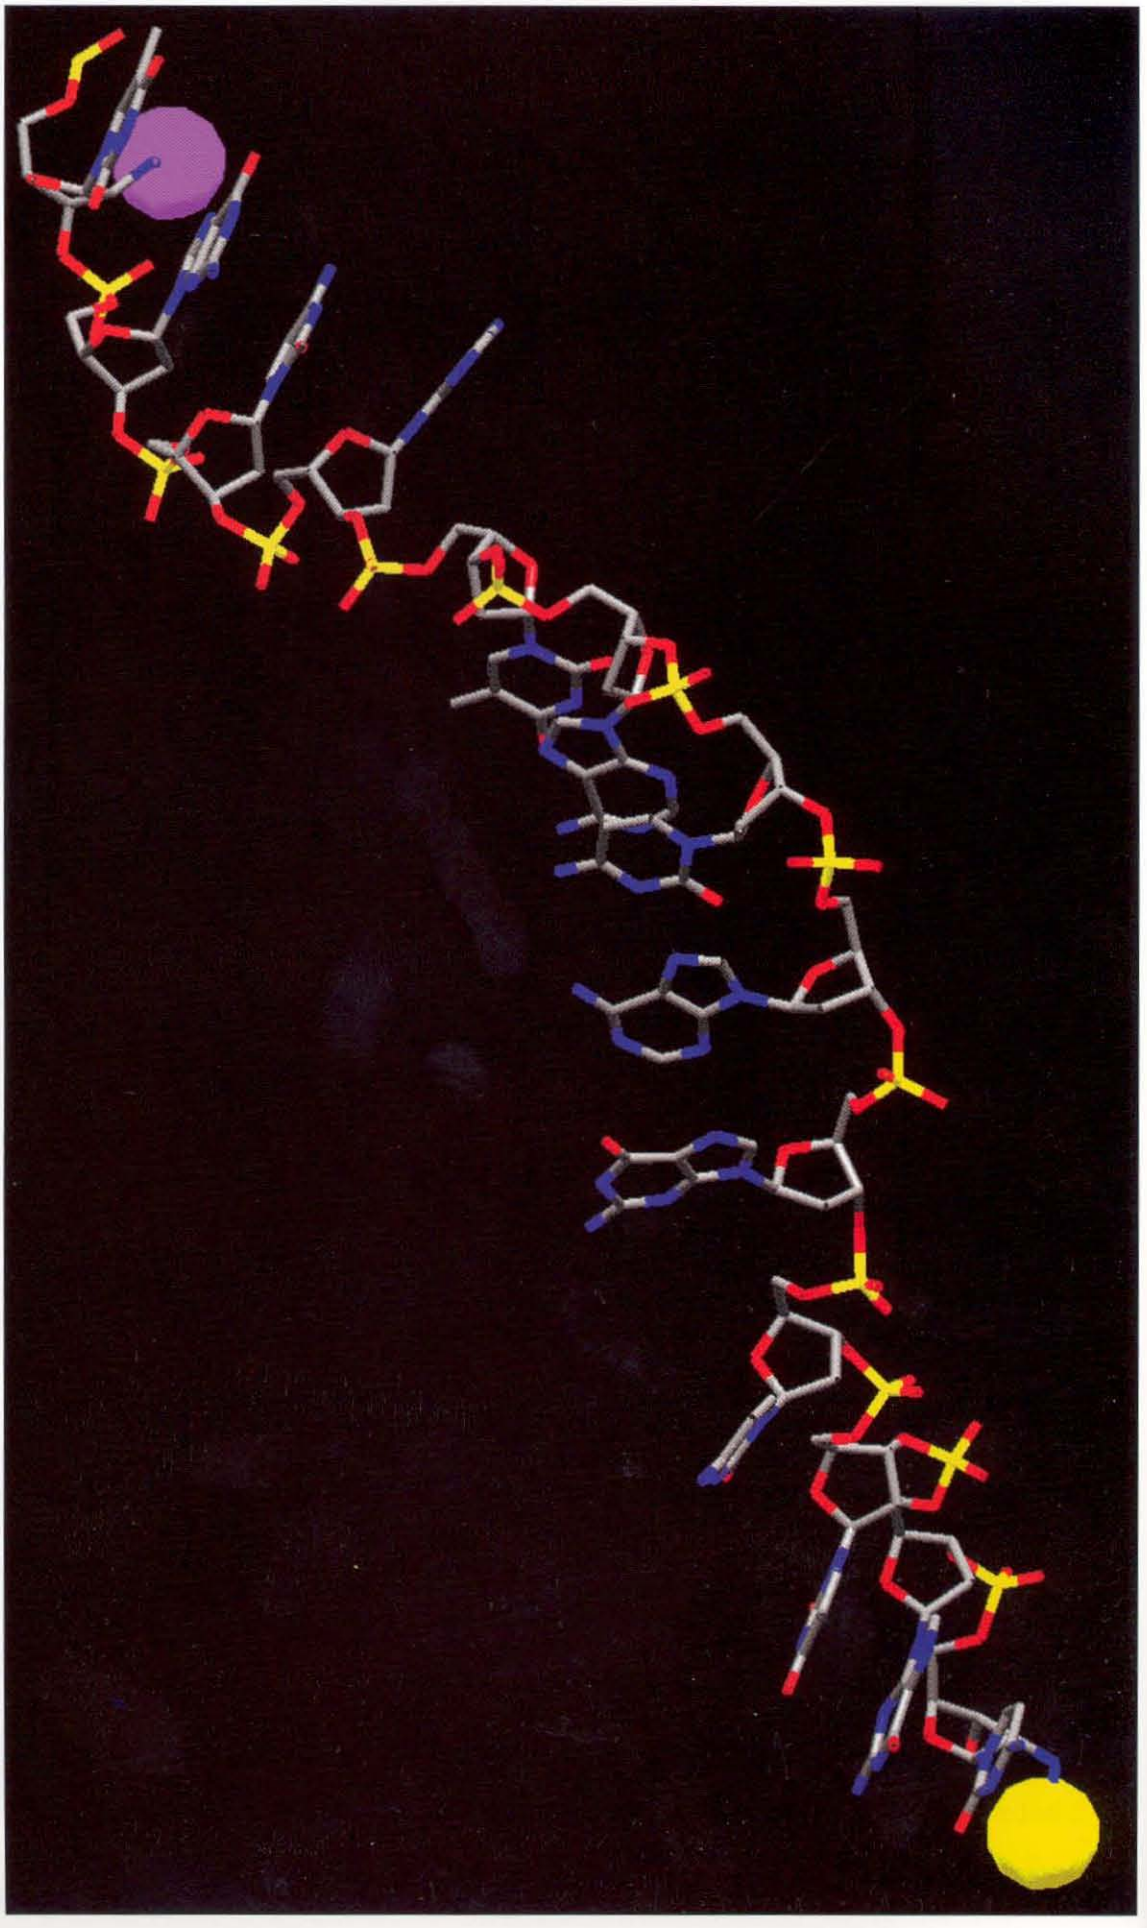

Heart Attack or Heartburn: New Chemical Diagnostics That Make the Call

If the creators of Star Trek *can build a medical tricorder capable of instantaneous diagnosis, why can't we?*

by Thomas J. Meade

Two-stranded DNA likes to coil up into the double helix whose shape has become emblematic of the Age of Biotechnology. But a single strand of DNA, like the one shown on the opposite page, flops about loosely like an overstretched telephone cord that's lost its curl. This lackadaisical behavior has pronounced effects on DNA's ability to conduct electricity, as Meade's lab has discovered. The researchers timed the flight of electrons from a ruthenium atom (yellow) at one end of the strand to a second ruthenium atom (purple) at the other end. A device that recognizes DNA strands by the way they conduct electricity could be developed for medical diagnoses.

An accurate diagnostic tool should, ideally, make the correct decision the first time every time. Consider a baseball game, where every call the umpire makes is a diagnosis. He has to make a rapid, accurate judgment based on limited information, and sometimes he's wrong. Do these diagnoses really matter? Apparently so. On August 10, 1995, the St. Louis Cardinals came to town to play the Dodgers. Moments after batter Raoul Mondesi and manager Tommy Lasorda got sent to the showers for arguing a called strike, about 4,000 fans who agreed that the ump had erred began throwing their souvenir baseballs (it was Ball Night at Dodger Stadium) onto the field. Los Angeles wound up forfeiting the game. While this type of diagnosis may ruffle the feathers of a few fans, medical diagnoses have a more far-reaching impact—a mistaken diagnosis can lead to something far worse than an early shower.

