You stick the pins in, and—thhhp!—the rubber just seals itself to them. This is a huge advantage, says Unger. "Imagine trying to epoxy a glass capillary the size of a grasshopper's shin onto a hole the same size—that's what people used to have to do."

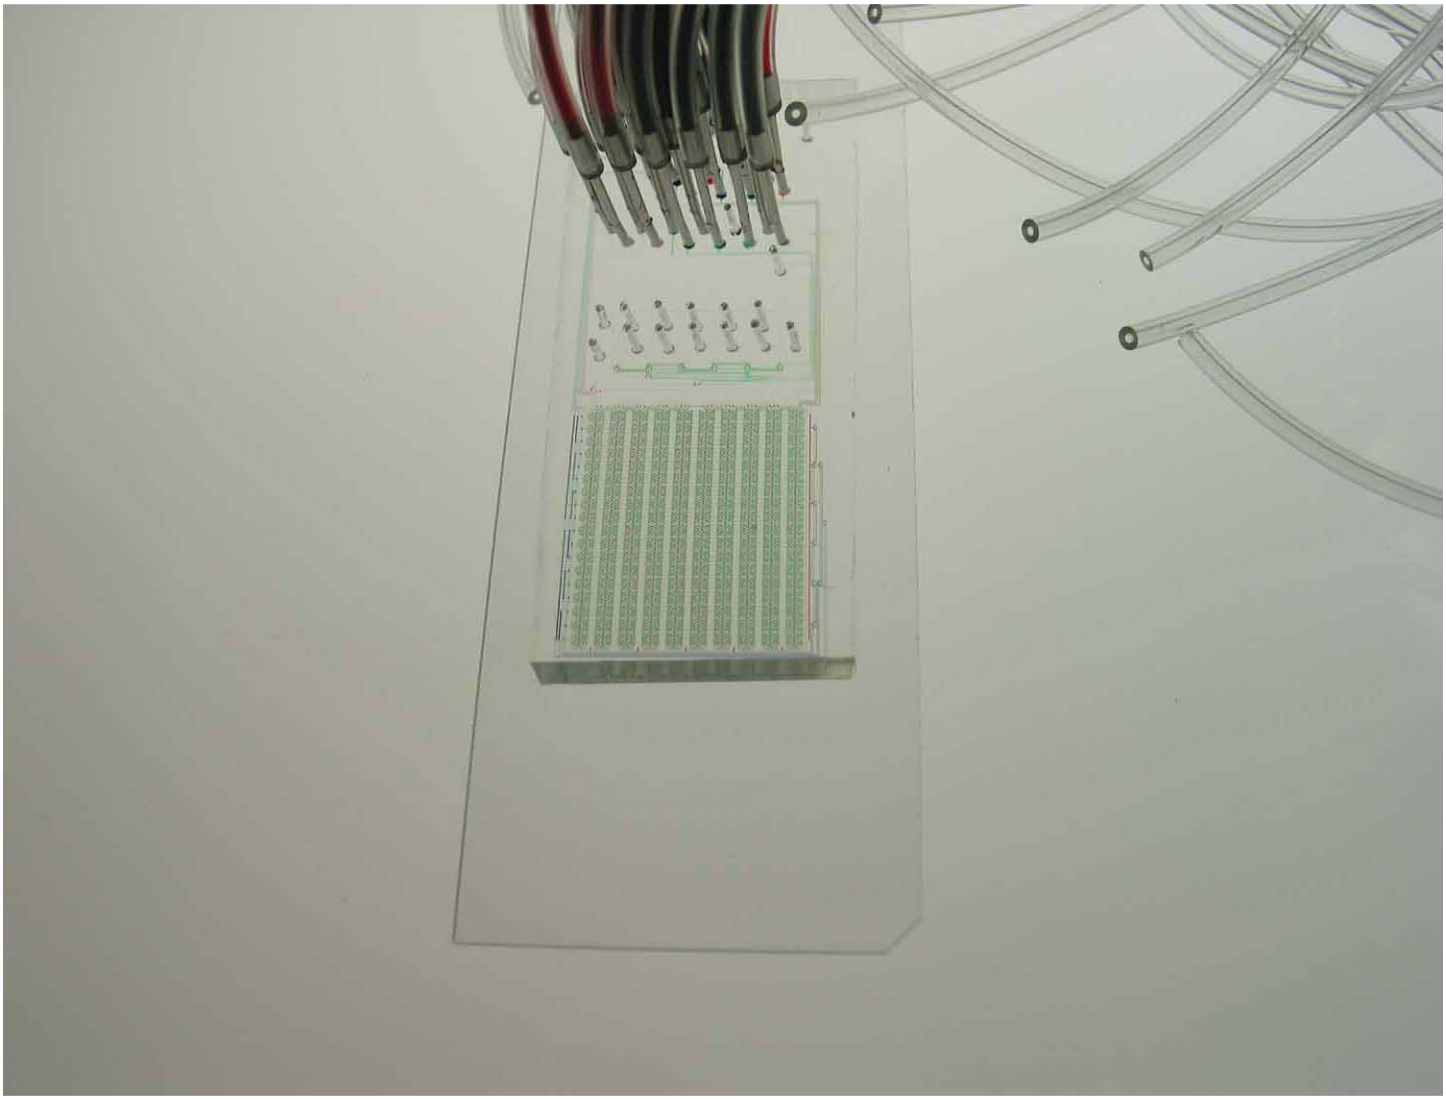

Rubber Layered Micropumpers

by Douglas L. Smith

Left: The very latest in protein-chemistry chips can handle 720 samples at once.

When you see the headlines—"Fat Gene" Found! DNA Solves Decades-Old Murder! Biotech Miracle Drug Announced!—you might think that biology has "arrived." Not so. By analogy to computer science, "biology is in the vacuum-tube stage," says Stephen Quake, associate professor of applied physics and physics. "An automatic genome sequencer or drug-discovery system fills a room, and requires a bunch of technicians to monitor it. It's roboticized large-volume fluid-handling, roughly equivalent to a vacuum-tube computer." So Quake and Axel Scherer, the Neches Professor of Electrical Engineering, Applied Physics, and Physics, are creating biology's equivalent of integrated circuits—the silicon brains in your PC, albeit not quite that sophisticated yet. Computers can process reams of data in parallel, to look for comparable gene sequences in different species, for example, but there's no way to do the lab work on even a remotely similar scale. It's all in the plumbing—dispense and mix, dispense and mix, over and over and over and over again—and, without the fluid equivalent of a number cruncher, "most biology students spend their career pipetting all day long," says Quake. "We're trying to free them for higher-level tasks." (On the consumer side, a "lab on a chip" the size of a flip phone could analyze the proteins in a saliva sample and tell you whether you have the flu or just a bad cold.)

