, so we'll be able to take a very complex phenomenon and change only one dial, leaving everything else the same. Any differences will then be attributable to turning that one dial. That's extremely valuable in science, so we're very excited by the prospect.

And finally, there's the issue of supersonic turbulence, whose nature is largely terra incognita. We need to make progress there if we're to fly faster than we do now. The usual theories of turbulence, even on the level of von Kármán, Taylor, and Kolmogorov, don't apply beyond the speed of sound, because the basic assumptions on which they rely are no longer valid. To study supersonic turbulence, one needs some way of recording the instantaneous velocity field—the simultaneous velocity of every point in the flow at one instant in time. This has not been possible, to date. A

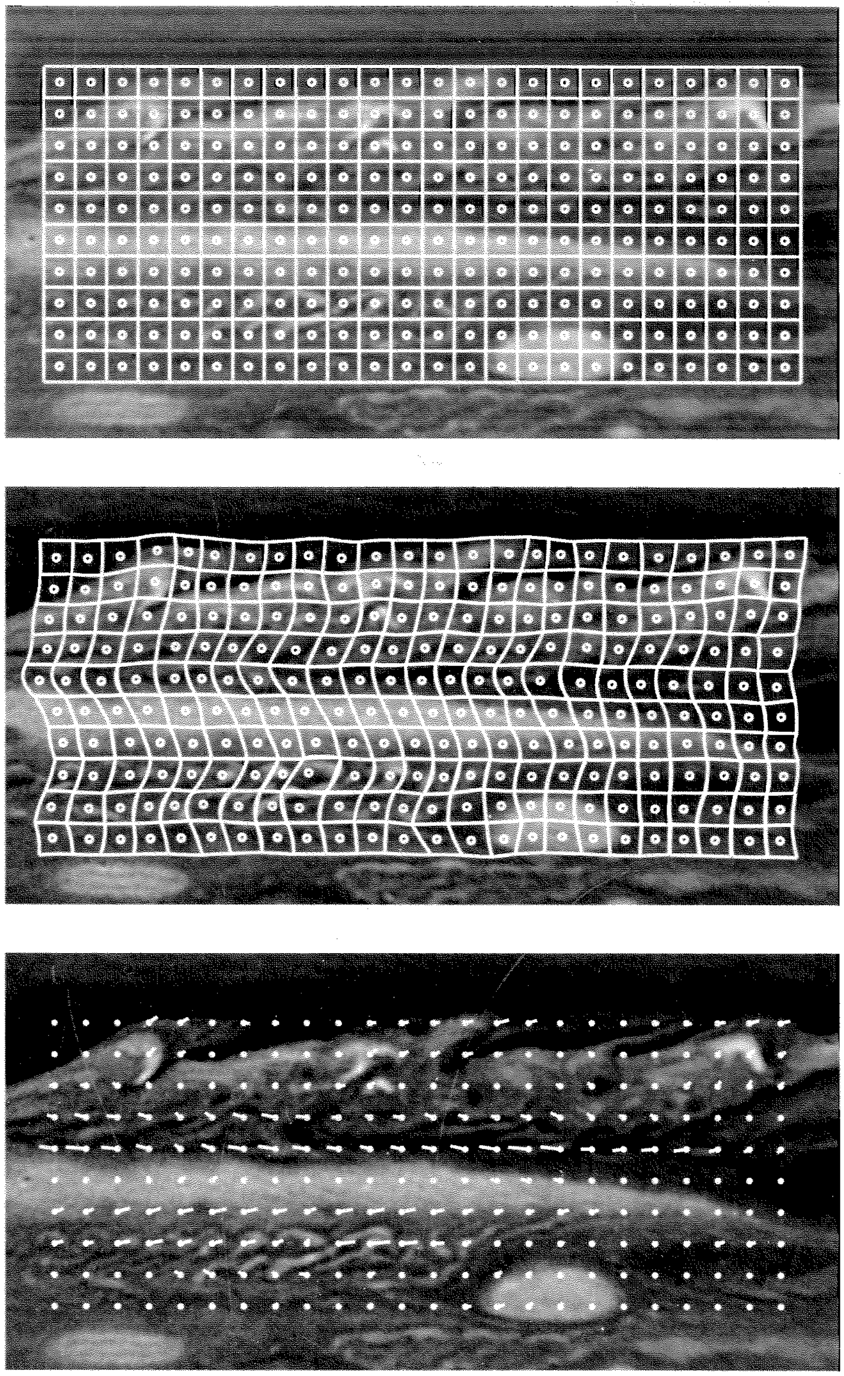

Mapping Jovian wind speeds by ICV. A grid superimposed on Voyager 2 images of part of Jupiter's southern hemisphere (top) is distorted one rotation later (middle), with the motion of each square's central dot indicated by a line segment (bottom). After Tokumaru and Dimotakis, *Experiments in Fluids*, 1995.

few years ago, Phil Tokumaru (MS '86, PhD '91) and I took some first steps in that direction by developing a method to deduce the velocity field from flow images recorded in quick succession. We call this method Image Correlation Velocimetry (ICV for short), and it looks for the mapping, or displacement field, that turns one image into the next one. The velocity field is then the displacement field divided by the time between successive images. Grad student Galen Gornowicz is now helping improve the method, which we've tested on a few toy flows: mapping the velocity field around an accelerating wing section whose performance is known, for example, and measuring wind speeds on Jupiter from a pair of images obtained from JPL. Our method works for the simple laboratory flows we've tested it on thus far, and we've been told by our friends in planetary science that our Jovian wind speeds are right. To use this method, however, one needs high-signal-to-noise-ratio images that are close enough in time to be reasonably well correlated. To do this in a supersonic flow, we have to solve the gas-phase-imaging problems mentioned above *and* record images as close as a few microseconds apart. So our JPL friends have designed and helped us fabricate a CCD that can record two high-signal-to-noise-ratio images with the requisite microsecond-scale spacing, which can then be read out and digitized at ordinary framing rates. We call this device the "Mach-CCD," after the flow speeds it is intended to decipher. We're bench testing it now.

