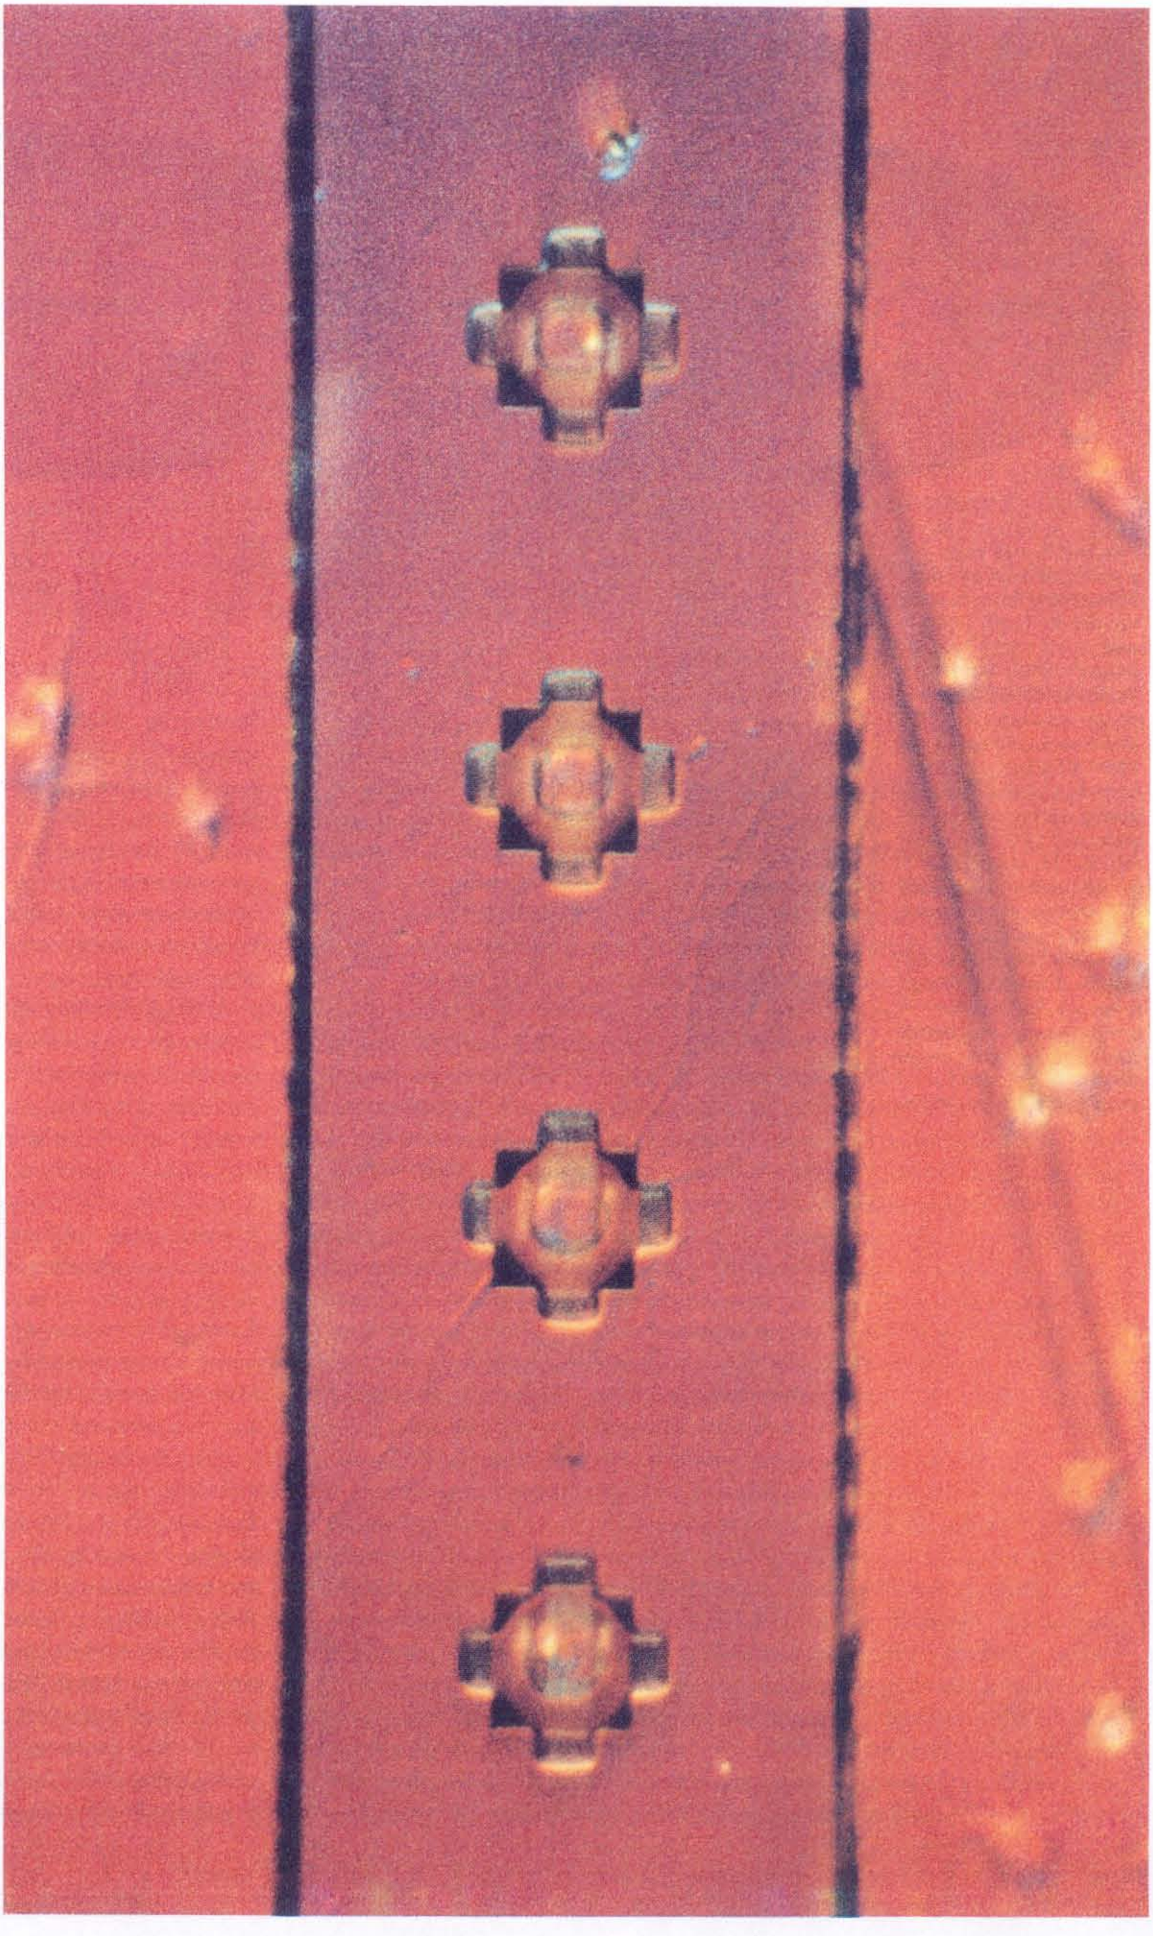

Report from a Small World

The trick to micromachining—and a big reason why the field is still in its infancy—is to figure out how to make things with moving parts, but using tools designed to manufacture immobile electronic circuits.

by Douglas L. Smith

What looks like a Navajo blanket here is actually a probe to study brain function. The reddish-brown strip down the middle of this photomicrograph is part of a silicon needle 0.15 millimeters wide. The four crosses down the center are wells the size of single nerve cells. Immature nerve cells have been implanted in the wells, and the probe will soon be inserted into a living brain, where the researchers hope that the probe cells will wire themselves into the brain's circuitry. (The second nerve cell from the bottom has already begun to send out "feelers" in search of other nerve cells.) The orange-red background is a nerve-cell culture medium; several nerve-cell bodies can be seen in it as light-colored blurs.