The integrated circuit shrank a gymnasium-filling computer to fit on a fingernail. For the last decade or so, people have been trying to create integrated microfluidics, using the same technology to carve teeny-tiny pipes and build itty-bitty valves. But water (and its cargo of cells, proteins, or DNA) has proven much harder to push around than electrons. The problem is the valves—it's not called solid-state electronics for nothing. Everything is carved out of a single chunk of silicon and generally needs to remain attached to it. Imagine trying to insert a tiny gate valve into a tiny pivot hole under an electron microscope; now imagine doing it ten thousand times on a single chip without running screaming from the factory. So instead of hinged valves, people tried cantilevers—think of a pool cover that's mounted like a diving board. Explains Scherer, "Silicon is rather stiff, so to move it, as in a valve, you need to push on a rather large surface area. Otherwise, you're going to have enormous problems trying to apply enough pressure to deflect it." And the valve is going to leak if it doesn't close against a compressible gasket to form a tight seal.

"We *tried* to make them out of silicon dioxide," recalls Scherer. "Then we tried to make them out of photoresist. Then we tried to make them out of polyimide, and then in the end we realized that the way of the future was bathroom caulk." "Rubber," Quake chimes in. Actually, it's PDMS, short for poly(dimethylsiloxane), a watertight sealant used on electronic components. Liquid PDMS has the consistency of maple syrup, so you basically make a mold with the fluid channels sticking up in relief from the bottom, pour the goop in, and bake it till it sets. Then you carefully peel the rubber off and reuse the mold. This method, called "soft lithography," was developed at MIT by George Whitesides (PhD '64).

But it took three innovations to make a functioning valve. Todd Thorsen (PhD '03), now at MIT himself, began working on a basic valve structure. The sample flows through a channel in the surface of the rubber, which is sealed, channel side down, onto a microscope slide. A control channel runs perpendicular to the one containing the fluid and very slightly above it, so that the thinnest of membranes separates them where they cross. "The

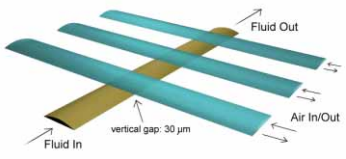

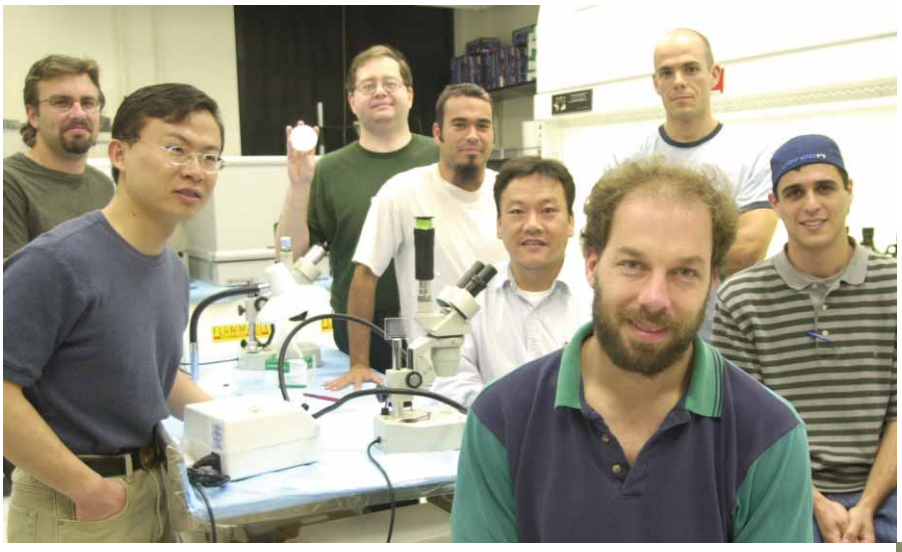

notion was that we could deflect the membrane and seal the bottom channel by applying pressure in the top channel. It's like stepping on a garden hose," says Quake. Making the control channel proved baffling, however, until Marc Unger (PhD '99) realized that each channel could be made in its own layer. PDMS comes as two components that have to be mixed, so Unger cast one layer with an excess of Component A and the other with too much Component B, cured them individually, and then sandwiched them together. A second heating then fused the two layers as the leftovers reacted.

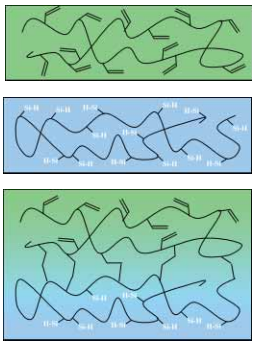

"Then," says Quake, "we couldn't get the valves to close all the way. And Hou-Pu Chou [MS '96, PhD '00] had a key insight, which was to fabricate rounded channels instead of square ones." Step on a big tin can with the top and bottom removed and it squashes flat; step on a one-gallon plastic milk jug and the corners tend to keep sticking up.

The molds are created with standard chipmaking techniques. You start with a blank silicon wafer, to which is applied 10 microns of a resin called photoresist, which will form the channels. (A micron is a millionth of a meter, about the thickness of the aluminized skin of a birthday balloon.) To ensure a nice, even layer you spin-cast the resin, pouring it onto the rotating wafer's center and letting centrifugal force do the rest. The faster the spinner, the thinner the layer—to as thin as one micron, with very precise control. (Ironically, this enabler of advanced technology is a dead ringer for a portable phonograph from about 1967. Remember 45s, man? Groovy.) A mask printed by a laser printer supplies the channel