There's a need for improved digital imaging in many fields. Chris Martin, professor of physics, and astronomy grad student Brian Kern have built a system that records optical phase-front distortions—the twinkle in starlight. Dan and I used our Cassini CCD camera system, in parallel with their system, to record 10- and 20-frame-per-second sequences of high-quality, short-exposure images on the 200-inch Hale Telescope at Palomar. In this collaboration, the hope is to understand how atmospheric turbulence causes optical distortion and test ideas for correcting it. Excited by these prospects, we've decided to up the specs of the kiloframe-per-second system that Dan is developing so that it will be able to run at that data rate for longer times, for both turbulence and astronomical applications. Scott Fraser, the Rosen Professor of Biology, and Professor of Physics Jerry Pine are interested in applying this capability to biological imaging, and we look forward to working with them. In these and other areas, the ability to follow high-resolution, high-signal-to-noise, two-dimensional data in millisecond or smaller time intervals would put

A fireball that the eye perceives as volume-filling in fact consists of complicated three-dimensional isosurfaces in constant random motion. But you've got to look fast to freeze them—this is a 1/1000th-second exposure.

So that's where we are. As with much of science, progress often awaits the development of a new technology, which allows a better view of nature, which improves our understanding, which begets new questions, which in turn await a new technology, which …

many important phenomena within direct reach of quantitative scrutiny.

So that's where we are. As with much of science, progress often awaits the development of a new technology, which allows a better view of nature, which improves our understanding, which begets new questions, which in turn await a new technology, which … So it is with our quest for a better description of turbulence. In this round, first there was the excitement of fractals, because they promised a description of complex geometry. Then came the disappointment when we realized that we couldn't test the idea because we couldn't record and analyze adequate data to check it. Then the technology arose to do so, followed by the disappointment when we found that turbulence wasn't a power-law fractal. And now there's the excitement of realizing that the mathematics of fractals can be extended to accommodate the behavior that our experiments have revealed. Fractal language gives us the proper tools to talk about turbulence, if you're not bent on fitting straight lines to things that are curved. The new scale-dependent fractal dimension contains a lot more information and is better able to describe turbulent mixing and combustion. But valuable as that is, it isn't enough. We need local velocity-field information along with the isosurface-geometry data, a need that has spurred the development of Image Correlation Velocimetry. Soon we'll be able to derive the local velocity field from the same set of images that will give us the isoscalar geometry—two birds with one stone!

Every now and then we make a little bit of progress in understanding turbulence, and then there's a long wait until the next step. We think this research will lead to a big jump in our understanding, and we're excited. This jump took the confluence of a new idea—Mandelbrot's proposal to apply the notion of fractals to turbulence—and two technologies—the advent of digital imaging, which generates large amounts of high-quality data directly in computer-manipulable form, and an astonishing increase in computing power. All three components are advancing today, and we're relying on their continued progress for the next jump. Will it be the last? Victory over turbulence has been declared on a semiregular basis, every time based on different means. And every time, turbulence has risen, undefeated, to mock us. The ancient Greek gods may well have left this piece of the classical world as their legacy to remind us of the perils and pitfalls of hubris. ☐

Paul E. Dimotakis arrived at Caltech as a freshman in 1964, and has been here ever since, earning his BS in physics in 1968, his MS in nuclear engineering in 1969, and his PhD in applied physics in 1973. He then rose through the professorial ranks, becoming the Northrup Professor of Aeronautics and Professor of Applied Physics in 1995. Besides turbulence, his research encompasses such other fluid-mechanical phenomena as cavitation and gas dynamics.

As a consultant, he has participated in the development of pilotless drone aircraft, high-power chemical lasers, the stealth fighter, the space shuttle, sealed computer hard disks, and helped with the fluid-mechanics design for the "Leap-Frog Fountain" at Disney's Epcot Center in Florida. An avid sailor, he was a member of the America^3 team and contributed to the sail design for their successful defense of the Americas Cup in 1992.

The work on turbulence, fractals, and digital imaging described in this article has been supported over the years by the Air Force Office of Scientific Research.

What Is Life?
A Closer Look

by Robert L. Sinsheimer

We can actually classify and enumerate the components of the machine, gene by gene, and discern their interrelated functional organization.

(On Seminar Day in April 1967, Robert L. Sinsheimer, then professor of biophysics at Caltech, delivered a talk entitled "What Is Life?" (published in the Caltech Quarterly, *a short-lived spin-off of* E&S*). At the dawn of the era of space exploration, many people were wondering how life was to be recognized if indeed it were found elsewhere than on Earth. Sinsheimer gave the biologist's definition. Now, 30 years later, he sticks by his answer, and fills in some gaps with data from current biological research.)*

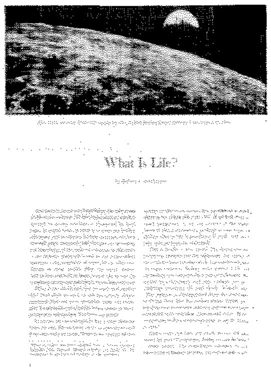
What Is Life?

The black box (or, more accurately, the colorless capsule) that we call a living cell is a specialized, intricate, highly evolved machine. Nearly three decades ago we were first able to describe qualitatively the essential elements of this self-perpetuating machine (*Caltech Quarterly*, Summer 1967) at the level of its molecular organization. At that time our understanding of the machinery of life was newly emergent, still quite incomplete and porous, but sufficient to replace the older vague theories and speculations.

Back in 1967 we were asking: What is the essence of this quality, "life"? Biology is the science of life, but no biology textbooks could provide a definition of what life is. There were two explanations advanced to account for the properties of living beings: one postulated that the substance of living matter was intrinsically different from that of nonliving matter; the other, that the properties of life were solely a consequence of an unusual organization of ordinary matter in living creatures. The latter argument had been hampered by the inability to describe that organization, but by 1967 biologists were able to describe the complex organization of the cell, the basic unit of life, in physical and chemical terms. I wrote then that "what distinguishes life is the presence of a persistent degree of structural complexity at a molecular level—unknown outside the sphere of life." Although higher organisms are composed of millions or billions of cells organized into a cooperative, interacting whole, a single cell can perform all the essential living functions.

Biology has progressed. Within the past year, biologists have anlyzed the complete hereditary instructions—the complete DNA sequences—of two species of microorganisms, *Haemophilus influenzae* Rd and *Mycoplasma genitalium*. Using these two sequenced species (and more are on the way), we can now describe the cell in much more detail. We can actually classify and enumerate the components of the machine, gene by gene, and discern their interrelated functional organization.