You may remember a photo of three intermeshed gears that *Time* magazine ran back in 1989. These gears, made at Bell Labs, were noteworthy in several respects: each tooth was the size of a blood cell; the gears, their axles, and their enclosure had been carved from a silicon chip with standard integrated-circuit-making technology; and they actually worked! Blow a puff of gas across the end one, and all three spun. The accompanying article described how several labs were making tiny springs, itty-bitty motors, and other microcomponents that might some day be assembled into microrobots that would cruise through your bloodstream like roving Public Works Department crews. "Dr. Iwao Fujimasa, a cardiac surgeon at Tokyo University, is building a robot less than one millimeter (0.045 inches) in diameter that could travel through veins and inside organs, locating and treating diseased tissue." The good doctor hoped to have a prototype to test on horses in three years, subject to the availability of parts—robotic, not equine.

Five years have come and gone, and if there *is* a microrobot jackhammering arterial plaque deposits somewhere out there, it's a safe bet that your HMO won't cover the procedure. Although microelectronic circuits are now as cheap as dirt and as pervasive as paper—you can even buy cards that sing "Happy Birthday"—the microfabrication techniques that sparked the electronics revolution have yet to ignite a mechanical one. Nevertheless, micromechanical devices—sensors, primarily—*are* making it out in the real world. The definitive sign that they've "arrived" is that they're now worth stealing—the theft of car stereos is taking a backseat to air-bag extraction as the hottest trend in auto burglaries; and the gadget that makes the air bag possible—the sensor that tells it to inflate when you slam into a tree, but not when you slam on the brakes—is a micromachined accelerometer.

You'd need an accelerometer to keep up with the growth of this field. It was all but nonexistent when Assistant Professor of Electrical Engineering Yu-Chong Tai was a graduate student a few years back. "I'd go to a conference and I'd basically know everybody. Nowadays, you go to a conference, and always more than 50 percent of the faces are newcomers. This society is expanding worldwide. It's like a disease, now—all the high-tech companies have it." If that's the case, then Caltech's biohazard lab is Tai's micromachining laboratory. The lab, currently located in Steele, will nearly double in size with the addition of space in the Moore Laboratory of Engineering, which is currently under construction.

The lab has several micromachine "viruses" in culture, as it were, but the one closest to being released is a micromotor for hard-disk drives. Lyndon Johnson was fond of saying, when told of a scientific advance, "How will this help Grandma?" Well, if Grandma has a computer, it will help her a lot. (Even if she doesn't, her gerontologist and pharmacist assuredly do.) As PCs give way to laptops, and laptops to notebooks, and presumably notebooks to wristwatches, more and more memory gets crammed into less and less space. The hard drive in your average PC is about the size and shape of a Kaiser roll, and stores 400 to 700 million bits per square

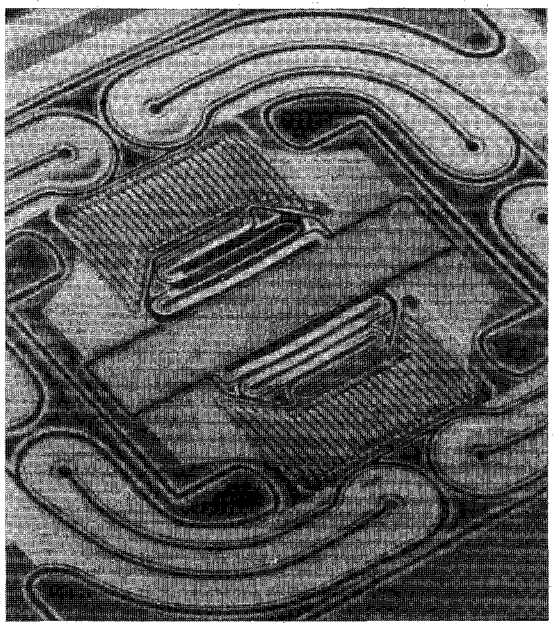

The complete microactuator. The large hairpin springs that connect the frame to the actuator curl around its four corners. The two sets of roughly vertical parallel lines just inside where the springs meet are the tops of the two motors' stator coils. Nested between the coils, and perpendicular to them, are the two thin hairpin springs that support the beam to which the read/write head is attached. The beam itself is the horizontal rectangle in the center of the picture. The entire microactuator is about three millimeters square.

inch of disk. (For comparison, the 44-million-word *Encyclopædia Britannica* runs 2.4 billion bits, not counting the index or the illustrations.) Since 1992, Tai and Denny Miu, an assistant professor of mechanical engineering at UCLA and this year a visiting associate in electrical engineering at Caltech, have been developing the technology needed to make credit-card-sized drives one centimeter thick that will hold one to two billion bits per square inch. The long-term goal is to keep scaling these drives down until they can be arrayed on circuit boards, the way memory chips are mounted, to make the massive storage space of disk drives as instantly accessible as chip memory.

A hard-disk drive works much like a phonograph. In both cases, the information is written on the surface of a disk that spins underneath a stationary arm—the tone arm in your stereo, or a stainless-steel suspension arm in your computer. The arm pivots to reach any part of the disk, from the rim on in. (The disk drive is actually a stack of up to a dozen platters, often with less than an eighth of an inch of space between them, spinning on a common shaft. Each side of each platter has its own suspension arm, so that the drive plays the A and B sides concurrently, without having to flip the record over.)

But whereas the sounds of *Sergeant Pepper's Lonely Hearts Club Band* are transcribed onto an LP as a wavy groove, which re-creates the music by vibrating a needle that's inserted in it, the data on a disk are encoded in puddles of magnetic polarity that an electromagnetic transducer—called a read/write head—interprets as a string of zeros and ones. The read/write head doesn't touch the spinning disk, but floats on an air cushion a couple of millionths of an inch thick. And whereas the cuts on a record are segments of one long spiral that takes up the entire album side, allowing any song to be played in its entirety once the needle touches down, the tracks on a hard disk are concentric circles one data bit wide. In order to retrieve a file, the read/write head skitters like a hockey puck from track to track, picking up file segments on the fly.

Current technology squeezes 5,000 tracks into an inch—in other words, 30 tracks would fit within the thickness of this page. Tai's group is initially aiming to cram 10,000 tracks into an inch in the credit-card-sized version, eventually upping that to 25,000 in the chip-sized one. So hitting the correct track is a lot harder than cueing up "Lucy in the Sky with Diamonds"—you can't just squint at the record and drop the tone arm into the dark space between songs. And if you miss your aim on an LP and drop into the song halfway through the third note, the skipped data merely jars your ears. A similar disk error would render the file unreadable. Moreover, the suspension arm is enormous, compared to the tracks—it's like trying to rotate the tower crane at the Moore Lab construction site to within two hundredths of a degree. Imagine trying to do this every 12 milliseconds—the amount of time the suspension arm has to find its next track.