In my research group, basic science is being harnessed to develop the chemical reagents needed to create the next generation of vastly more sensitive and discriminating diagnostic tools—tools that will give us a lot of detailed information about a patient in a cost-effective manner. Ideally, we'd like to make these tools portable, so that instead of you having to visit a huge, expensive machine in some clinic somewhere, the machine would come to you. Hollywood, of course, has gotten there before us, but if the creators of *Star Trek* can build a medical tricorder capable of instantaneous diagnosis, why can't we?

This work is a direct consequence of the collision of disciplines that the Beckman Institute, where I work, was designed to foster. The Beckman Institute houses people from diverse fields—cell biologists, developmental biologists, chemists, physicists, and engineers—all contained in one building, so that we're constantly tripping over one another. For example, my group, which is part of the Beckman Institute's Biological Imaging Center, consists of nine people. No two of them have degrees in the same field. I was born and raised in a large chemistry department, where the inorganic chemists didn't talk to the organic chemists, let alone the biologists. But this is a special place. It's a very active, exciting environment, where research projects cut across disciplinary lines. Our work is emblematic of the kind of research that results when you're forced to speak lots of different scientific languages simultaneously.

I'll focus on two types of new diagnostic tools that are emerging as a result of the interdisciplinary work that is taking place here. The first consists of a new variation on Magnetic Resonance Imaging, or MRI, which is a technique that's widely used to take three-dimensional pictures of the inside of a specimen (and sometimes that specimen is you). While traditional MRI provides millimeter-scale anatomical information, our new chemical tools allow the same instrument to provide equally detailed information about the physiological and metabolic functioning of a specimen as well. The second new diagnostic technique is an entirely new way of analyzing DNA that exploits some fundamental properties of electron-transfer reactions. The method is designed to be rapid, require no

Medical X rays are the diagnostic equivalent of stone knives and bearskins—I can imagine a time when we won't have to bombard somebody with high-energy particles just to take a picture, even for dentistry.

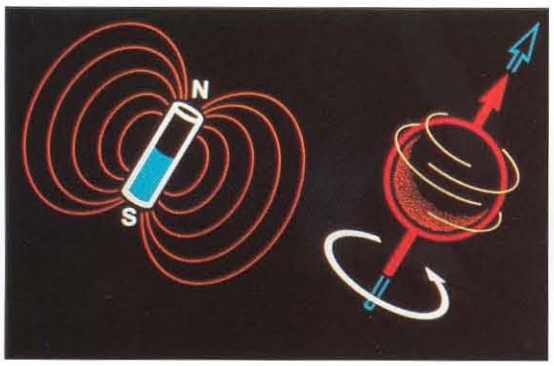

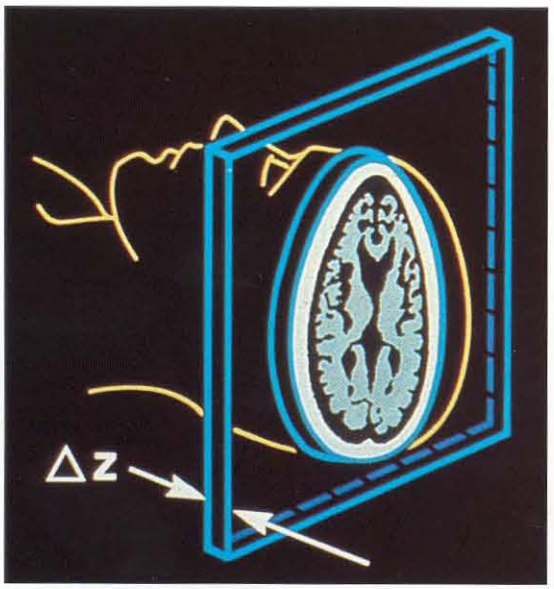

Top: A charged, spinning sphere (in this case, a proton) generates its own tiny magnetic field, as indicated by the north-pointing arrow. Bottom: So if you stick something containing a lot of protons (in this case, your head) into a larger magnetic field, they will tend to line up with that field. A radio-frequency pulse will then excite all the protons, but computer manipulation of the resulting signals allows the protons in a particular slice (Δz) of the brain to be singled out, and an image of them to be constructed.

sample purification or amplification, and is virtually automatic, unlike current methods of DNA testing.