But first let me return to sum up the essential features of this specialized molecular organization that I described in 1967:

1. A flexible, self-made bounding membrane, which can be replaced if it is torn and can be enlarged as the cell grows, which defines the ordered, integrated space of life and controls ingress from and egress to the external world. It also provides a two-dimensional surface on which agents involved in particular reaction sequences can arrange themselves.

2. Coordinated groups of specific catalysts (usually proteins) which, by accelerating a myriad

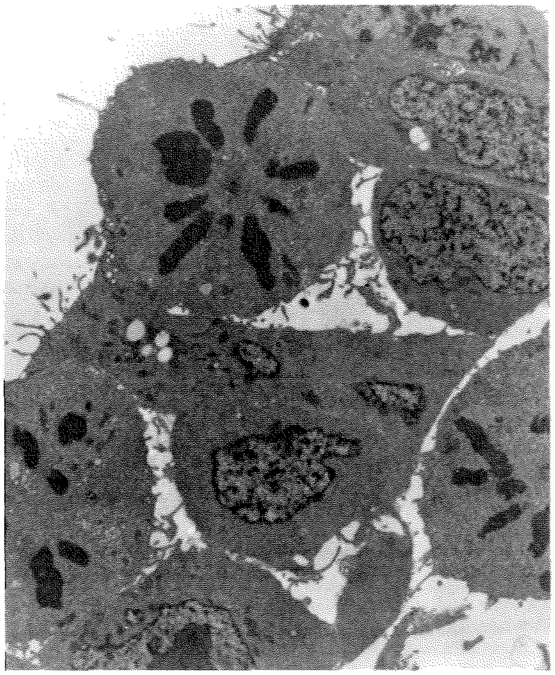

One of the essential features of a living cell is its ability to produce new cells by division and to convey hereditary instructions. This micrograph shows cells in a mouse embryo at 8.5 days of gestation, a time when rapid cell division is taking place. The three round cells with a rosette-like pattern (two of them cut off on the sides) are in different stages of mitosis (cell division). In less than an hour each will have divided into two daughter cells. The dark, sausage-shaped bodies arranged in the "rosette" are the chromosomes into which each cell's DNA is packaged for equal distribution to the daughter cells. Apprixmate magnification: 3000x. (Jean-Paul Revel, Stefan Offermanns, Mel Simon)

In 1967 this qualitative functional description of the elements essential to the organization of a living cell was the best we could do. But now we have, as illustrations, these two completely sequenced microorganisms.

of specific reactions in specific directions, determine the pathways of biosynthesis and biodegradation within the cell.

3. Molecular machinery to convert an external source of energy (either solar energy or the energy stored in the chemical bonds of nutrients) into forms suitable to drive reactions within the cells.

4. A hereditary system of information storage (DNA), which provides in each generation the coded formulas for the structures of the various catalysts and structural members of the cell; also the means to repair the DNA structures if damaged, utilizing the redundancy of information in the double helix; and molecular machinery to convert (by transcription and translation) the inherited coded information in DNA into the specific molecules and structures of the cell. The cell must also have the capability to produce new cells by division, so as to multiply and to permit the trial of modifications of the hereditary instructions.

5. Such division must accurately provide each daughter cell with a full set of the hereditary instructions (occasionally modified by mutation), together with a sufficient endowment of machinery to transcribe and translate these instructions. It must also provide a supply of usable energy adequate to permit each daughter cell to flourish.

6. Because this organization must be flexible and adaptive to changing environments, it has interlocking control systems that automatically regulate the varied functions to keep the cell on course or, more accurately, on its many courses. In addition, the living cell has a capability beyond and distinct from that of the conventional cybernetic system, which can only adjust the rates of action of its component parts to maintain a level of function or output. Since a cell makes all of its component parts, it can also adapt the number of component parts to the task—creating more if needed—and it can also salvage and repair damaged parts if the damage is not too great.

In 1967 this qualitative functional description of the elements essential to the organization of a living cell was the best we could do. But now we have, as illustrations, these two completely sequenced microorganisms.

Haemophilus influenzae Rd is a small, nonmotile bacterium, a pathogen that causes respiratory infections in humans; one strain of it causes meningitis. *H. influenzae* is able to exist in a relatively simple medium and has a genome of 1,830,177 nucleotide pairs (DNA's basic molecular units, whose signifying letters, A, C, G, and T—for adenine, cytosine, guanine, and thymine—make up the genetic code), containing 1,743 regions that are equivalent to genes and that code for specific protein structures. (This work was published by Robert D. Fleischman et al. in *Science*, vol. 269, No. 5223, pp. 496–512, 1995.)

The second species, *Mycoplasma genitalium,* is a tiny, parasitic bacterium, lacking a rigid cell wall and found in association with ciliated epithelial cells of primate and human genital and respiratory tracts. *M. genitalium* has a "stripped-down" genome of 580,070 nucleotide pairs containing only 470 protein-coding regions—likely a near-minimal content for a self-replicating organism, albeit one dependent upon a host for many complex nutrients and environmental stability (published by Claire M. Fraser et al. in *Science*, vol. 270, No. 5235, pp. 397–403, 1995).

Most, although not all, of the defined gene sequences can be related to proteins of known or likely function, and I thought it would make an interesting exercise to try to fit them into the categories I had described in 1967. The result is shown at left. This doesn't cover the complete organisms: in *Haemophilus*, probable functions have been ascribed to 1,007 of the 1,743; in *Mycoplasma,* to 374 of the 470. (Many of the products of the currently unknown genes are likely to be involved in interactions, pathogenic or otherwise, with host organisms.) Many of the known genes can actually be clustered into groups with more narrowly defined missions, but for the sake of this illustration I am grouping them in the more general categories. Assignments of genes to these categories are of necessity incomplete, and in some instances somewhat arbitrary, but I hope they are consistent.

One important additional function that was not considered in 1967 is the provision of means for defense—for survival in a dangerous world, for potential encounters with toxic chemicals, invading viruses, and other cell-eating cells. Both species possess genes for at least rudimentary defense and counter-force measures.

A comparison of the gene distributions for the two microbial species indicates the economies made possible by the parasitic lifestyle of *Mycoplasma genitalium.* Reliant upon host cells for many complex nutrients, this species has a greatly reduced biosynthetic (production) machinery; stealing energy from the host permits much simplification of the energy-producing apparatus; adaptation to life in the relatively controlled environment of a host cell permits great simplification of the means of border control and of the adaptive systems for regulation, control, and repair. Defense agents are trimmed deeply, and even the machinery for information storage and expression is reduced to an essential minimum.