But if the read/write head were mounted on a micromachined actuator, which in turn was attached to the suspension arm, it wouldn't need such exact control—you could just move the arm close, then jockey the actuator to the right track. (Compact-disk players, which pack 18,000 tracks per inch, use such a two-stage gadget, but it's much too big to wedge between the hard drive's platters.) Tai and Miu's actuator is carved from a silicon slab, yet has an almost lacy quality. The read/write head hangs from a beam supported by two impossibly delicate springs—flat, hairpin-turning squiggles that zigzag back and forth. Flanking the beam are two micromotors that pull the read/write head from side to side. The micromotors are what's called variable-reluctance motors. They work in the same way that an electromagnet made by wrapping copper wire around a nail picks up another nail. "In our case," says Miu, "the nails are permalloy, which is 80 percent nickel and 20 percent iron. One nail is the stator, which is fixed to the actuator and has the copper coils; the other nail is the rotor, which is fixed to the beam and moves the read/write head." The whole business is slung by four more hairpin springs—relatively big ones, this time—

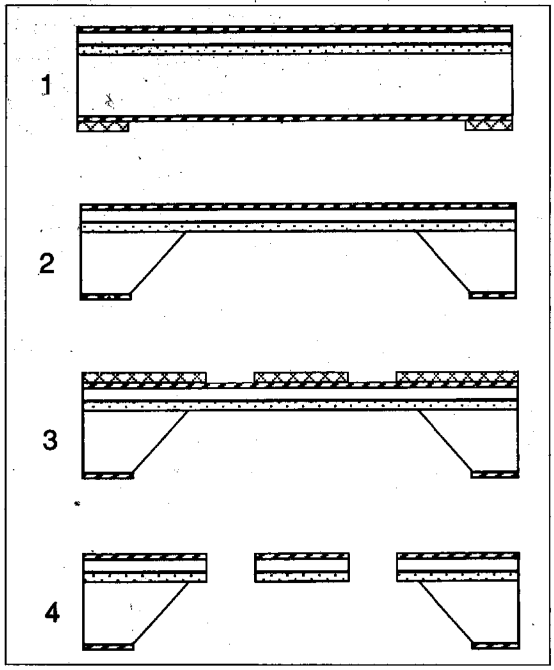

Right: A simplified cross-sectional schematic of how the springs are made. The materials are crosshatched according to the key below. (Si is silicon, B is boron, SiO_2 is silicon dioxide, PR is photoresist, Cu is copper, and NiFe is permalloy.)
1.) The composite wafer's underside is patterned with photoresist in the shape of the diaphragm to be etched.
2.) The etchant eats up through the wafer to the silicon-boron zone.
3.) The wafer's top surface is patterned with photoresist in the shape of the springs.
4.) The springs are cut from above with reactive ion etching.

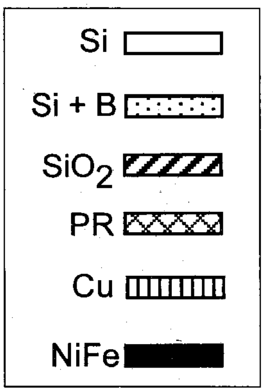

within a frame that's still part of the same piece of silicon, and the frame is glued under the suspension arm's tip.

Now the trick to micromachining—and a big reason why the field is still in its infancy—is to figure out how to make things with moving parts, but using tools designed to manufacture immobile electronic circuits. The technology is the same—you cover the chip with a mask, then add a layer of something to (or strip a layer of something from) the parts of the chip exposed through the mask. The trick within the trick is planning ahead so that succeeding steps don't mess up what you've already done. Conceptually, there are two basic processes for making the actuator: one for carving the springs, and the other for building the motor. In reality, the two processes are interleaved. The entire procedure requires 20 masks—the equivalent of a memory chip. It's the most complex structure Tai's lab has built.

Cutting the springs is the simpler process. It starts with a silicon wafer 500 microns thick. A thin layer of a silicon-boron mixture is applied to the top surface by chemical vapor deposition silicon epitaxy, meaning that the silicon and boron atoms form a single crystal that blends seamlessly with the pure silicon below. Then comes another 20 microns of pure silicon—which again continues the single crystal—followed by a thin film of silicon dioxide, which acts as the plastic wrap on the sandwich and is applied to both the top and bottom surfaces. Next, the frame is masked off on the wafer's underside and etched from below. Applying the mask is a darkroom process exactly like printing a photograph: you shine a strong light through the negative to project the image onto light-sensitive paper. In this case, a photoresist—a light-sensitive chemical—is spin-coated onto the wafer, and the negative carries the frame pattern. (Spin coating is a neat way to get a very uniform layer of something without much fussing around—you hold the wafer horizontally and put a puddle of the coating in the center, then spin the wafer at several thousand revolutions per minute; centrifugal force does the rest.) When the photoresist is developed, the illuminated stuff doesn't stick to the chip any more and washes off, exposing the areas to be etched. The chip is bathed in hydrofluoric acid, which removes the silicon dioxide in the exposed areas, transferring the mask to the silicon dioxide layer. (The photoresist itself can't stand up to the etchant that follows, but silicon dioxide can.) The photoresist is rinsed off with a solvent, and the chip is then dunked in the etchant (ethylene diamine/pyrocatechol), which eats up through the wafer to the silicon-boron zone. The etchant can't digest the silicon-boron mix, leaving the 500-micron-thick wafer framing a 20-micron-thick diaphragm—the silicon-boron layer and the stuff above it—into which the springs will be carved. Their pattern is masked off on the diaphragm's top surface, using another layer of photoresist, and is cut by reactive ion etching. In this technique, the wafer is bombarded with a sulfur hexafluoride plasma, which consists mostly of fluorine ions that just tear into the unprotected silicon. Once the diaphragm is cut all the way through to make the hairpin springs (which takes about half an hour), another solvent rinse removes the photoresist.

The construction sequence for the motor, using the same crosshatching code.
1.) A "seeder" layer of copper is applied to the wafer's top surface.
2.) A photoresist mold is patterned in the shape of the coil's bottom half.
3.) Electroplated copper fills the mold.
4.) The photoresist is rinsed off, and the exposed (unplated) seeder etched away.
5.) A fresh layer of photoresist is applied, which heat transforms into a permanent insulator.
6.) Another copper seeder layer is deposited, followed by the photoresist mold for the permalloy core.
7.) The core has been plated on, the mold and seeder removed, and a fresh layer of baked-on photoresist insulation added.
8.) A new seeder layer goes on over the insulation.
9.) Another photoresist mold for the top and sides of the coil follows.
10.) The remainder of the coil is plated on, and the photoresist mold and the remaining seeder layer removed.