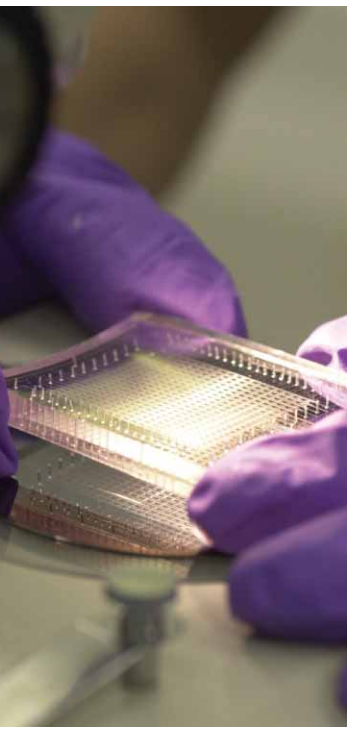

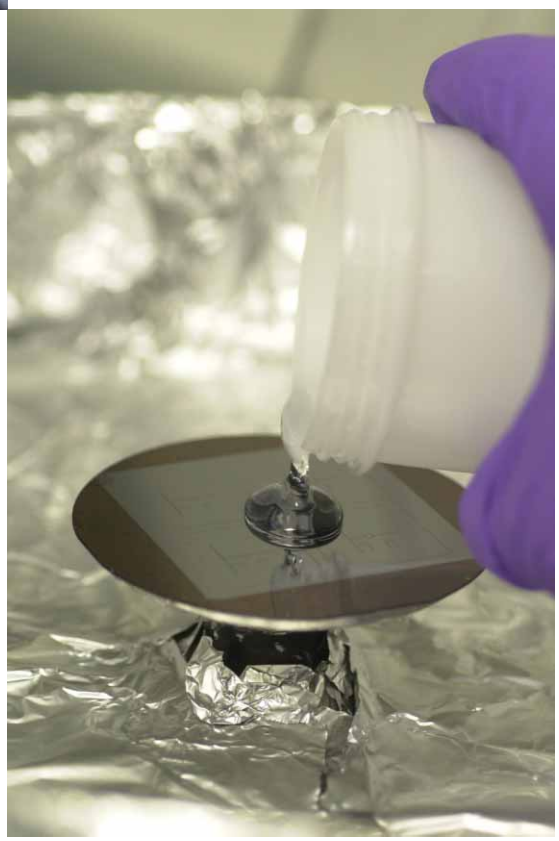

Top left: The microfluidics portion of Quake's research group. From left: grad students Michael van Dam, Jian Liu, Emil Kartalov (BS '98, with wafer), and Sebastian Maerkl; postdoc Jong Wook Hong; Quake (foreground); grad students Carl Hansen (background) and Joshua Marcus. Above: Pouring the goop on a mold before revving it up. Aluminum foil lines the spin-caster's turntable well, for obvious reasons.

Left: The rubber layers really do flex!

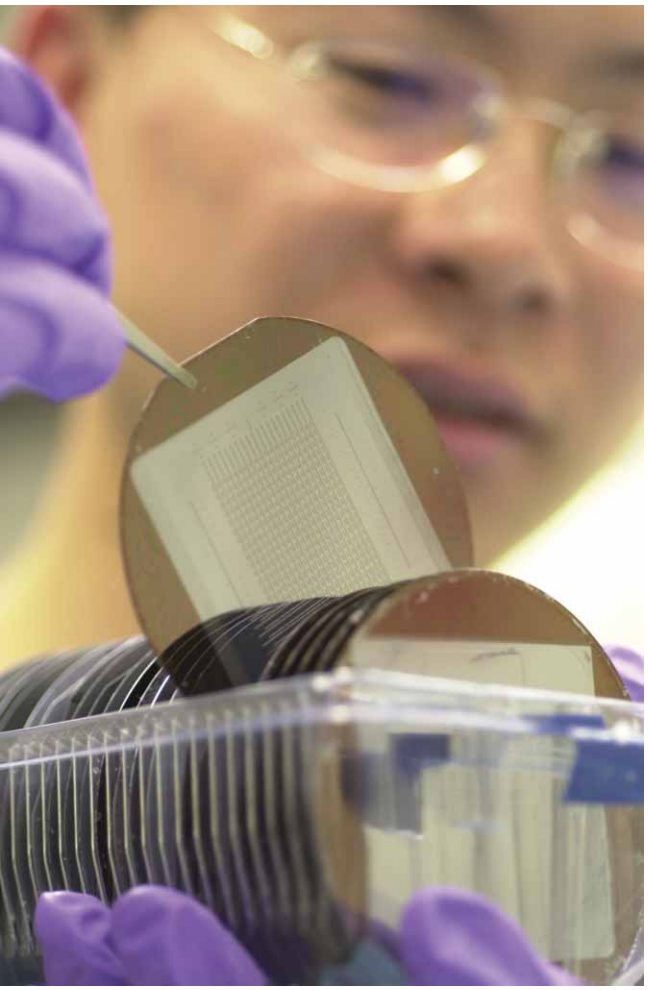

Ironically, this enabler of advanced technology is a dead ringer for a portable phonograph from about 1967. Remember 45s, man? Groovy.

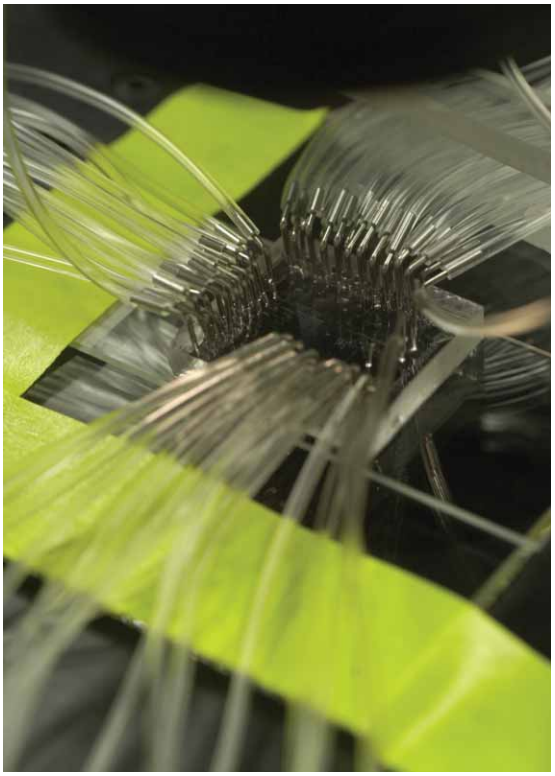

Above: Liu plays disk jockey, selecting a mold from the collection. The lab has enough of them to stock several jukeboxes. Right: A working chip, with all its fluid and control lines plugged in.

pattern in a process called photolithography; the resin is exposed to ultraviolet light through the mask, "developed," and rinsed to leave only the raised, square-sided lines of hardened photoresist. Then a quick heat treatment softens the photoresist just a tad, rounding the lines' corners. The fluid-layer rubber, which is perhaps 20 microns thick, is also spin-cast, but the control layer, which can be half a centimeter thick, is just poured by eye.

After curing, the two layers are aligned under a microscope before their second baking seals them to each other and to the slide below. Hollow steel pins—the same stock used for syringe needles—form the completed chip's connections to the outside world. You prepunch the pinholes in the control layer before making the sandwich; holes going into the fluid layer are punched through the assembled stack. Then you stick the pins in, and—thhhp!—the rubber just seals itself to them. This is a huge advantage, says Unger. "Imagine trying to epoxy a glass capillary the size of a grasshopper's shin onto a hole the same size—that's what people used to have to do." And aside from the mold making, which is best done in a clean room, "it's technology you could do in your garage," says Scherer. Assuming, of course, that there's room among the half-finished projects on your workbench for a record player, a microscope, and a small oven.