As we unveil the "secrets" of life even in these simple cells and thus define and measure their complexity, this understanding only deepens our marvel at their very existence.

		Haemophilus influenzae	*Mycoplasma genitalium*
1.	cellular envelope, surface structures transport	223 (12.5%)	57 (12.1%)
2.	metabolic reactions	246 (14.1%)	38 (8.1%)
3.	provision of energy	105 (6.0%)	31 (6.6%)
4.	DNA replication, cell division	111 (6.4%)	37 (7.9%)
5.	information transfer	147 (8.4%)	101 (21.5%)
6.	regulation, control, repair	94 (3.6%)	31 (6.6%)
7.	survival, defense	44 (2.5%)	3 (0.64%)

An interesting comparison emerges from the allocation of the genes of two newly sequenced microorganisms into categories devised 30 years ago for defining a living cell. At right, Robert Sinsheimer during his Caltech years.

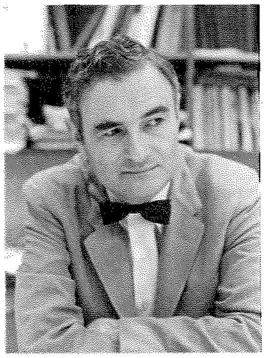

Robert Sinsheimer was professor of biophysics at Caltech from 1957 and chairman of the Division of Biology from 1968 until leaving to become chancellor of UC Santa Cruz in 1977. He is currently an emeritus professor in the Department of Biological Sciences at UC Santa Barbara.

Books

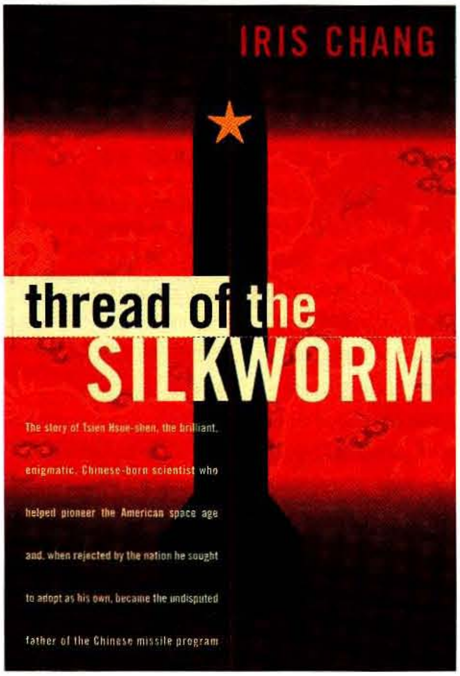

Thread of the Silkworm

Iris Chang
BasicBooks, New York; 1995

by Judith R. Goodstein

The period we know as the McCarthy era touched many aspects of American life, not the least of which was science and the scientific community. Weighed down by Cold War fears, Americans took refuge in the anti-Communist movement that spread across the country after World War II. In this atmosphere, spurred on by alumni, trustees, and congressmen, college administrators often dealt harshly with faculty whose politics became a public issue. Professors accused of being Communists lost research grants, passports, and travel privileges, and were forced to appear before government committees investigating subversive activities. Some university professors lost their jobs; others had worse things befall them.

At Caltech, the Red scare also played havoc with individual lives. Jet Propulsion Laboratory cofounder Frank Malina was one example. In 1946 Malina suddenly took a leave of absence from the lab, joined UNESCO in Paris, and remained overseas. Sidney Weinbaum, a gifted mathematician, musician, and chess player, was less fortunate. Weinbaum received his PhD at Caltech, worked in Linus Pauling's laboratory for 12 years, and had just started working for Bendix Aviation in 1941 when he was first accused of membership in the Communist Party. To his own surprise, Weinbaum was cleared for secret work from 1941 to 1949, by which time he had a permanent job at JPL. His real troubles started in 1949, when he got a telephone call from JPL asking him to fill out a new security form, which included listing all the organizations he had ever belonged to. The FBI wanted him to supply the names of people he had known and associated with in the late 1930s; when he refused, he was tried and convicted of three counts of perjury and one of fraud in denying he was a member of the Communist Party. Then 52, Weinbaum spent four years in jail; when he left prison in 1953, his scientific career was finished.

Tsien Hsue-shen, a distinguished scientist with impeccable credentials in aeronautics and jet propulsion and a friend of Malina and Weinbaum, was another casualty of the Red-baiting fifties. *Thread of the Silkworm*, by Iris Chang, a freelance writer, is the first full-scale biography of Tsien's life. In her well-documented book, Chang draws on all her skills as a reporter to flesh out his story, using interviews with Tsien's former students, classmates, and colleagues in this country and in China, as well as FBI reports and Army Intelligence records in U.S. Customs files and the National Archives, recent articles about Tsien published in China, and documents and letters in presidential libraries and university archives, including Caltech's.

The starting point for anyone studying Tsien is the writings of Milton Viorst, and Chang relies heavily on them to anchor her story. Tsien himself, who lives in Beijing, declined to be interviewed by the author. "One should never write a book until he is on his deathbed, because he won't live to regret

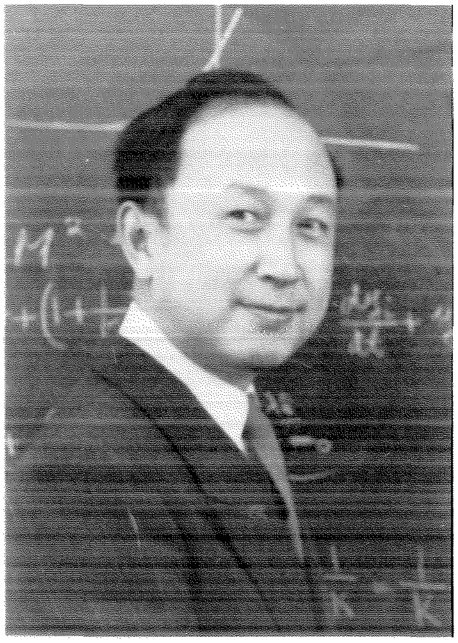

Tsien Hsue-shen at Caltech in the early 1950s.

it," Tsien is quoted as having once told a graduate student.