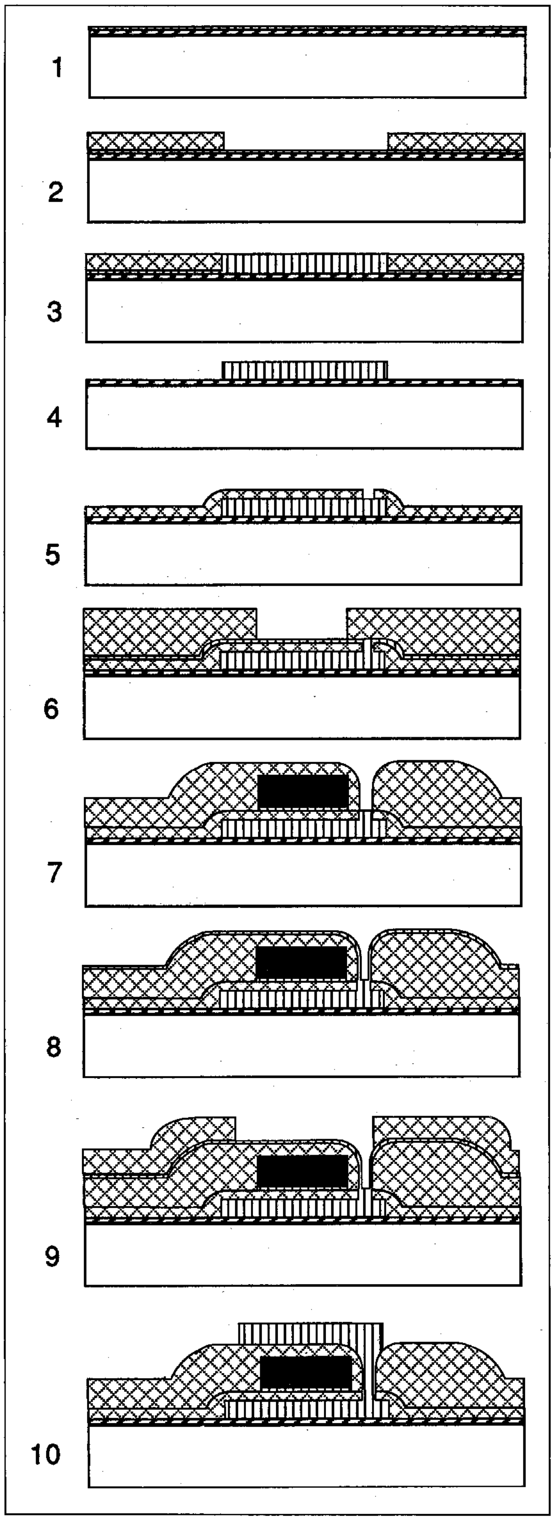

The springs are flat, but the motor is three-dimensional; consequently, making it is considerably more complicated. You have to wrap a copper coil around a permalloy core, and since the motor is embedded in the actuator, you can't just pick up the core with tweezers and wind wire around it. So the construction proceeds in three stages: first the bottom part of the coil, then the core, and finally the coil's top and sides. The metals are deposited through a process called mold electroplating. Electroplating is commonly used to coat one metal with another—you clip an electrode to a hubcap, for example, dunk it in a bath containing chromium ions and the other electrode, run a current through the circuit, and—zap!—a chrome-plated hubcap. But silicon doesn't conduct electricity very well, so the first step is to apply a "seeder" layer of metal to the whole surface. (This requires yet another technique, called vacuum thermal evaporation, in which you place in a high-vacuum chamber the wafer and a small crucible of the metal to be deposited. The crucible is heated electrically until the metal evaporates, and the vapor then deposits itself on the relatively cool room-temperature wafer like shower steam on your bathroom mirror. Of course, the vapor also deposits itself all over the rest of the vacuum chamber's interior, but oh, well...)

"Electroplating different metals takes different seeders," Tai explains. "For example, for copper we put down 100 Ångstroms of chrome and 1000 Ångstroms of copper." Then comes the photoresist, etc., leaving the seeder exposed where the copper is to go. After copper fills the photoresist mold, the solvent strips the mold away and an acid etch gets rid of the unplated seeder layer. The acid takes a wee bit of the copper, too, but since the copper layer is some 10 times thicker than the seeder, it doesn't matter. What's left is a set of parallel copper lines, slightly slanted, which will be the bottom part of the coil, as shown on the opposite page. (It generally takes about 10 working days from the time you started on the springs to get this far.)

Now you need insulation—if the copper windings touch each other or the core, the motor will short out. It turns out that photoresist is a really good insulator, so a fresh layer of photoresist takes care of that. You have to plan ahead at this point, and remember to pattern this insulating layer to create the holes through which the two halves of the coil will eventually connect. Now you're ready to spend another day mold-electroplating the permalloy core. But how do you remove the *mold* photoresist (and etch off the seeder beneath it) without stripping off the *insu-*

lating photoresist too? The answer is that you very cleverly baked the insulating photoresist before starting work on the core. The heat turns the photoresist into a long-chain polymer that can withstand the solvents and etchants. "There's a lot of materials science going on here," says Tai. "A lot of these processes are intimately related to the mechanical and electrical properties of the materials. These are the details that actually decide whether the process works or not." Once the core has been laid down, there's another layer of photoresist insulation (again patterned with the holes for the coil's electrical connections). Then another round of mold electroplating for the top half of the coil, and you're done. "That's often a trick I pull on my students when we start a project—'See, it's so easy! That's the way you draw it—now, go make it!' But we know how hard it is. There are a lot of tricky steps."

Below: The evolution of a micromotor, as seen from above.
1.) These slanted parallel lines of copper will become the bottom part of the coil.
2.) The permalloy core runs up the middle of the coil.
3.) Other parallel lines of copper arch over the core and connect the slanted copper lines into a continuous coil wrapped around the core.

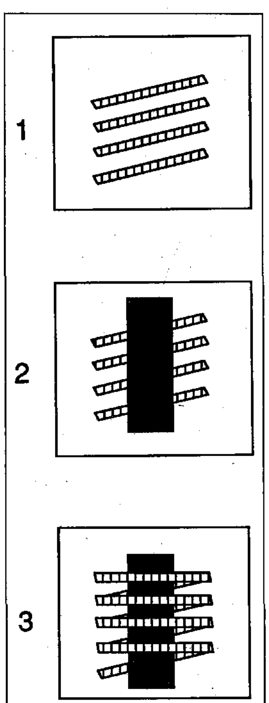

Making integrated circuits is actually easier, because they're only skin deep. The wafer is still 500 microns thick, but the lower 495 just sit there. The working parts don't penetrate any deeper than five microns into the chip, nor do they stick up any higher than five microns above it. "But when we do micromachining, we dig in. We often cut all the way through the wafer. So although the number of masks are the same, the technical issues are very different." For example, you have to treat a 20-micron-thick diaphragm with utmost care to avoid breaking it. "Also, a little bit of force can distort these structures, so we have to develop special expertise to handle them."

Tai's lab has built prototype actuators (a consortium of hardware companies are working on the drives), but there are still some kinks to be worked out. Says Miu, "right now, the arm is still supported on mechanical bearings, which gives you a certain amount of slop. Also, the wire leading to the micromotor behaves as a spring at very small deflections, so we have to account for that error." And there are other subtleties, too—the motors can't be too powerful, for example, or their magnetic fluxes can confuse the sensor that tells the read/write head what track it's on.

Another project has attracted more attention of late, even though it's much further away from practical application. In collaboration with a group of engineers from UCLA, the micromachine lab has demonstrated a "smart" skin designed to reduce turbulent drag on airplane wings. This isn't the kind of turbulence that makes the pilot turn on the Fasten Your Seat Belt sign and the flight attendants wheel their beverage carts back to the galley just before they get to your row. Instead, it's caused by the airplane itself—the wing's passage spins off swirling pockets of air, called vortices. These vortices start out lying flat against the wing, but they rapidly stand up to become miniature tornadoes, and that's when they cause trouble. Once upright, they pump high-speed air that's trying to rush past the wing down to the wing's surface. Thus the shear stress on that piece of the wing—stress caused by the air moving in one direction while the wing is moving in another—increases. (There's a certain amount of shear stress on the entire wing anyway. The wing essentially peels off a thin layer of the adjoining air and drags it along, but that's just part of the cost of doing business.)