Besides not needing a high-tech vacuum chamber and a good eye with the epoxy, rubber chips have several critical advantages over silicon. You can do the whole process in a day, from designing the masks to testing the product, so it's easy to evolve designs. Or, you can reuse the same mold indefinitely, says Quake, "until you drop it and crack it." But most important, PDMS is gas-permeable—as the channels fill, the trapped air just seeps away. On a silicon chip, every dead end needs a vent line, and you can *still* wind up with channel-clogging bubbles. And caulk is cheap—

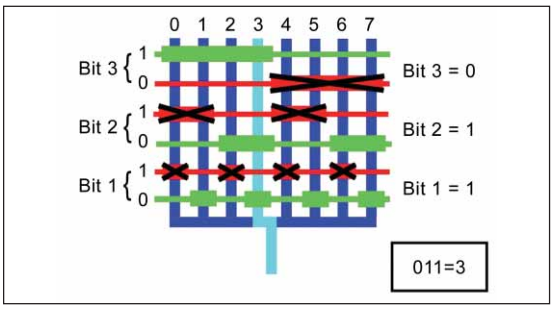

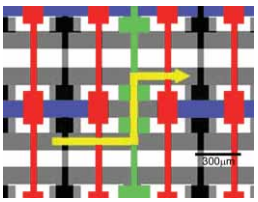

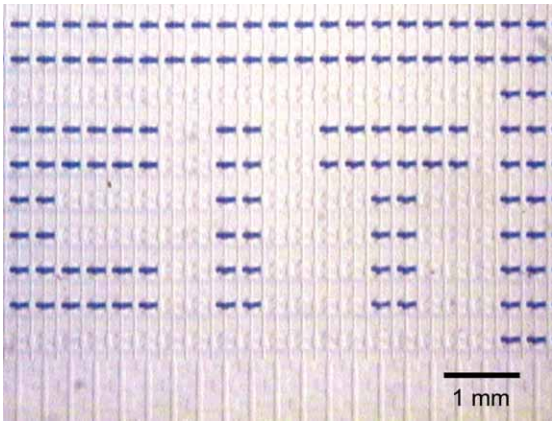

Clockwise, from the top: A. How a multiplexor works. The red and green lines are the control lines, with the red lines under pressure and the Xs marking the closed valves. The blue lines are the fluid lines, with the light blue one (number three, binary 011) being the only one open. In general, *n* fluid lines can be worked by $2\log_2 n$ control lines. B. To demonstrate selective addressing, blue dye was loaded into the memory chip and then individual chambers were purged with clear water to spell out CIT. Each chamber holds about 250 picoliters. C. The entire memory chip can be loaded (blue) in one shot by opening the red valves. To retrieve a sample, the row multiplexor sends pressurized water (yellow arrow) into the fluid line (gray) below the desired sample row, and the column multiplexor opens the green valves above and below the proper chamber.

about 50 times less expensive than silicon. So you could crank popular chips out by the truckload, but making custom ones isn't prohibitively expensive either.

Recalls Scherer, "Once we developed a valve and a pump, Steve ran with it." (A pump is just three sequential valves, opened and closed in the proper order.) "You can do a *lot* with two layers," says Quake. "However, we've shown that we can do up to eight, just by alternating A and B. I don't think there's really much of a limit." Adds Scherer, "It's just a matter of aligning them on top of one another." Two layers are enough to make chips that can store or process many subnanoliter samples at once, in layouts that rather resemble their silicon counterparts. (A nanoliter is one billionth of a liter—about one-thousandth the size of a sneezed aerosol droplet.)

A single chamber can have several valves, so if each valve needed its own control line, the plumbing nightmare would seriously limit the number of chambers that could be put on a chip. One control line can shut many valves at once, which simplifies things. But if you want to shut a specific valve in the grid's interior, the control line may have to cross many fluid lines you don't want to affect. Fortunately, it's easier to make a wide channel bulge than a narrow one, so the control lines look like piano keys laid end to end, with the wide parts being the valves and the narrow parts merely crossovers. This ability to step on some hoses while striding over others is the key to managing complexity.

Even so, as you scale up the grid, the number of valves quickly gets out of hand. Quake and his cohort designed a multiplexor that allows all the valves in the grid to be controlled by a handful of valves on the periphery. A computer uses binary numbers—strings of ones and zeros—to "address" specific locations. The multiplexor does the same, except that it needs two control lines per digit. The first line represents the "one" state, in which, for example, all the even-numbered valves are closed. The other line represents the "zero" state, in which the even valves are open and the *odd*-numbered valves are closed. As a demonstration, Thorsen, grad student Sebastian Maerkl, and Quake cast a 1,000-chamber memory chip—a 25 × 40 grid—addressed by a mere 20 lines. By sending the appropriate pair of binary numbers to its row and column multiplexors, you can fill or flush any desired chamber without disturbing the others.

The trio also built a prototype 256-unit processor consisting of four pairs of columns of 64 chambers each. The contents of the chambers in adjoining columns get mixed pairwise, and the result from any one pair can be pumped out. As a test, one column was loaded with *E. coli* bacteria containing a mutant enzyme, at a bacterial density such that there was, on average, one bacterium every five chambers. The other column was loaded with a dye that, when oxidized by the enzyme, fluoresced bright green. By draining only the chambers that lit up, the mutant cells were collected in a highly concentrated solution. (An earlier cell sorter built by Anne Yen-Chen Fu, PhD '02; Charles Spence, PhD '02; Frances Arnold, the Dickinson Professor of Chemical Engineering and Biochemistry; and Quake used a T-shaped channel with a valve on each arm of the crossbar. Fluid was pumped up the T's leg, and the fluorescing cells were diverted one by one into the proper arm by opening and closing the valves.)

Besides checking for biological activity or concentrating samples, the processor can also split them up—perhaps dividing a diverse cellular stew into tiny subsamples that can be analyzed independently. ("Simplification by partitioning," Quake calls it.) It can also do chemical reactions in parallel, including "combinatorial synthesis," in which you mix and match, say, amino acids to make all possible protein sequences of a given length at once. In fact, grad student Michael van