I will begin with the MRI method. In recent years, MRI has emerged as a powerful clinical tool because it is noninvasive and nondestructive and renders visible the entire three-dimensional volume of the subject. MRI measures the differences in the local environments of all the water molecules in your body. This is an excellent way to examine the human body, which is mostly water.

MRI works because the hydrogen atom's nucleus is a single proton, whose charge and spin cause it to behave like a tiny magnet. So if the doctor sticks the patient in a large magnet, those protons will tend to align themselves with the magnetic field. An image is created by imposing one or more magnetic fields upon the specimen, while exciting the protons with radio-frequency pulses. Each pulse flips the protons' spin axes, briefly inverting them before they "relax," or flip back into their original alignment, emitting another radio signal. The rate at which the protons relax is very sensitive to their local environment in several ways. Thus, the signal intensity from a given unit of volume in the specimen is a function of the local water concentration and of two relaxation-time parameters called T_1 and T_2. Local variations in these three properties provide the vivid contrast seen in magnetic-resonance images. For example, the low water content of bone makes it appear dark, while the short T_2 of clotted blood affords it a higher signal intensity than nonclotted blood—

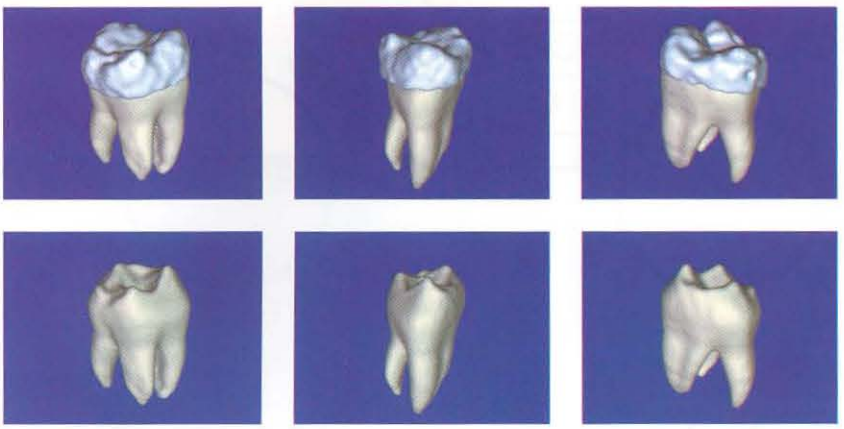

Above: Six computer renderings based on an MRI scan of an adult human molar. The top frames show the tooth's exterior: enamel (white) extends down to just below the gum line, while dentine (beige) lies beneath. The enamel has been "peeled" away in the bottom views, revealing the growth points at the dentine-enamel interface—the cusps from which new tooth tissue springs. It's normally difficult to get good magnetic-resonance images of teeth, because enamel and dentine contain very little water, but research fellow Pratik Ghosh and Member of the Beckman Institute Russell Jacobs have developed a method for taking MRIs of dry solids. Postdoc David Laidlaw, grad student Kurt Fleischer, and Associate Professor of Computer Science Alan Barr developed the image-processing software that distinguishes the rock-hard enamel from the slightly softer dentine and generates three-dimensional images of them both.

thus blood clots appear bright. Moreover, the image may be acquired in a variety of different ways that emphasize the variation in one or more of those three properties. In any case, we collect intensity data as the specimen is exposed to a variety of fields, and then a mathematical technique called deconvolution yields a one-, two-, or three-dimensional image of the specimen.

MRI scans have already replaced X-ray photos in many ways. Magnetic resonance images are much more detailed, and, unlike an X ray, you can have an MRI taken of you every day. Medical X rays are the diagnostic equivalent of stone knives and bearskins—I can imagine a time when we won't have to bombard somebody with high-energy particles just to take a picture, even for dentistry. We can take an MRI of a tooth, for example, and pictorially peel off the enamel and look directly inside.