Until now, the fuel that was burned fighting turbulent drag was unavoidable overhead, too. (And it's not trivial—one aerospace industry analyst estimates that a one-percent drag reduction for all commercial aircraft would save the airlines a billion bucks a year worldwide.) The vortices stick to the wing for less than a second before detaching themselves to go sailing harmlessly away, so there's not much time to react. You could shed them sooner by putting a ramp, such as a lifted flap, in their path—they'd hit it like a ski jumper and go roiling off into the wild blue yonder. But at typical wind-tunnel speeds, these vortices begin life larva-sized—about two millimeters wide and one centimeter long—and at jet airplane speeds, they're even smaller. So the sensors that will detect them and the flaps that will punt them need to be small, too. ("In order to test this idea in the wind tunnel, we had to make relatively large devices, and that's not very pleasant," says Tai. "This would actually have been easier on a real airplane. Micromachining technology simply isn't designed to make things that big—it's the inverse of trying to use an enormous power saw to cut a very small part.") And the vortices are all over the wing—the wind-tunnel model looks like it's crawling with maggots. Thus, the entire surface needs to be able to detect them, but only the affected regions should react to them, because if there aren't any vortices, raising the flaps will create them.

The grand design is to tile the wing with four-inch-diameter silicon chips, each of which would incorporate sensors, control circuitry, and flaps. The sensors measure shear—the proximate cause of the drag—by running a steady current through a silicon "wire" whose resistance rises rapidly with increasing temperature. The wire heats up, but the onrushing air carries the heat away. The high shear within a vortex cools the wire faster than usual, causing its resistance to drop below that of its neighbors. The controllers

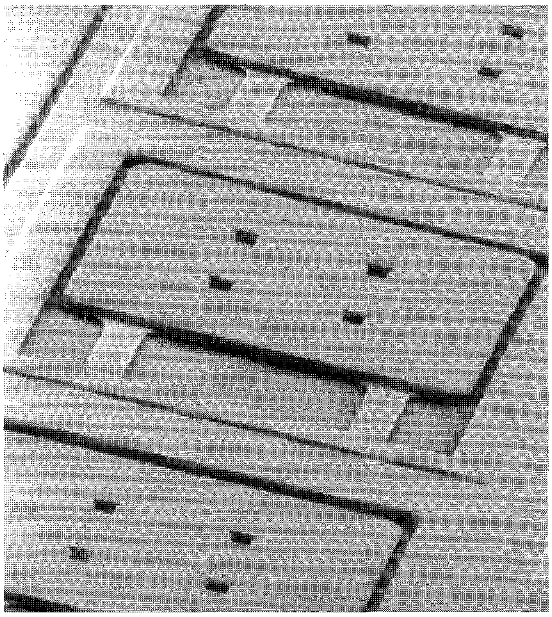

Right: What all the flap's about. An aerial view of a portion of a flap array, seen from the hinged side. The four holes in the flap allow the etchant to undercut the flap outward from the center as well as inward from the edges, minimizing the time it takes to free the flap. Both this array and the flap shown below are from the steering project.

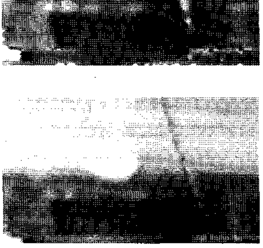

Below: Three frames from a video of a flap flapping. In the top image, the magnetic field is turned off, and the flap lies flat. (The hinge is to the right.) In the middle picture, the field is at about half strength, and the flap sticks up at a 45-degree angle. The field is at full strength in the bottom frame, and the flap is almost standing straight up.

compare the sensors' outputs to decide where the vortices are, and lift the flaps in that general area. The electronics are still being designed, in collaboration with Professor of Electrical Engineering Rod Goodman's research group, but the lab has built prototype models of the sensors and flaps.

And it's the flaps that are in the limelight. They're thin, flat, multilayer sandwiches that cantilever out over pits etched in the silicon beneath them. One of the layers is a permalloy coil, which, when electrified, raises the flap magnetically—up to a good 65 degrees from the horizontal—by pushing against the field created by another magnet on the floor of the pit below. (The magnets, which operate at 80 gauss, or about the strength of a refrigerator magnet, exert a force some 20 times stronger than gravity on a typical one-millimeter by one-millimeter flap.) Each flap can be raised or lowered individually. Most remarkable of all are the hinges—there aren't any. Instead, two tiny silicon beams connect one side of the flap to the pit's brink. In our world, silicon structures—glass windows and ceramic pots, for instance—are stiff and brittle. They resist stress until they shatter. But in the microworld, silicon behaves differently. If you make thin enough beams of it, they're quite amazingly flexible. This is actually true of most materials, because as you make smaller and smaller crystals of something, the number of lattice defects—places where the atoms don't quite line up, and where fractures can start easily—gets smaller, too. Other people had verified this with millimeter-sized hunks of silicon, says Tai, "but we've gone down to microns, and even nanometers, and we've definitely confirmed the trend. And, of course, we're enjoying it. It's a happy result."

These hinges are not only flexible, they're beefy. In a parallel project, Tai's lab and the UCLA group are building flap arrays that exert ten times more force per flap than the anti-turbulence ones—enough muscle to actually steer the airplane. A wing is normally steered by large flaps, called ailerons, along its trailing edge. The microflaps go along the leading edge instead. Both kinds of flaps work by deflecting the boundary layer—the airflow along the wing's surface that causes lift (and drag). The boundary layer is wedge-shaped—very thin along the wing's leading edge, and thickening toward the rear. As the wedge thickens, it contains more air and gets harder to move, so manipulating it from the leading edge makes a lot of sense, says Tai. "You use much less energy to achieve the same degree of control. It's like a transistor—you put a little signal into the leading edge, and it will be amplified automatically as the boundary layer goes back over the wing." Tai's group has demonstrated this approach in UCLA's wind tunnels, using a generic delta-wing model. Aerospace engineers are very interested, but the prospect of aileronless planes is probably too much for the flying public. Don't look for jumbo jets of this design any time soon.

At the moment, each of the three components in the turbulent-drag project—sensor, controller, and flap—are still separate units connected by old-fashioned copper wires. Tai expects to have the three on one chip within six months, but figuring out what sequence to make the magnets and control circuitry in is a chicken-or-egg problem: making the sensors requires heating the wafer to 800°C, which melts the aluminum connections between circuits; making the circuits entails depositing layers of silicon atoms, which clog up the flaps. "The more things you put on, the more headaches you have," Tai says ruefully. "Whenever you try to put a lot of different kinds of devices together, that means you are combining all these processes into a big, long, complicated one. We're constantly thinking about how to solve problems like this."

Any hobbyist who has ever been reduced to howling fury while trying to tweezer a balky antiaircraft gun into its mounting on a 17-inch replica of the battleship Missouri *will appreciate the frustrations of trying to do the same sort of thing on something a hundred times smaller.*

Tai sees this project as pushing the envelope, not of aircraft design, but of micromachine design. "This may never be used on a real airplane—who knows? The point is that it demonstrates a new technology that combines microsensors with microactuators and microelectronics—what I call M-cubed." Once you've integrated those three components—the eyes, hands, and brain, as it were—there's no mechanical system you can't build, at least in principle. "If we demonstrate that the technology can be developed to include all three things on one chip, we have defined the boundary of microfabrication. That's the ultimate challenge." The *real* ultimate challenge will be to figure out what undreamed-of things you can create with M^3.