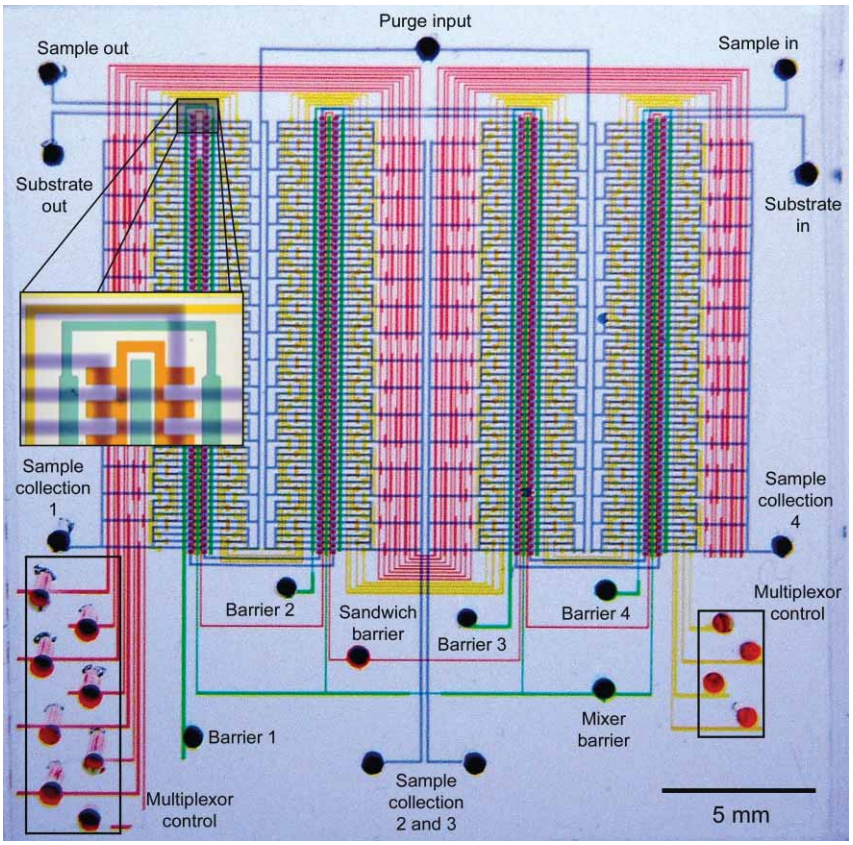

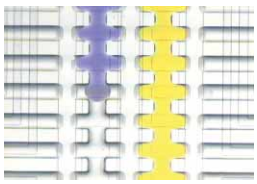

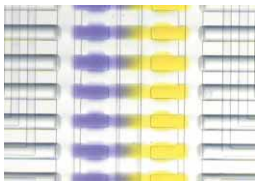

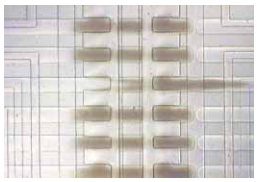

Right: The plumbing diagram for the processing chip, photographed by injecting food coloring into the various lines. ("Substrate" refers to the material the samples are going to react with, and the numbers identify the column pairs.) Far right: With all the vertical valves closed, a sample column is loaded with blue dye and the adjoining substrate column with yellow (top). The barrier valves separating the two columns are opened, and the dyes mix (middle). The product from any given reaction pair can be purged to the sample collector (bottom).

Dam wants to make a universal gene-detection chip that would contain samples of all possible single-stranded DNA sequences. When the gene you're looking for gets turned on, it would start cranking out RNA copies that would bind to the complementary DNA somewhere on that chip. But to conclusively isolate a gene, the DNA would have to contain enough letters so that the RNA only binds to one sequence. Depending on the complexity of your organism, this number ranges from 10 to 16 letters, or 1 million to 64 million sequences—rather more chambers than can be put on a chip at the moment, but perhaps attainable within the lifetime of a grad student.

It may come as no surprise that a start-up company has been formed; Fluidigm's first product, a protein crystallizer, hit the market in March. Proteins are a cell's molecular machines, but what a protein does—or fails to do—depends on the structure's excruciating details: one hydrogen atom out of place can kill it. And the best way to determine a protein's precise 3-D structure is by X-ray diffraction, which requires a high-quality crystal about 100 microns on a side. But there's no way to predict the conditions under which a protein will crystallize, so trial and error is the order of the day. Finicky is the word—crystallization frustration is the leading cause of hair loss among structural biologists, not to mention carpal tunnel syndrome from all the pipetting.

Fluidigm's design, based on one by grad student Carl Hansen; postdoc Emmanuel Skordalakes and professor James Berger, at Berkeley; and Quake, has 48 units, each of which can be loaded with a different set of crystallizing reagents. Further, each unit contains three pairs of mixing chambers of assorted sizes to give a range of mixing ratios. When you open the valves separating each chamber pair, the contents mix by diffusion. This is how crystals grow on the space shuttle, but it's well-nigh impossible to do on Earth because any sample much larger than these falls prey to convection, whose turbulent motion can jar the protein molecules out of solution into a noncrystalline glop. The slower the growth, the better the crystal, and gentle diffusion lulls the protein into remaining in solution long after it should have fallen out. It's like Wile E. Coyote running off a cliff—as long as he doesn't look down, he can keep going. Sometimes the chips even grow beautiful, diffraction-ready crystals under conditions that give glop in conventional experiments. And they do this with minuscule amounts

A protein crystal. If you see one you like, just slice open its chamber, suck it out with a micropipette, and pop it in the X-ray diffractometer.

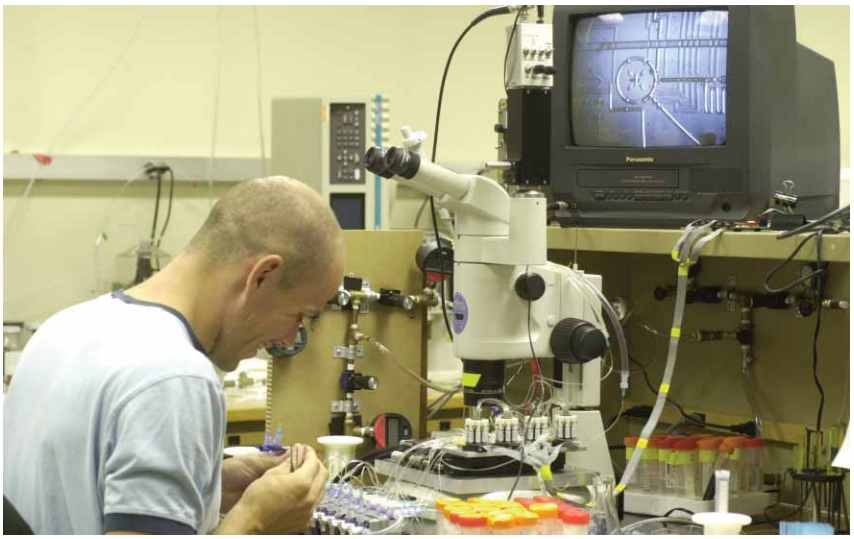

Right: Hansen at a microfluidics lab station. The chip is under the microscope, whose view is displayed on the monitor. Far right: The clustered cylinders that look like firecrackers are computer-driven controllers, developed by Fluidigm, that provide compressed air to pressurize the water in the chip's control lines. The array of white-handled valves in the foreground supply the fluids the chip is processing.

of protein—three microliters will supply all 144 experiments on the chip.

So—does anyone outside the world of biotech care? Well, the cops might. Grad student Jian Liu, then-postdoc Markus Enzelberger, and Quake have developed a potentially handheld PCR reactor. PCR stands for polymerase chain reaction, which allows you to make millions of copies of a single piece of DNA quickly and easily and which won Kary Mullis the 1993 Nobel Prize in chemistry. Conventional PCR machines are as big as toaster ovens and use microliters (millionths of a liter) of fluid; depending on the procedure used, one complete cycle can take from a few minutes to a couple of hours. (The first cycle yields one DNA copy; the second, four; then eight, sixteen, and so on.) It takes 30 cycles or more to get a usable amount of DNA from a single drop of blood, and Caltech's chip, which used a record-setting 12 nanoliters of sample, can run at about 30 seconds per cycle. Thus a readout could be ready in 30 minutes or so, far less time than *CSI*'s Catherine Willows spends at the average homicide. And PCR is morbidly sensitive to cross-contamination, so a sealed "lab on a chip" you could take to the crime scene, use once, and discard would make positive matches much more positive. The coroner's office could save some big bucks into the bargain—PCR reagents are very pricey, so consuming them by the bottle cap instead of the bottle would make the budget go a lot farther.