Tsien's story goes something like this: Born in China in 1911, he received his BS degree (in railway engineering) in Shanghai in 1934, his master's degree at MIT in 1936, and his PhD in aeronautics at Caltech in 1939. A protégé of Theodore von Kármán, Tsien worked alongside Malina on military rockets during the war, consulted for Aerojet on rocket engines, and served as a member of the air force Scientific Advisory Board from its inception in 1945.

After the war, Tsien joined the faculty of MIT, then returned to Caltech in 1949 to become the first director of the school's new Jet Propulsion Center. He also took the necessary steps that year to become a U.S. citizen. By then, a series of political events—ranging from the trial of Alger Hiss and the Russian detonation of an atomic bomb to Mao Tse-tung's victory over Chiang Kai-shek in China, and the Korean War—turned America's obsession with anti-Communism into an international crusade.

The first hint that Tsien's future was on the line came in the spring of 1950, when he learned that FBI agents had been on campus asking questions about him. In June he lost his security clearance. It was now revealed that Tsien, Malina, and Weinbaum, back in the 1930s, had participated in what they called a social group but which actually had turned out to be Professional Unit 122, the local Communist group.

In 1950, accused of having concealed membership in the Communist Party, Tsien was arrested by the FBI, held without bail for two weeks at the Immigration Detention Center at Terminal Island, then released on bail until the time of his hearing. After reviewing the file on Tsien, President Truman's assistant secretary of the Navy, Dan Kimball, wrote Caltech President Lee DuBridge a melancholy note about the government's case against Tsien. "It is nothing but a witch-hunt," he told DuBridge.

Tsien was ordered deported but prohibited from leaving the country. Free on bail, he continued to teach and work at Caltech until August 1955, when the Immigration Service notified him that he was free to leave the country. Tsien, accompanied by his wife and children, sailed for China that September, and he has never returned to the United States. The deportation order was rescinded in 1984; the charges are still on the books.

America's loss was China's gain. Chang describes how Tsien, a first-rate engineer, presided over the rise of China's missile program, building it from the ground up into the formidable military enterprise it has become in recent decades. To Chang, it's just a simple extrapolation from Tsien's training and work in America to the Silkworm antiship missile used against the United States during the Persian Gulf War.

Chang starts out determined to rescue Tsien's reputation. Tenacious to a fault, she has interviewed everyone even remotely connected with the events leading up to his arrest and detention: Malina's first wife, Weinbaum's second wife, the owner of the Bekins Van and Storage Company in Pasadena. Was Tsien a Communist in the 1930s? Did it matter then? Does it matter now? Chang solemnly tells us that "an independent investigation conducted by the author revealed that it was unlikely that Tsien had ever joined the party." She praises DuBridge for his efforts on behalf of Tsien, but faults the aeronautics department for not trying hard enough to vindicate him.

But Chang can't quite make up her mind what the focus of her book is. Two-thirds of the way through, she does an about-face. She lectures the reader about Tsien's shortcomings: he suffered from the sin of too much pride. "Most likely," Chang writes, "if Tsien had kept a low profile during the McCarthy era . . . he . . . would have suffered a decade of lost clearance . . . and reclaimed his clearance at a later date." Even prison, she seems to suggest at one point, would have been better than serving the Communist regime in mainland China. Tsien is no longer the foreigner caught in the FBI's web; he has become our enemy, the scientist who shares the blame for building weapons that can cause destruction on a global scale.

At one point the author muses: "If Tsien had died in 1955 and had never gone to China, his life would not have merited a first-rate biography." Can anyone blame Tsien for not wanting his biography written just yet?

Judith Goodstein, who holds a PhD in the history of science from the University of Washington, has been Caltech's archivist since 1968 and registrar since 1989, and is also a faculty associate in history. She has long had an interest in the Red-scare era of Caltech's history.

Books
continued

Depth Takes a Holiday
Essays from Lesser Los Angeles

Sandra Tsing Loh
Riverhead Books, New York; 1996

Sandra Tsing Loh is 1) a reluctant Caltech graduate (BS Physics, 1983); 2) a musician and performance artist who once played her piano atop a downtown building while scattering dollar bills (her own) on the audience below; 3) a very funny woman, as this collection of essays, most published in *BUZZ* magazine, hilariously demonstrates. Loh was pushed by her father, also a Caltech graduate, to be an aeronautical engineer: "He believed I was destined to shine in the Advanced Tactical Weapons Division at Hughes Aircraft Company," she says. "He was wrong." Hughes' loss is our gain. Loh's weapon of choice is clearly the keyboard (piano or computer), and her strategic target LA, including the San Fernando Valley ("the Grid"); beginning screenwriting classes (". . . the screenplay is a thong bikini, exposing all structural flaws. I and my pear-shaped musings were advised to cover ourselves in the loose old bathrobe of the novel"); and most of all, her own "futon set"— those downwardly mobile, arty, 30-somethings squeezing out meager livings from strange part-time jobs while they wait for MacArthur Genius Grants, lust after IKEA furniture ("an enlightened person . . . understands that self-assembly is the key to affordability") and consume room-temperature Trader Joe's products (not being able to afford French brie, they console themselves with the Trader's "canny invention: Canadian brie!")

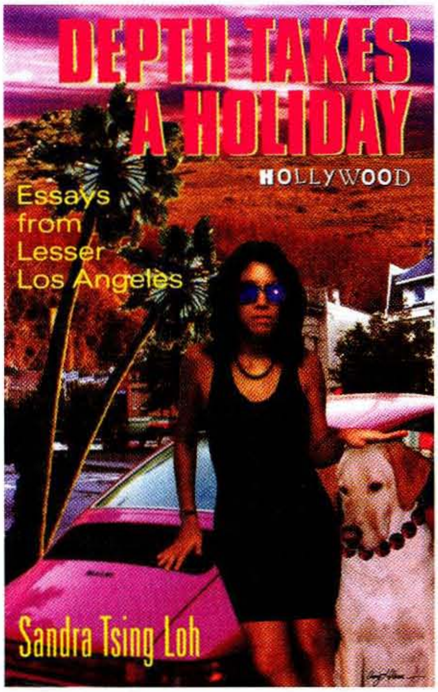

A few casual allusions to computer expertise and an admission of an addiction to Nintendo do hint at a technological bent that Caltech may have fostered. But Loh's father was probably wise not to wait for that first check from Hughes. Instead he dove for the rooftop dollars, reportedly exclaiming, "Finally I get some of my money back!"
—*Rebecca Rothenberg*