For starters, here are a few things that people *have* dreamed of. Like the flap projects, these are distributed systems in which little neighborhoods of components operate independently within large arrays. First, you could use a flap array to create turbulence where you want it—in the combustion zone of a turbine engine, for example, where fuel and air have to mix fast and thoroughly. Or consider active soundproofing, in which a wall detects the sound waves hitting it and adjusts itself to damp them out. Or an array of micromirrors that, properly illuminated, would form a flat-screen TV of unlimited size. Or dish antennas that continuously adjust their surface curvature to focus a signal.

And speaking of communications equipment, the micromachine lab has joined forces with Caltech's Jet Propulsion Laboratory to demonstrate the manufacture of waveguides for millimeter- and submillimeter-wave antennas. These waves fall between microwaves and infrared light, and JPL wants to use them for deep-space communication, radar, and spectroscopy. Waveguides are essentially speaking tubes for electromagnetic waves—tunnels with reflective metal walls down which the waves travel. The waveguide's cross-sectional dimensions need to be within 10 percent of the length of the wave in order to guide it. Accurately machining a metal channel the width of a gnat's eyelash is an art that computers haven't mastered yet, and it takes months for a skilled human to make a submillimeter waveguide that works. Then, to make it into an antenna, you have to glue a transducer on it, which is also done by hand. Any hobbyist who has ever been reduced to howling fury while trying to tweezer a balky antiaircraft gun into its mounting on a 17-inch replica of the battleship *Missouri* will appreciate the frustrations of trying to do the same sort of thing on something a hundred times smaller. Micromachined waveguides avoid these problems. The channel's width is precisely set by the mask, and the depth by the etchable layer's thickness. And the transducer can be micromachined directly into the channel. Silicon doesn't reflect microwaves, but coating the channel with a reflective layer of metal atoms is standard technology, as we've seen.

Robotic spacecraft with silicon hardware are worlds less complex than live mice with protoplasmic circuitry, but a brain's a brain. The micromachining lab is using the construction techniques of the former to help study the workings of the latter. Since 1980, Professor of Physics Jerome Pine has been studying how nerve cells, or neurons, interact in networks. The idea is to grow a small array of neurons connected to one another in their normal fashion, so that you can stimulate one cell and listen in on what it says to its fellows. Growing the arrays in culture is relatively easy, but wiring them for sound is a lot harder. First of all, you can't just jab electrodes in them if you want them to live very long. Pine's first plan was to lay an array of electrodes in the bottom of a Petri dish, and then grow the neurons on it. This was fine in principle, but it was difficult to communicate with a single desired cell after the network grew a cobweb of processes—the filaments that connect nerve cells—all over the array. The next refinement was to make tiny diving-board-shaped silicon electrodes that could be wheeled up to the cell bodies. This proved awkward, but it got Pine thinking about micromachining.

In 1988, Pine's lab began making arrays of shallow wells, each of which was just the size of a mature nerve-cell body and whose floor was an electrode. An immature neuron was injected via micropipette into this dungeon, the ceiling of which was a grating that admitted nutrients and allowed the neuron's processes to grow out. As the cell matured, its body filled the entire volume of its prison and pressed tightly against the electrode in the floor, making a solid contact. The unfettered processes, meanwhile, slipped through

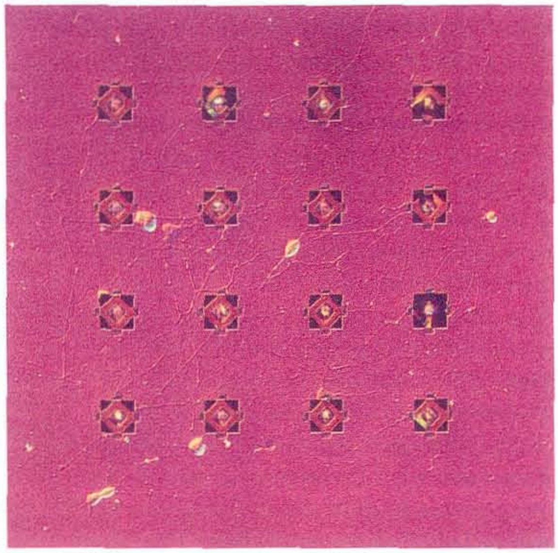

Right: An individual neurodungeon, 30 by 30 microns square and 16 microns deep. The neuron will be injected through the central hole in the diamond-shaped lattice. This cell isn't quite ready for occupancy yet, as there's no electrode in the floor.

Far right: A four-by-four cellblock. The tendrils wandering all over the surface are processes that have formed a neural network. The blobs are nerve-cell bodies, some of which probably escaped from the pen. In fact, if you look carefully, you can see a couple of them in the act of going over the wall.

Below: A couple of months after implantation, probe neurons (stained dark, and marked with arrows) have grown several hundred microns into the surrounding brain tissue. The probe (partly visible as the straight-edged black shape) runs down the photo's right side.

the grillwork and connected with the ward's other inmates. But the fabrication problems were too challenging, says Pine, so he helped recruit Tai to Caltech to collaborate on building a better neurotrap.

The collaboration is now making 16-neuron cellblocks—arrays of four cells by four cells—in which embryonic nerve cells from rats are incarcerated. The group's record for keeping neurons alive in captivity is about a month, long enough to form a network and start recording its behavior. But Pine would like to keep them alive for about three months, in order to study each network thoroughly—like snowflakes or fingerprints, no two networks are completely alike. The trouble is, the neurons climb through the bars and escape. "They squish like water balloons," says Pine. "It's astonishing how small a hole they'll get through. A 20-micron-diameter neuron can crawl through a one-by-three micron slot. They'll stay alive for three months, easy, just not where we want them." The neuron's growing processes cling to the silicon for support, and one process in particular, called the axon, is known to exert a lot of traction on the cell body—enough, apparently, to pull it through the lattice. The next design will replace the grillwork with narrow channels up to 30 microns long, down which the processes will have to grow. The hopes are that the cell bodies won't be able to stay squeezed long enough to worm through.

Tai and Pine are building similar probes to study neural activity in real, live brains. "We'd love to plant spies in brain tissue to tell us what's going on," says Pine. Multichannel electrodes are a basic tool of such studies, but driving spikes—even wire-thin ones—into the brain tends to kill or maim the cells in the immediate vicinity. This in itself is not bad, as the brain has cells to spare and there are no pain receptors in it, but the signals from the healthy cells on which the researchers wish to eavesdrop are muffled by the dead zone surrounding the probe. And the probe picks up the chatter from everything in its vicinity, while we may only be interested in the conversations over a specific phone line, as it were. Tai and Pine hope that a micromachined probe with a line of neurodungeons will minimize these difficulties. The probe neurons should send their processes out in search of healthy cells to connect to. And, by stocking the probe with a cell type peculiar to the circuit the researchers wish to wiretap, the probe might be encouraged to wire itself into that circuit as the captive neurons instinctively seek out their compatriots. (How nerve cells "know" which connections to make remains one of the great puzzles of neurobiology, but we can still take advantage of the fact that they do.) Of course, all this depends on the assumption that the imprisoned neurosnitches can survive for months or even years in the probe without special attention and "mainstream" themselves into the brain cell population.