The design is based on a ring-shaped mixer developed by Chou, Unger, and Quake back in the early days. As the liquid courses around the circle, it passes over tungsten heating elements set to the proper temperatures. (PCR methods vary, but there are two or three steps that run at different temperatures, including one near boiling.) In the current design, the reagents swirl with the sample. But if the DNA polymerase—a heat-sensitive enzyme—could be confined to the chip's middle-temperature region, the reaction could use faster polymerase strains that are even less stable when heated. In fact, pretty much any medical and most biotech applications you can imagine, like van Dam's gene detector, would benefit from being able to attach proteins, DNA, or what-have-you to the chip. This can be done with avidin, a protein found in egg whites, and biotin, a growth factor—also known as vitamin B_7—that comes from the yolk. Avidin and biotin bind strongly and exclusively to each other and, says Quake, "there are tons of enzymes and other proteins that have been 'biotinylated,' and you can biotinylate DNA molecules. So if you have a way to attach avidin to a surface, you can catch all these things. It's like the Krazy Glue of biology." It works the other way, too—you can put biotin on the glass

The layout plan (right) and assembly diagram (far right) for the PCR chip. The red line is the fluid channel, which can be made in varying widths so that the sample lingers for the correct time over each heater (blue). Liu designed the S-shaped pumps (yellow) after noticing that a control line inflates from one end to the other, like those long, thin balloons used to make balloon animals. One S thus does the work of three parallel lines pressurized in sequence, helping reduce the plumbing's complexity.

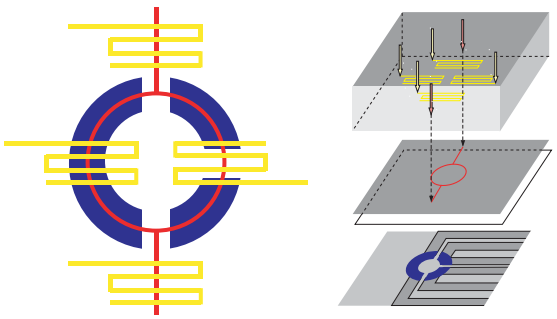

(or silicon) to affix avidin-anchored antibodies. Either way, you just make a rubber layer whose channels take the avidin or biotin to where you want to attach it. Once it's bonded, you peel the rubber off and put the real chip together.

Quake, whose background is in biophysics, came to Caltech to work on ways to manipulate individual biomolecules, such as DNA strands; meeting Scherer crystallized his interest in using microfluidic chips for the job. Scherer, a solid-state physicist, came to Caltech in 1993 after eight years at Bellcore, where he coinvented a surface-emitting microlaser—essentially a five-micron-tall, one-micron-diameter tower of hundreds of semiconductor layers stacked like poker chips. When a current passes through the stack, a laser beam shoots out the top. Until Quake's arrival in 1996, Scherer was developing microlaser arrays for communications networks and, perhaps, optical computers. "Axel helped mentor me when I got here," Quake recalls. Says Scherer, "Initially, a lot of the photolithography was done in my lab." Laughs Quake, "We wore out our welcome." "They were monopolizing our optical mask aligner," Scherer shoots back. "He was overrun with grad students," Quake agrees. "So it was better to make a parallel effort," Scherer concludes, "and it's worked very well."

"The original idea was to make ultrasensitive analytical tools using single-molecule spectroscopy," says Quake. "As we started moving farther up the food chain, we split efforts—I tried to optimize the plumbing part, and Axel's been trying to optimize the sensor part, and now we're in the process of knitting them back together."

A typical sensor includes a light source and a detector—you shoot light through the sample, which either absorbs some or fluoresces. Either way, the particular wavelengths involved fingerprint the sample, and the signal strength tells you how much of it you've got. So the goal is to take a solid-state laser and a digital camera and make a silicon sandwich, with the plumbing being the peanut butter.

The goal is to take a solid-state laser and a digital camera and make a silicon sandwich, with the plumbing being the peanut butter.

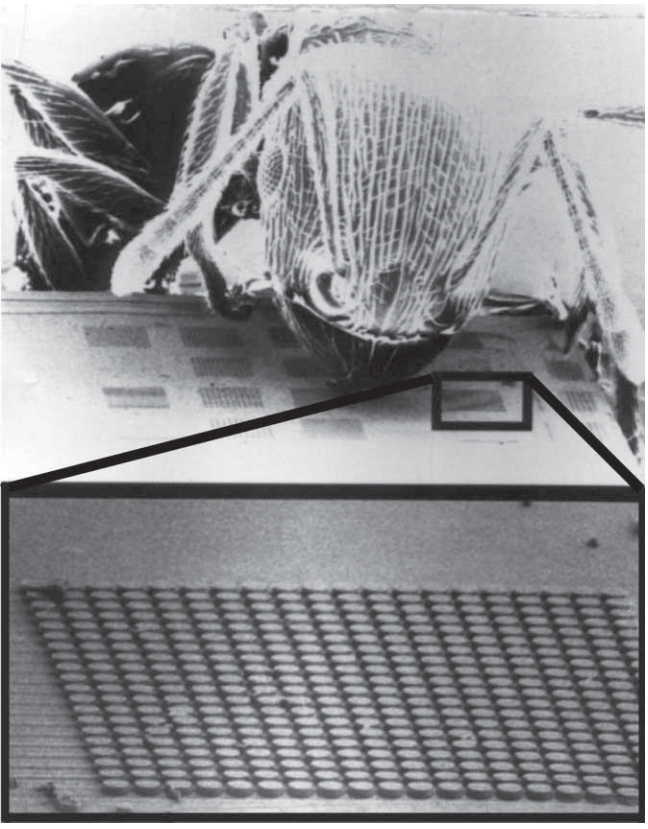

Left: An Argentine ant—those little guys about three millimeters long found in every back yard in L.A.—inspects a chip containing several arrays of Scherer's surface-emitting microlasers.