Linus Pauling
A Life in Science and Politics

Ted Goertzel and Ben Goertzel
Basic Books, New York; 1995

A second biography of Pauling published late last year (*E&S* carried a chapter from Tom Hager's *Force of Nature* in a previous issue), this one covers much of the same ground but less thoroughly—it's less than half as long. Correspondingly, Pauling's science is covered in less depth, and Caltech readers will probably find the explanations of scientific background superfluous. Sociologist Ted Goertzel's parents, Mildred and Victor Goertzel, started work on this book in 1962 as part of their work on the childhoods of eminent people, and did so with Pauling's cooperation. It does not, however, claim to be an authorized biography, and it is quite critical of the famous chemist, in particular of some of his actions in his later years. And, since the original authors were most interested in the personality of their subject, this emphasis continues to dominate the book, which includes in the appendix several fascinating current interpretations of a Rorschach ink-blot test that Pauling took in the 1950s.

Information Proficiency
Your Key to the Information Age

Thomas J. Buckholtz
Van Nostrand Reinhold, New York; 1995

A man lassoing a tornado of sheets of paper graces the dust jacket of this book, whose author offers advice on how to accomplish this difficult task—more generally defined as keeping up with the challenges and fast-paced changes of today's information technology. Caltech alumnus Buckholtz, who earned his BS in mathematics afrom Caltech in 1967 (and a PhD in physics from UC Berkeley) gives practical insights and techniques for both organizations and individuals on how to get a grip on the Information Age. Buckholtz, called "the information wizard" by *InfoWorld*, developed his ideas during a stint from 1989 to 1993 with the federal government as commissioner of the Information Resources Management Service of the General Services Administration.

Oral History
Norman Davidson

Norman Davidson, 1966

Norman Davidson, the Norman Chandler Professor of Chemical Biology, who was recently awarded the National Medal of Science (see page 43), began his career as a physical chemist. He was recruited to the Institute in 1946 by Linus Pauling and later joined in an honorable Caltech tradition of switching from the physical sciences to biology. In this brief excerpt from his 1987 oral history (he was interviewed by Heidi Aspaturian for the Caltech Archives), he describes his conversion.

HA: Who was your main conduit?
ND: Linus was one. Linus was *the* example of a person who had the intellectual courage—you could even say the "chutzpah"—to think, "Well, if I know basic chemistry I can apply it to biology." Of course, since he was a genius, where some other people might not have done it so well, he did do it with extraordinary skill and made extraordinary contributions, as the record shows.

Delbrück was easy to talk to. Delbrück had been a physicist, and at this time was not interested in biochemistry—or in molecules. In fact, it's said that he vetoed the suggestion that John Singer, then a senior research fellow here and a very good protein physical chemist, should be on our faculty, because he didn't think that the field had any future. He thought genetics and virus phenomenology was the way to go. Later, he realized he was wrong and he changed his mind. But Caltech was not strong in the biochemistry of DNA at the time. . . .

To some extent, I knew that very exciting things were going on in biology, but at least right now I can't remember specifically what I knew in detail. The Watson-Crick structure had been discovered, and it was realized that this was going to be central to the understanding of genetics and would found the subject of molecular genetics. But, to my recollection, there wasn't an awful lot of that nature going on at Caltech at that time. . . One day, I remember, a guy named Frank Schmidt came to visit. Schmidt was a professor at MIT and a great organizer and promoter—in the good sense of the word—of what was then called biophysics. He was a crusader for converting physical scientists into biophysical scientists. He'd heard that I was interested in this, and I remember having lunch with him at the Athenaeum. He said, "We're going to have this big four-week conference at Boulder, sponsored by the Biophysics Study Section of the NIH. The idea is to educate bright young physical scientists about what's going on in the new biology and what contributions they can make."

I went to the Boulder conference. It was the summer of 1958, I think. It was marvelous. It was a typical kind of a meeting of that type. In addition to the people who were supposed to be the educatees—the students—a tremendous number of leaders in the fields were there. Basically, they gave lectures, and then there were workshops in which they really talked to one another more than to us, and we were supposed to try to find out what was going on. But I have this mental picture of Leo Szilard, who after World War II, with his student Aaron Novick, had gone into one area of biology from physics. He was kind of a senior statesman. . . . At every lecture Szilard would sit in the front row and listen to the first three or four minutes of the lecture. The titles all seemed fascinating, and I was sitting with anticipation in the back. Sometimes, though, the first three or four minutes were just dull, and I kept thinking, "Gee, this is supposed to be an exciting topic. When's it going to get exciting?" But after three or four minutes, if it wasn't exciting, Szilard would get up and walk out. He didn't leave like some

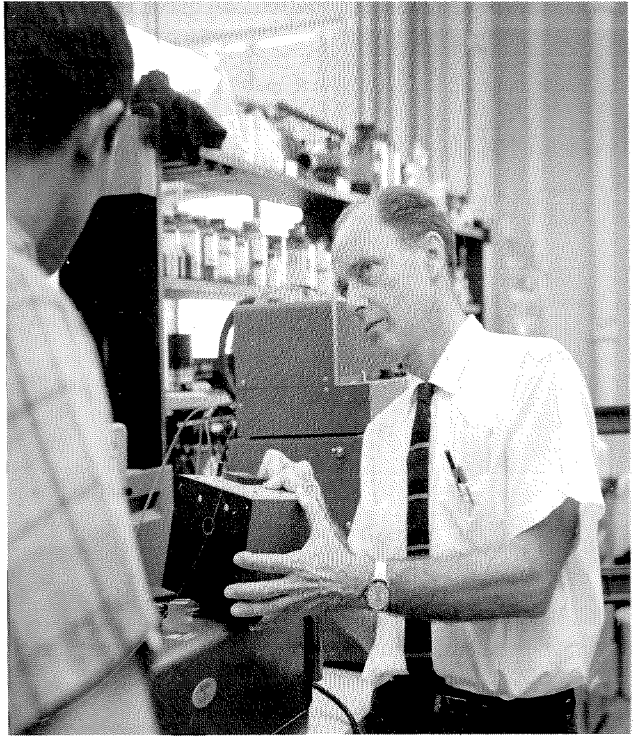

Former chemist Davidson uses a spectrophotometer to measure DNA concentrations (1966).

people do—wait till the room is dark, then hunker down and sort of sneak out. He just stood up and slowly walked out. And by God, he was never wrong. Every time he stayed, the lecture was good; every time he left, the next 57 minutes were as bad as the first 3. And I never had the guts to walk out when he did.