But, as some doctors know, we can make better images for better diagnoses through chemistry. Consider a patient with a brain tumor, for example. An ordinary magnetic resonance image gives the surgeon a lot of vital information—where the tumor is, how big it is, and what brain functions might be affected by it or by its removal. However, by injecting an MRI contrast agent—basically a magnetic-resonance dye—into the patient, the surgeon can delineate that same tumor in much more detail. In the brain, where every millimeter matters, I'd certainly prefer my surgeons to have the highest-resolution images possible.

The principle behind a contrast agent is the same as that of tracing a leak in your sewer line. Your nose will tell you there's a leak, but not where it is, and the plumber doesn't want to dig up your entire yard to find it. So he brings a can of dye with him, flushes it down the toilet, and looks out in the yard—wherever the dye comes up is where the leak is. In other words, the dye says "Dig here." Standard MRI contrast agents work in a similar way—they go where the pipes go. A tumor generally has thin-walled, leaky blood vessels, so the contrast agent leaks out and pools there. Contrast agents are good for determining the type and scope of any injury in which the circulatory system is damaged—brain trauma suffered in a car accident, for example.

The area where the contrast agent congregates appears brighter than the surrounding tissue, because the contrast agent allows the protons in the neighboring water molecules to flip back into alignment with the magnetic field faster. Thus, each succeeding pulse will find those protons back in alignment, and can tip them against the field again, while the protons elsewhere may still be tipped against the field from the previous pulse. Unlike a radioactive tracer, such as a barium milk shake, you don't see an MRI contrast agent directly, but rather its effect on its neighboring molecules.

The best contrast agents contain large numbers of unpaired electrons, which cause the protons to relax by essentially siphoning off their extra energy magnetically. (T_1 becomes much shorter, in other words.) Most researchers, including us, use gadolinium ions (Gd^{+3}), which have seven unpaired electrons each—the highest number in the entire periodic table. Unfortunately, gadolinium is also very toxic—if it's not chelated, or caged, you might get a great picture, but you'd end up killing the subject. Several ways to cage gadolinium safely have been developed. These cages also limit the number of water molecules that can snuggle in the gadolinium ion's relaxing embrace at any given time, but the water molecules exchange in and out so fast that it doesn't really matter.

We have been working on "smart" MRI contrast agents that don't simply go where the pipes go and that don't simply report their anatomical location. These new agents report on the metabolic state of cells and organs in a way that shows up under MRI. This provides, for the first time, a means to obtain high-resolution, three-dimensional magnetic resonance images based on the metabolic and physiological function of living systems. We've developed two ways to make a contrast agent smarter. The simpler one, which we developed last summer, is a cell-specific reporter that's designed to seek out a specific type