The laser technology revolves around "photonic crystals." At the turn of the 20th century, the father-and-son team of Sir William and Sir Lawrence Bragg invented X-ray diffraction crystallography, for which they shared the Nobel in 1915. As mentioned earlier, this is the method of choice for determining protein structures, and it works because an X ray having a wavelength roughly the same as the spacing between the atoms in a crystal will be diffracted by them into patterns that reveal their arrangement. More generally, electromagnetic radiation of any wavelength can be reflected, diffracted, or focused by a lattice of "atoms" of the proper size and spacing—a photonic crystal. So a properly constructed silicon wafer with islands of some other material embedded in it can trap and concentrate light into a volume 100 times smaller than a cubic

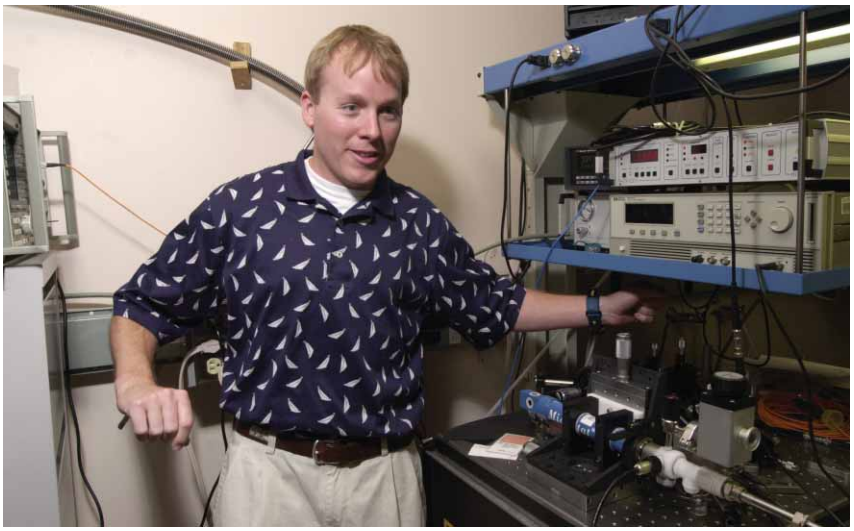

Adams with the apparatus (left) used to test the defect-cavity lasers (below).

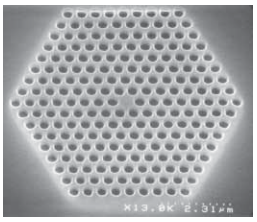

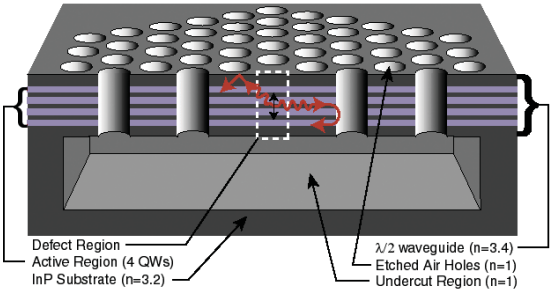

Left: A cross section through a defect-cavity laser. Light gets trapped in the waveguide's central region, because it's reflected wherever it meets a sharp change in the refractive index (n). (QW stands for quantum well, of which there is one in every red band in the active region, and λ/2 means one-half a wavelength.)

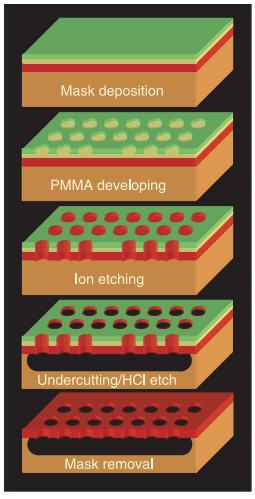

Right: Making a microlaser chip is a bit more complicated than making a microfluidic chip, but it's still all standard technology. The green layer is polymethyl methacrylate (PMMA), a photoresist that is patterned by a scanning electron microscope's electron beam. Then a highly reactive beam of fluorine or chlorine ions drills through what will be the photonic crystal (red) to the silicon base (brown). A nice acid soak then opens up the air space underneath. (The yellow layer is a second kind of mask.)

wavelength. You just make the wafer half a wavelength thick, with air above and below it. The silicon-air interface acts like a mirror, confining the light within the crystal, where Bragg reflection does the rest. (Of course, the trapping material has to be transparent, so for silicon this only works in the infrared, which is to silicon as visible light is to glass.)

Oskar Painter (MS '95, PhD '01), now an assistant professor of applied physics; Reginald Lee (MS '96); Scherer; and Amnon Yariv, the Summerfield Professor of Applied Physics, realized that it would be a lot easier to make the entire crystal out of silicon-air interfaces—all you needed to do was drill a bunch of holes in it. The resulting "defect cavity" is a hexagonal array of holes, not unlike a honeycomb, surrounding an un-drilled-out space in the center. That missing hole is the "defect," and it traps light. It's a "cavity" only in the optical sense, because the light within it behaves as if between a set of mirrors. The light resonates, amplifies, and, as with the microlasing pillars, eventually shoots out the surface. Voilà—a nice, flat laser that could be sealed to a rubber layer.

Meanwhile, postdoc Enzelberger and Scherer's grad student Mark Adams (MS '00) were laying rubber on the latest spaceflight-quality camera chips provided by Robert Stirbl at JPL's Microdevices Lab. But the narrower the channel, the shorter the path light takes through the sample and the less sensitive the sensor becomes. The simplest way to keep the sample in the beam longer is to make a hole in the cavity, redundant as it sounds, in order to collect the fluid. But would a defective defect still act as a laser? Nobody knew, and the odds didn't look good, but Marko Loncar (MS '98, PhD '03) took on the challenge. Says Scherer, "that was a two-year design process all in itself, trying to make a high-resonance cavity with a hole in it." Amazingly, it worked, and it created a third way of analyzing the sample beyond fluorescence and absorption. The fluid

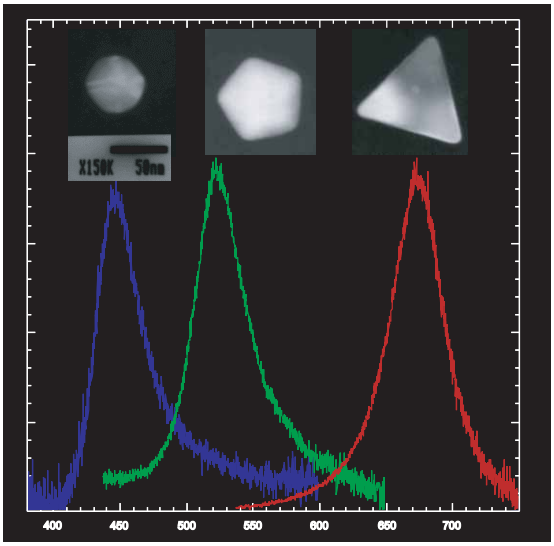

Right: Zooming in with high-magnification TEM reveals the shapes of individual particles. (The scale bar is 50 nanometers.) Each particle's visible-light spectrum is shown below it. Wavelengths are in nanometers.

Below: The plasmon particles are awfully pretty when seen by a dark-field microscope (top), but are barely visible under a transmission electron microscope (TEM) at the same magnification (bottom).