I forget who all was there among the physical scientists. Charlie Townes, who invented the laser, was there. He toyed with the idea of becoming a serious biophysicist, but never did make the switch. Bruno Zinn was there, along with a lot of other people who had basically made the conversion, although they weren't fully established yet.

But the point I want to make is, so far as I know, the only real hard-core physical scientist who became a hard-core, exclusively biological scientist as a result of that conference was Norman Davidson. In a certain sense, they spent $500,000, or whatever, to convert me. . . .

After this conference, I came back to Caltech, determined not to build a new shock tube but to change fields. There are several ways you can make a major change, especially from physical science to biology. The most courageous way is the way Max Delbrück and Seymour Benzer did it, in which they said, "We are not going to use any of the specific techniques and approaches we have learned in our work as physicists"—Max in theoretical and nuclear physics, Seymour in semiconductor physics. "The only thing we're going to bring from physics to biology"—and this was Delbrück's *real* contribution—"is the habit of looking for systems where you can ask specific questions, preferably with quantitative evaluations of the answers." Clear-cut qualitative answers are really just as good; but Delbrück's major contribution by consensus agreement—this is not an original idea of mine—was to select bacteriophage for that purpose. Benzer picked a specific genetic locus in T4 bacteriophage and made major contributions to the nature of mutations, but he didn't use any solid-state physics.

I was smart enough to realize that for the major questions, much of the fast-reaction technology and intellectual approach that I had developed wasn't really very useful. You could do good experiments and publish papers, but they really weren't central. On the other hand, I did decide to continue to use physical chemistry and inorganic chemistry to try to study DNA. I realized early that there were several important questions you might be able to study. I learned that somebody had done some simple initial experiments on mercury and DNA. I knew enough about mercury ions and their complex chemistry to realize that these had the potential of being very clean complexes, which might be useful. I thought they might be useful for x-ray diffraction in order to do structural work using the principle of heavy metal substitution. That turned out to be completely wrong because the structures became completely disorganized on binding mercury, and it's never been used usefully for that. But it did turn out to be useful for other purposes because of its unique and simple chemistry, and I recognized that. . . .

HA: I was interested in hearing more about some of your colleagues' reactions when you decided to move out of physical chemistry and more into the molecular biology area.

ND: In general, this is a place that respects independence and initiative. I can't recall anybody making any critical remarks. I can recall a number of questions about how I was going to do it. But the important point is that Caltech is an environment that understands and appreciates interdisciplinary research and science. As I said previously, there were precedents in Delbrück, Pauling, and Vinograd. I think the main thing is it really was a very supportive environment. Even people who don't know anything about it appreciate people moving into new and exciting areas. There are some instances around here of people who haven't been terribly successful in trying to make comparable switches; so that in a certain sense, the proof of the pudding is how the pudding tastes, how things actually work out. In my case, they clearly did work out well, both in the objective scoring of what happens to your research grants under peer review, and in the more valid subjective scoring of how your work is perceived by colleagues in your field. ☐

Random Walk

Norman Davidson Wins National Medal of Science

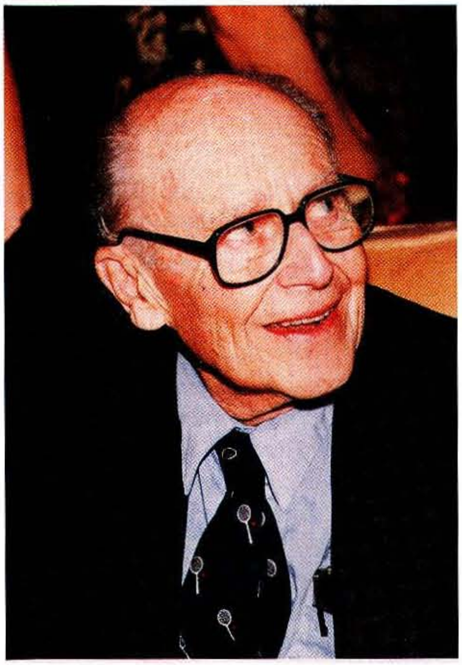

Norman Davidson, the Norman Chandler Professor of Chemical Biology, Emeritus, is among eight winners of the 1996 National Medal of Science. He will receive the award from President Bill Clinton at a White House ceremony later this summer. Davidson is the 20th member of the Caltech faculty to be honored with this award.

In his research, Davidson created innovative methods to bridge the gap between the physical and biological sciences. He pioneered new methods in physical chemistry, specifically for the study of fast reactions behind shock waves and by flash photolysis. Later, he developed new techniques, including electron microscopy, for genetic mapping and for exploring the informational properties of DNA and RNA. In his current research, Davidson is working on creating methods for studying electrical signaling in the nervous system and the ways in which it changes during learning and the formation of memories.

Davidson received his PhD from the University of Chicago in 1941. He came to Caltech in 1946 as an instructor in chemistry, became a full professor in 1957, and was appointed the Chandler Professor in 1982.

Honors and Awards

Yaser Abu-Mostafa, professor of electrical engineering and computer science, has been awarded the 1995–1996 Feynman Prize for Excellence in Teaching. The honor, presented annually to a Caltech professor who has demonstrated "unusual ability, creativity, and innovation in teaching," consists of a $3,000 prize, matched by an equivalent increase in the awardee's salary.

Michael Alvarez, associate professor of political science, has been chosen by the Midwest Political Science Association to receive the Sprague Award for his paper with John Brehm entitled "Are Americans Ambivalent About Affirmative Action?" The paper was considered to be the best delivered at the 1995 meeting to apply quantitative methods to a substantive problem in political science.

Michael Aschbacher, professor of mathematics, was elected vice president of the American Mathematical Society for the 1996–98 term.

Jacqueline Barton, professor of chemistry, has been awarded the Paul Karrer Gold Medal by the University of Zurich, and delivered the Paul Karrer Memorial Lecture there in June.

Mory Gharib, professor of aeronautics, has been elected a Fellow of the American Society of Mechanical Engineering.