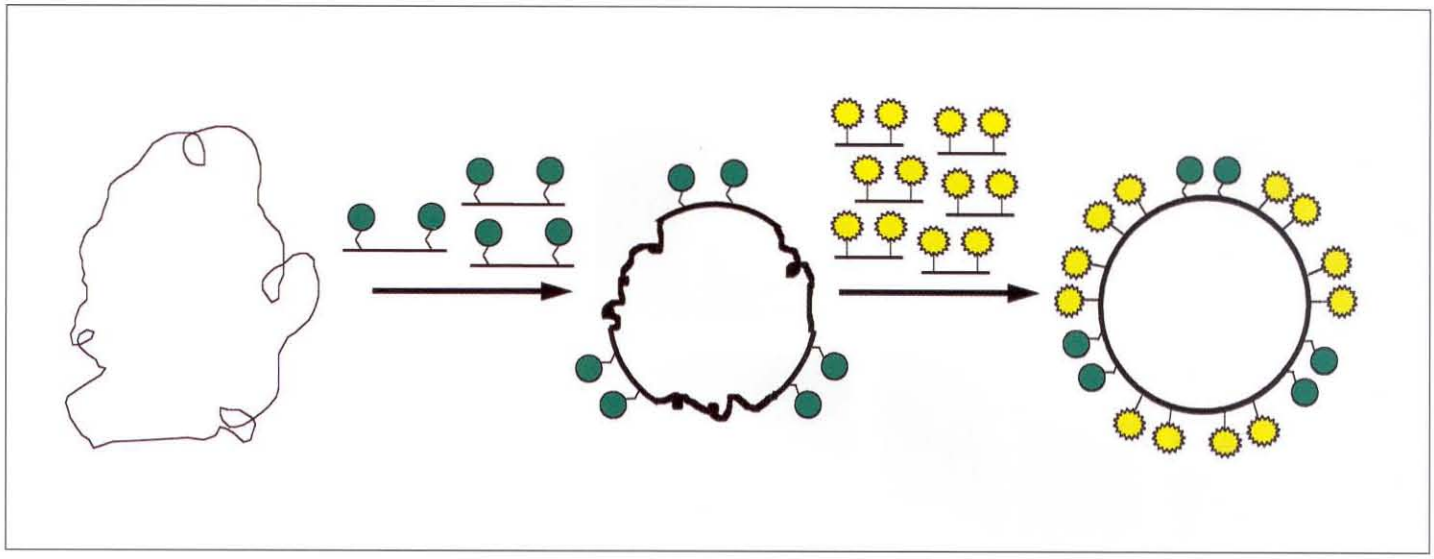

A schematic representation of the contrast-agent-delivery vehicle assembly line. The DNA backbone of the plasmid (the tangled loop at left) contains numerous negative charges, so the transferrin (teal spheres) is chemically bound to a molecule called polylysine (the short, straight line segments), whose backbone has many positive charges. A molecular version of static cling causes the polylysines and the plasmid to stick together when mixed. Not enough polylysine is used to neutralize all of the DNA's negative charges, so the gadolinium cages (gold stars) are bound to other polylysine molecules that, when added to the mixture, also cling to the DNA and soak up the rest of the charge.

of cell. Those cells, and only those cells, will light up wherever they are in the body. The second type, the functional reporter, is even smarter. It stays dark—or off—to the MRI until some metabolic or physiological event of our choosing turns it on, that is, lights it up.

Let's begin with the agent that's smart enough to find a specific cell type. In order to deliver our contrast agent reliably to a specific cellular address, we've borrowed a technique from gene therapy. Gene therapy is getting a lot of press nowadays, and the basic idea is this: Say you're missing an enzyme because your body's copy of the gene that makes that enzyme is defective. If a doctor could insert a working copy of the gene into your cells, your body would start making the enzyme, and you'd be cured. This technique has only very recently been attempted in humans, and one major technical hurdle is quantifying how much of the gene is actually getting into the target cells.

The "truck" that gets the gene into the cell is called a plasmid, which is a little ring of DNA containing the gene, and the truck's "driver" is a protein attached to the plasmid that recognizes a receptor on the cell's surface. When the protein docks with its receptor, a sequence of events is triggered that causes the plasmid to tumble into the cell like a truck into an opening sinkhole. The cell surface dimples underneath the plasmid, folds over it, envelops it, and pulls it in. So if the plasmids are loaded up with large numbers of our contrast agents before being injected into the subject, we can collect magnetic resonance images to see where the plasmids go. It's akin to having a radio-equipped moving van. We've been using an iron-containing protein called transferrin as our truck driver because it was one of the first recognition proteins used in gene therapy, but there are many other kinds of receptors that work in a similar way, and potentially we could use any one of them.

To verify that our contrast agent is going where we want it to, we did an experiment whose results are shown at the top of the opposite page. We loaded three capillary tubes with K562 leukemia cells, which have transferrin receptors. Tube A also contains the plasmid-transferrin-gadolinium particle, and the cells light up. To prove that they didn't light up for some other reason, tube B holds the plasmid-transferrin particle without the gadolinium. And, finally, to show that it's the transferrin that's actually responsible for getting the contrast agent in, tube C gets a large excess of free transferrin molecules—enough to monopolize all the transferrin receptors on the cell surface. You can see that tubes B and C remain dark in the MRI image. (The contrast agent doesn't light up outside the cells in tube C because we rinse them off before we do the MRI.)