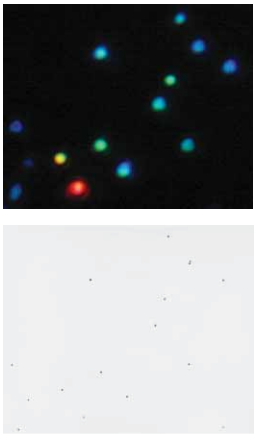

alters the laser's wavelength in very specific ways—alcohols are different from water, and proteins are different from one another. "You couldn't do this by drilling a hole in a relatively big laser, like the one in a laser pointer," says Scherer, "because there are just too many states available to the system. But here there are only a few available states, so you can deconvolute it."

Another method may work with visible light. Postdoc Mladen Barbic is experimenting with flecks of silver some 50 nanometers (about one-tenth the wavelength of green light) in diameter. Through a phenomenon called "plasmon resonance," their shapes govern the colors of light they absorb and reemit—circles turn blue, pentagons green, and equilateral triangles red. When a molecule from the sample attaches itself to one of the metal particles, it alters how the light behaves by a process called surface-enhanced resonant Raman scattering (don't ask). When you hit the metal-molecule combo with a laser, you get a spectrum containing many sharp peaks that identify the molecule, and the particle amplifies the spectrum so that even single molecules can be seen. Barbic currently makes what is essentially very small pocket change by chemical means, but the particles come out in assorted shapes and, when seen on a darkened microscope stage, look like the world's tiniest Christmas lights. He'll shortly carve them to order out of a silver layer deposited on a silicon wafer, using the brand-new, state-of-the-art clean room that Scherer and Professor of Physics Michael Roukes have just gotten built.

Quake and Scherer are close to putting the optics, fluidics, and electronics all on one chip. One needs to be clever planning the plumbing, of course, so that the only hole the fluid channel passes over is the one in the defect, but this is a minor detail. In a year or so, a rubber multiplexor could be sandwiched between a camera array and a laser array, with each laser drilled to a different wavelength. The multiplexor would shunt the sample to the appropriate lasers, and you'd have a microanalyzer. Another year to build in a processor as well, and a true general-purpose lab on a chip is born.

Meanwhile, word is getting out. Says Scherer, "Our biggest problem right now is that we've become *too* successful. We're making structures that are in high demand." "People are banging on our doors," Quake agrees. "And not just from on campus, but actually from around the world." So rather than open up a sweatshop filled with grad students, the soft-lithography fab lab is available to anyone on campus. And part of the recent Moore gift has been earmarked for a "foundry," where a full-time technician will mass-produce chips, or make them to order based on Ath-napkin doodles. Says Scherer, "We're very excited about having this technology transferred to the biologists on campus."

The current designs have fluid channels 100 microns wide and handle samples of a couple dozen nanoliters. Scherer and Quake are aiming for one-micron channels, about the size of an *E. coli* bacterium, which translates into femtoliter (trillionths of a liter) volumes. Such fine masks can be made with off-the-shelf equipment—one micron is as wide as a highway, by silicon standards. So there's plenty of room at the bottom, as Richard Feynman famously remarked in these very pages. Says Quake, "These devices obey a Moore's-law-type scaling—in fact, they beat the conventional semiconductor Moore's law by quite a bit." (Moore's law says that advances in technology allow the number of transistors, or in this case valves, on a chip to double every 18 months.) "So we can now start to count on this happening, and we should start planning what kind of devices we can make with that. On the other hand, it's worth spending the effort in technology development to make sure we stay on track." Adds Scherer, "The exciting part is that so little has been done that

Right: Scherer's and Roukes's new clean room is rated Class 100, meaning it has less than 100 dust particles per cubic foot of air. (Typical Pasadena air might contain a million particles per cubic foot; if you have an indoor air filter, you might be breathing Class 50,000 air.) The equipment is still being broken in, but the air samples are already in the Class 10 range, and they hope to get to under four. Loncar grips the access door to the e-beam writer, which can aim a 13-nanometer-diameter electron beam to 0.6-nanometer accuracy anywhere on the surface of a standard six-inch wafer, allowing you to write several successive patterns in perfect register. The entire system is mounted on its own concrete foundation pier so that people's footfalls don't jar it.

Scherer the silicon chef.

PICTURE CREDITS:
8 – Sebastian Maerkl; 9, 10 – Fluidigm; 10, 11, 14, 16, 18 – Bob Paz; 14 – Doug Cummings; 12 – Doug Smith; 13 – Carl Hansen; 15 – Axel Scherer

you can get a lot of mileage out of even small details."

Eventually, of course, they'll hit the wall—literally. The layer of water molecules next to the channel wall tends to stick to it, so as the walls get closer and closer together, the free-flowing fluid region gets narrower and narrower, and at some point the pumps will no longer be able to force the passage. This doesn't occur in the one-micron channels that have been made as demos, so grad student David Barsic (MS '01) is trying to see just how narrow a channel can be. But Shapiro's law of cell sorting says that a 49-micron cell will plug a 50-micron channel, so for some uses there's no point in going smaller anyway.

"The tools are now here," says Scherer. "But the applications are in front of us. And that will drive the development of the next generation of tools. Caltech has a lead right now, but a lot of infrastructure has to be built, and we have to invest in order to take advantage of this moment." Adds Quake, "We've taken a five-year detour in technology development, and now it's mature enough to do science. We have a *lot* of things planned. In the near term, my group plans to look at unculturable bacteria. Ninety-nine percent of what surrounds us can't be grown in the lab, and therefore is sort of invisible. It's the biological equivalent of cold, dark matter." Taking a tack analogous to the protein crystallizer, Quake will collaborate with Jared Leadbetter, assistant professor of environmental microbiology, and David Relman at Stanford to learn what living conditions these little bugs like, to try to find out what they can teach us about the spectrum of life. "And we want to look at the human body's rarest cells, stem cells and such. It's difficult to analyze them with conventional techniques, because they occur in such small numbers. But we should be able to get detailed molecular and genetic characterizations of them with integrated microfluidics." For this, he's collaborating with W. French Anderson, director of the Gene Therapy Laboratories at USC.

"Integrated circuits automated the process of computation," says Quake. "During World War II, people wanted to solve differential equations in order to compute missile trajectories. They did this with teams of people with adding machines." So ENIAC, the world's first electronic digital computer, was built at the University of Pennsylvania in 1946. Weighing over 30 tons, including its power supply and air-conditioning units, ENIAC contained 19,000 vacuum tubes and 1,500 relays, and drew about as much power as 200 households. With that, it could add, subtract, multiply, divide, and do square roots on twenty 10-digit (base-10) numbers simultaneously, and there was much rejoicing. Then the transistor came along, followed by the integrated circuit and eventually the PC revolution. "And all of a sudden people realized that automated computation was not just useful for solving math problems, but could be used for word processing, spreadsheets, e-mail, the World Wide Web, and Tomb Raider. Nobody anticipated that when they started this program of automating math. In comparison, our lab is now in the '70s. We have specific large-scale integrated circuits for certain tasks, but we don't yet have a general-purpose programmable microprocessor." But with Moore's law holding sway, the '90s aren't far off, and who knows what the fluidic equivalent of a Pentium will bring? □