The probes are shipped to Rutgers, where Professor of Biology Gyorgy Buzsaki's research group implants them in rat hippocampi. The architecture (although not the function) of the hippocampus is well understood, and collecting embryonic hippocampal neurons and integrating

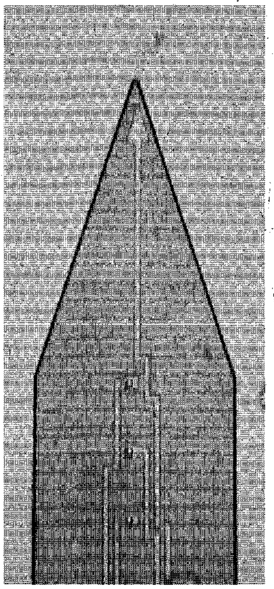

Top: The Art Deco zigzags on the probe's back side are the leads connecting the electrodes (the small squares) to the outside world. The arrowhead is a reference electrode. Bottom: The probe has 15 neurodungeons spaced 50 microns apart at the tip of a 2-millimeter-long shank that's 20 microns thick.

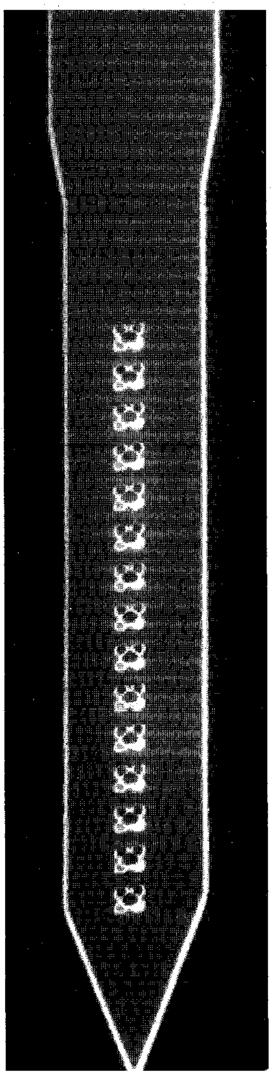

them with host neurons is Buzsaki's specialty. The Rutgers contingent has proven that the probe neurons do, in fact, grow connections to the host cells. Buzsaki's next step is to figure out exactly where those connections go, by stimulating an individual probe cell and monitoring the neighborhood's reaction, or waking the neighbors and seeing which probe cell responds.

Although these probes are strictly for basic research at the moment, Tai sees them eventually getting out into the larger world as controllers for prosthetic limbs. The probe could tap into the brain circuits that would normally move the limb, and send the electrical outputs to servomotors that could flex an artificial knee or clench synthetic fingers into a fist. Such experiments have been going on for 20 years with metal probes, but the neuroprobe offers the chance to make permanent, one-on-one connections. And you wouldn't have to intercept exactly the right circuits—probably an impossible feat in any case—since the patient's brain would automatically rewire the connections as the patient learned to use the prosthesis.

In fact, Tai sees a growth industry in biomedical microdevices of all kinds—not Dr. Fujimasa's *Fantastic Voyage* robot, but less grand schemes. For instance, one company has been making micro blood-pressure sensors for a decade, says Tai, and another is making microvalves "that could revolutionize biomedical instruments. Micromachining can make small systems that function as well as the big ones, or even better. That's terrific for biomedicine, because people want smaller and smaller devices."

We're not talking about teeny-weeny heart valves for preemies here, but something much bigger: a laboratory on a chip. When you visit your doctor for blood work in a few years, you may get away with depositing a few drops, instead of leaving what seems like a gallon's worth. Several organizations are working on scaling down the equipment needed for an arsenal of standard analytical procedures. A technique called capillary electrophoresis, for example, which is used to identify proteins or DNA sequences, separates the constituents in a sample by dissolving them and drawing them through a narrow tube via an electrical gradient. The components pass through the tube at rates depending on their size and charge, allowing each one to be identified when it emerges. Right now, such systems take a lot of fancy plumbing squeezed into a unit about the size of a home bread-making machine. Add the laser sample-detection system that goes with it, and you have another unit the size of a toaster oven. And the workhorse of biotechnology, a technique called polymerase chain reaction (PCR) that takes a snippet of DNA and copies it many times over—a critical step in screening for assorted genetic diseases—requires heating and cooling the sample over and over again, while adding different reagents at specific steps in the cycle. This also means lots of plumbing, *plus* a programmable oven. The current ones are about the size of microwaves, but instead of getting popcorn in five minutes, you get PCR in an hour. Doing the procedure on a chip, with just a smidgen of sample to heat and cool, might cut the processing time to 15 minutes. Eventually, one could design special-purpose chips to do specific blood tests while you wait—can drive-through service be far behind?

And there are a legion of applications beyond the biomedical. For example, self-contained laboratories on a chip could be used as process controls in industries from brewing beer to refining gasoline. Beyond the factory gates, such sensors could form the basis for rugged yet compact air- or water-pollution monitors.

Along with the usual M^3 problems of component integration, these projects are hampered by a lack of fundamental knowledge of what goes on in machinery of cellular dimensions. "There are so many promising applications that everybody has been spending their resources developing new devices," says Tai. "But we're neglecting the study of fundamental micromaterial properties, which we need in order to keep advancing. I can't overemphasize the importance of fundamental research, and I feel that academia, rather than industry, has an obligation to do it because it benefits everybody." Tai and Miu are therefore running a silicon microproperties lab, too.

As we saw in the case of the airplane-wing flaps, specks of silicon can behave quite differently than silicon in the large. One of the questions the microproperties project is trying to answer is just how small you can make, say, a hinge—at some point, there are simply going to be too few atoms to accommodate the bending force. The project is studying static properties such as tensile strength (how much you can stretch a sample) and fracture strength, as well as dynamic properties such as fracture propagation. The project is also looking at composites, in which the silicon has been coated with a metal, an alloy, a polymer, or even a ceramic. Most silicon microgadgets incorporate other materials, if only as the metal lead to an electrical connection. Says Tai, "Composite materials have been a big research topic in materials science, but microcomposites are relatively new and there's no general theory describing them. Microcomposite materials open up a whole new range of properties and behaviors that we can use in ways we can't even imagine because we don't know enough about them. We've already found a lot of interesting things we don't see on the macro scale." They've discovered, for example, that applying a layer of metal to the top of a silicon beam markedly alters its fracture behavior. Whenever you do a set of fracture experiments, there's always a certain amount of statistical scatter in the results. But the metal layer reduces that scatter—the results cluster more tightly around a single value. Furthermore, the alloy's exact composition strongly affects the clustering.

There's also a pressure spike at bottleneck number two, where the channel narrows to 40 microns. This could mean that the gas molecules pile up like a mob of Keystone Kops running full tilt at a narrow doorway.