But confirming that the gene was delivered to the right address doesn't necessarily mean we've fixed the cell. Simply because the cell swallowed the gene doesn't guarantee that it will be taken to the nucleus, where the cell's own DNA is kept, or that if it gets to the nucleus that it will work correctly in its new environment. Currently, gene therapists just dispatch the plasmid trucks and wait to see if appreciable amounts of the gene's

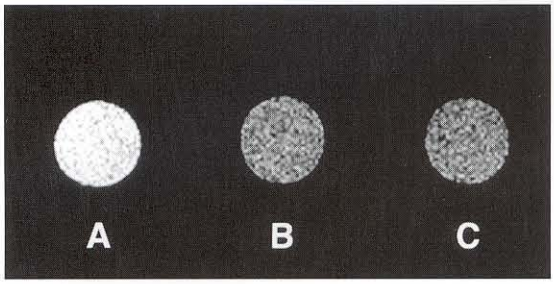

Above: The bright spots in tube A are MRI return receipts from aggregates of cells that have absorbed the smart contrast agent. Tube B has cells and plasmids but no gadolinium, and remains dark. Tube C's darkness says, "Return to sender/ Address unknown"—the tube contains the same ingredients as tube A, plus enough extra transferrin to block all the binding sites on the cells' surface and prevent the delivery from being made.

product begin to appear. It often takes weeks, or even months, to know whether the procedure has succeeded. (We knew immediately when our gene had arrived in the nucleus, because we were using the luciferase gene that puts the fire in fireflies. Our cells glowed in the dark the moment the gene started working. However, this gene is of no therapeutic value in people.)

This brings us to the functional reporters—agents smart enough to be turned on only by a physiological or metabolic event of our choosing. Such an agent would give us two pieces of information—the agent's location and the fact that the desired event has taken place there. How does this work? Remember that our magnetic dye, our gadolinium ion, makes water molecules light up, but that only a few water molecules can get to it at a time because it's in a cage. The gadolinium ion is big enough that nine water molecules could bind with it if it were uncaged, so we started with a cage design that blocked eight of the nine sites. (This cage, called 1,4,7,10-tetraazacyclododecane-N,N',N'',N'''-tetraacetic acid, or DOTA for short, is approved in France for human use, but has not yet been approved here by the FDA.) DOTA looks like a picnic basket, with the gadolinium ion inside. So postdocs Rex Moats and Andrea Staubli put a lid on the basket that blocks the water molecules from getting to the last available site. The lid's hinge is a sugar called galactopyranose, which is digested by an enzyme called β-galactosidase. Our hope was that the enzyme would still recognize the galactopyranose lid, even though it's been built into the picnic basket, and break it

Below: The cage's chemical structure, with the lid (left column) and without it (right).
In the diagrams in the top row, a carbon atom sits wherever two line segments meet. The heavy segments represent parts of the molecule that stick out in front of the page; light segments lie behind the page. The sphere in the middle is the gadolinium, which lies in the plane of the page.
In the 3-D models in the bottom row, the cage has been rotated 90° toward you to show what a water molecule sitting above it would see. White spheres are hydrogen atoms, red are oxygen, gray are carbon, blue are nitrogen, and the purple one is the gadolinium.