Harry Gray, the Beckman Professor of Chemistry and director of the Beckman Institute, has been awarded the Sierra Nevada Distinguished Chemist Award of the American Chemical Society.

Michael Hoffmann, the Irvine Professor of Environmental Science, has been honored as a Distinguished Lecturer in Environmental Chemistry by the São Paulo State Foundation of the University of São Paulo in Brazil, and by the Hebrew University in Jerusalem.

Jeffrey Hubbell, professor of chemistry and chemical engineering, received the Clemson Award of the Society for Biomaterials at the 5th World Biomaterials Congress in Toronto.

Jonathan Katz, assistant professor of political science, has been awarded a 1996 Haynes Foundation Faculty Fellowship to pursue his research project entitled, "Why Did the Incumbency Advantage Grow in U.S. Congressional Elections?"

Daniel Kevles, the Koepfli Professor of the Humanities, Senior Trustee Ralph Landau, and Nelson Leonard, faculty associate in chemistry, have been elected to the American Philosophical Society, founded by Benjamin Franklin in 1743. America's oldest learned society recognized the three both for significant contributions within their own fields and for a broader range of interests.

Richard McKelvey, professor of political science, has been awarded the Rochester Distinguished Scholar Medal from the University of Rochester.

Thomas Palfrey, professor of economics and political science, has been selected a Fellow of the Econometric Society.

Everhart to Step Down as Caltech President

Caltech President Thomas E. Everhart has announced his intention to step down from the Institute presidency on or shortly after September 1, 1997. In his letter to Caltech faculty, students, and staff, he noted: "I will have served a decade as president by that time, and I have always thought that institutions, like people, need renewal: new ideas, new vigor, possibly new directions. In that sense, it is time for a change." Everhart said that he was notifying the Caltech community of his plans now to ensure time for a smooth transition and to avoid any ambiguity about the timing of presidential succession.

"I am proud of this institution and all that has been achieved during my time here, both on campus and at JPL," Everhart said in his letter. "Although Caltech, like all research universities, may face uncertain times in the days ahead, we have the traditions, the people, and the facilities to face them with optimism."

In his 10 years in office Everhart has overseen the construction of Beckman Institute, the Keck Observatory in Hawaii, Braun Athletic Center, Moore Laboratory of Engineering, Avery House (which will open this fall), and the Fairchild Library; and the successful completion of the Campaign for Caltech, which raised close to $400 million. Caltech's Board of Trustees will initiate the search process for a new president in the near future.

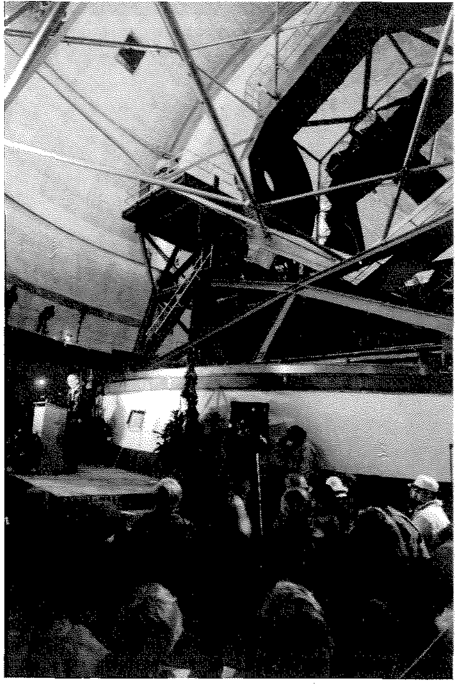
Ed Stone welcomes guests to Keck II's vast dome, where the dedication ceremony was held.

Keck II Twin Telescope Dedicated in Hawaii

The second 10-meter Keck Telescope was dedicated May 8 on Mauna Kea in Hawaii. The site of the W. M. Keck Observatory is "as close to space as you can put a telescope without going into orbit," said Edward Stone, Caltech vice president, director of JPL, and chair of CARA (the California Association for Research in Astronomy) as he welcomed the approximately 200 guests, who had ascended the 13,600-foot mountain. It's also surrounded by a thousand miles of ocean, Stone explained, has no mountains or city lights to disturb the atmosphere, and is cloudless for 300 days a year—perhaps the most perfect place on Earth from which to explore the heavens.

A Hawaiian kahuna opened the occasion with a chanted blessing. In addition to Stone, who acted as master of ceremonies, brief remarks were offered by the presidents of the three universities involved in the project—Tom Everhart of Caltech, Richard Atkinson of the University of California (joint partner in CARA), and Kenneth Mortimer of the University of Hawaii (which donated the site)—as well as NASA Chief Scientist France Cordova (PhD '79) and Robert Day, president of the W. M. Keck Foundation.

Howard Keck, the foundation's board chairman, was unable to attend because of illness, but did watch the ceremony via satellite broadcast to his home and was able to witness the naming of an asteroid in his honor. Keck got the telescope project off the ground with a $70 million pledge to Caltech in 1984, and in 1991 the Keck Foundation made another pledge of up to $75 million to fund the second instrument. NASA has contributed additional funds, partially for developing optical interferometry technology that will ultimately yoke the two telescopes electronically. Combining their light into one signal will produce the resolution of a single mirror 85 meters in diameter, the distance between the two instruments.

Keck II, which will be optimized for infrared astronomy, was constructed on the same revolutionary segmented design (by Jerry Nelson, BS '65) as Keck I, which has already been making significant discoveries out on the edges of the cosmos in the comparatively short time since it began operation in 1993 (see *E&S*, No. 1, 1996). The two instruments are the world's largest optical and infrared telescopes—a massive project, conceived and constructed over the course of a dozen years. But, as Gordon Moore, chair of the Caltech Board of Trustees, noted at the dedication, the 12 years during which the Keck Foundation had supported the observatory project seems short compared to the 15 billion years that the telescopes will be able to look back in time—almost to the birth of the universe.

Keck II's first-light image, made on April 27, the first night of observing, shows a nearby bright barred-spiral galaxy, NGC 5850, which is located approximately 150 million light-years away in the constellation Virgo. Although Keck II, like its twin, will be looking at galaxies billions of light-years away, studying relatively close ones helps astronomers understand how such galaxies formed and evolved over time.

Engineering & Science

California Institute of Technology
Pasadena, California 91125