And if the quintessence of rock-solid silicon changes with its bulk, it should come as no surprise that more evanescent phenomena are mutable as well. Take fluid—gas and liquid—flows, for example. The vast literature on fluids in enclosed channels (the sort of thing you use to design natural-gas pipelines or chemical plants) tends to streamline the calculations by concentrating on what's happening in the middle of the pipe and neglecting the complexities, called edge effects, that occur along the walls. But you can't do this in a microchannel, where the channel's height is comparable to the mean free path—the average distance a fluid molecule travels before colliding with another fluid molecule. At that scale, everything is edge effect. "If you don't have micromachining technology, it's very hard to do these experiments, and there's really no need for them. Now, suddenly, we have this technology, and people are showing that many useful microfluid devices can be made. But in order to

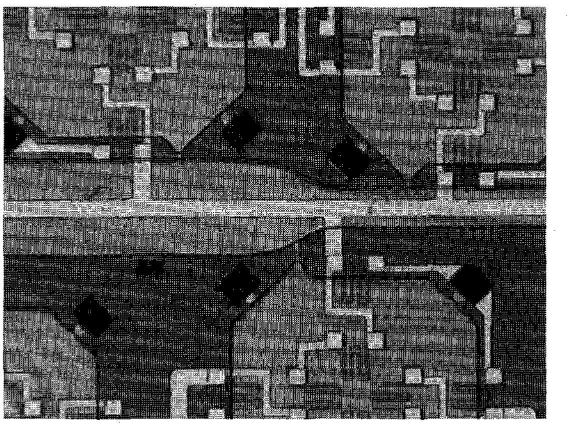

properly design micropumps and so forth, you have to know how fluids behave on this scale."

So the micromachine lab and the UCLA engineers are building wind tunnels on chips. This has required developing a micro pressure sensor that can be integrated into a channel so that the ensemble can be built as one unit. The flow at various points in the channel—which is what you really want to know—is then derived from the pressure data.

The first wind tunnel looked at the simplest possible situation—a pure gas (helium or nitrogen) in a straight, rectangular channel. And, says Tai, "we found that no theory, even when we modified the famous Navier-Stokes equations, could explain the differences we saw between helium flow and nitrogen flow." These equations, which work very well at macro scales, say that the two gases will behave differently. The gases, however, didn't behave differently in the way that the equations said they would—they behaved differently in a completely different way. None of the fluid mechanists that Tai talked to were able to explain what was going on, so the group eventually just published the data in an article that said, "Here, theorists—what do you make of this?" The group also discovered that the pressure distribution in the channel was nonlinear. In a big pipe, like a gas main, the pressure is high at the inlet, drops at a steady rate—linearly—as the gas flows down the pipe, and reaches its lowest value at the outlet. This pressure drop forces the gas through the pipe, just as an elevation drop forces water down an aqueduct. But in the microchannel, the pressure

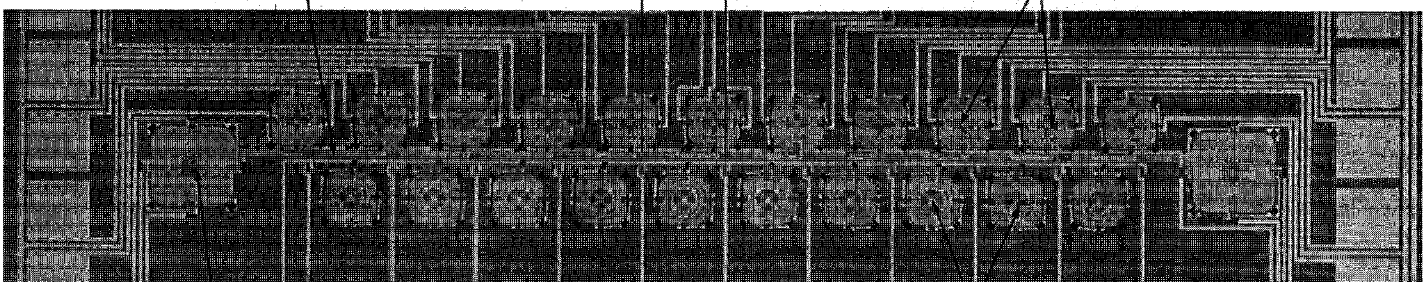

Above: The uniform—flow micro wind tunnel. The channel is 4.5 millimeters long by 40 microns wide by 2 microns deep. The structure at one end is labeled "Gas Inlet/Outlet" because the tunnels are designed to accommodate flow in either direction. The "dummy" sensors along the bottom side of the tunnel provide for leak checks during the fabrication process. Opposite page: The section where the newest wind tunnel narrows from 100 to 40 microns. Portions of three pressure sensors (two above and one below the tunnel) can be seen, as well as the very narrow channels that feed them. The sensors are cavities beneath thin diaphragms that flex as the pressure changes. These distortions are measured by the zig-zag structures visible on the diaphragms. The thick, light-gray stripe down the center of the tunnel is an electrical lead.

didn't drop very fast in the first portion of the pipe, which may indicate that the gas molecules are clogging up the channel. "There are ideas as to why this should happen, but the bottom line is that we still don't understand the physics yet."

The lab's latest wind tunnel has three choke points in it. It starts off 100 microns wide, narrows to a 60-micron-wide throat, then expands back to 100 microns. Later on, the tunnel funnels down to 40 microns, and then, later still, there's an 18-micron-wide neck in the 40-micron channel. There's a micropressure sensor before, after, and near each choke point, as well as at the channel's inlet and outlet. "We see even stranger things in the nonuniform channel. We're more puzzled there than the day we started! That pretty much sums up the current state of our research." In addition to the high-pressure region in the early part of the channel, there's a pressure drop at the first and third bottleneck, which might confirm that the system doesn't have enough oomph to force many gas molecules toward the outlet. There's also a pressure spike at bottleneck number two, where the channel narrows to 40 microns. This could mean that the gas molecules pile up like a mob of Keystone Kops running full tilt at a narrow doorway.

Tai is now expanding these studies into liquid flows. Liquid flows will no doubt act even odder, because liquids are more viscous than gases. For one thing, it will be harder to force a liquid through a microchannel, which means that the channel will have to resist substantially greater stresses. Fortunately, we've already seen that microstructures actually get stronger as they get smaller. And if the knowledge needed to design sturdy, efficient microplumbing systems emerges from Tai's research, then hand-held blood-sample screening devices become more plausible, which brings us back to Dr. Fujimasa and to Lyndon Johnson's grandma.... □

Yu-Chong Tai earned his BS in electrical engineering from National Taiwan University in 1981, and his MS and PhD in electrical engineering from UC Berkeley in 1986 and 1989, respectively, inventing the first electrically spun micromotor along the way. He came to Caltech as an assistant professor in 1989.

Numerous people have contributed to the work described herein: Amish Desai, Raanan Miller, Wei-Long Tang, Viktoria Temesvary, and Shu-Yun Wu to the microactuators; Charles Grosjean, Fu-Kang Jiang, Chang Liu, and Tom Tsao to the flap projects, with Bhusan Gupta and Sarah Bates of Goodman's lab; John Wright and Svetlana Tatic-Lucic (MS '90, PhD '94) to the waveguides, with JPL's Bruce Bumble, Henry LeDuc, and William McGrath; Wright and Tatic-Lucic to the neuron projects, with Hannah Dvorak and Michael Maher in Pine's lab; Michael Debar, Grosjean, and Wen Hsieh, to the microproperties studies; Jian-Qiang Liu (PhD '95) and Xing Yang to the wind tunnels. Tai's UCLA collaborators are Chih-Ming Ho, Jin-Biao Huang, T. S. Leu, John Mai, Kin-Cheok Pong, and Steve Tung. Tai's work is primarily funded by the Advanced Research Projects Agency, the Air Force Office of Scientific Research, the National Institutes of Health, the National Storage Industry Consortium, and Hewlett-Packard.