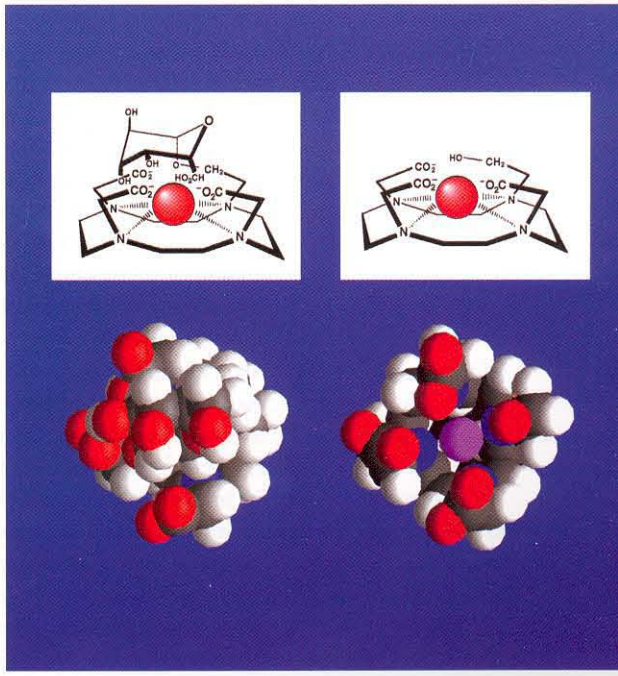

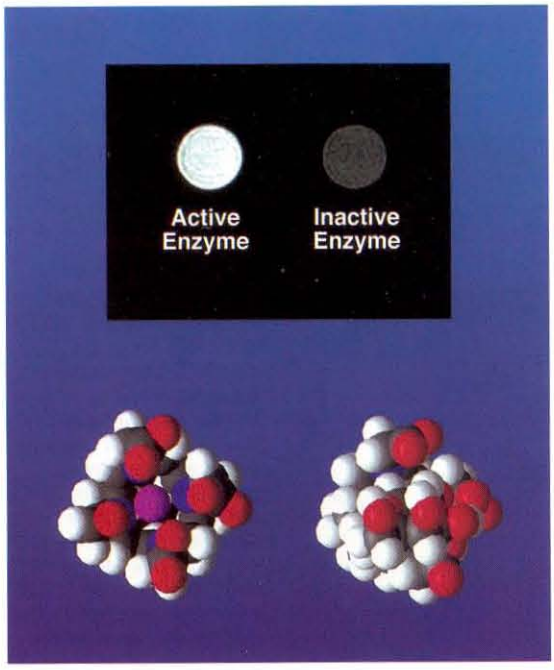

In an MRI experiment similar to the one on the preceding page, the left tube contains the active enzyme and the contrast agent; the enzyme breaks the lid off and the tube lights up. The right tube has an inactive form of the enzyme; the lid stays attached and the tube remains dark.

down. If the hinge did break, the lid would fall off, the gadolinium would be exposed to water, and the area would light up.

Above is an experiment that verified that this is what actually happened. The right tube contains our contrast agent, plus a chemically inactivated form of the enzyme that can't break the hinge; the left tube has the normal enzyme. Only the left tube lit up.

In principle, this idea could be adapted to almost any enzyme, by incorporating whatever the enzyme feeds on into the basket's hinge. We could track the progress of gene therapy in real time by injecting the patient with a contrast agent whose hinge is the target of the enzyme we're trying to impart. More broadly, we could map the pattern of activity of any enzyme throughout an organism, and follow how that activity changes over time.

The important uses of these agents in clinical and laboratory settings are several. We could diagnose many classes of brain disease, and I'll come back to this shortly. We could tell the difference between myocardial infarction (a heart attack caused by the complete blockage of a coronary artery) and ischemia (a partial blockage of an artery) in real time. This is important, because infarcted tissue is dead and gone beyond any hope of revival, but ischemic tissue can be saved by prompt action. We could identify and locate the binding sites of drugs and toxins anywhere in the body. And we could rapidly screen physiological responses to drug therapy. These agents are an enormously powerful addition to the established diagnostic technique of MRI.

We could use functional contrast agents to diagnose any organ, not just the brain and heart. As I said before, one limitation to MRI diagnoses is that ordinary, dumb contrast agents only go where the pipes go. If the patient is suffering from liver disease, for example, you can send an ordinary contrast agent to the liver to see if it's swollen, or shriveled, or otherwise abnormal looking. But you can't see how well the cells are functioning. However, a smart contrast agent that's keyed to an important liver enzyme could easily tell the quick from the dead. Healthy cells will light up, dead tissue will be black, and diseased cells will be shades of gray depending on how sick they are. You could locate the worst damage quickly and easily, without having to do biopsies or exploratory surgery.

But the biggest clinical application could be in brain diseases. We've recently modified our contrast agent to detect, not enzyme activity, but the presence of calcium—a so-called secondary messenger that transmits chemical signals between nerve cells. So by mapping calcium levels, we're mapping brain function. In this variation, the basket has a floppy lid that can grab hold of a calcium ion. When there aren't any calcium ions around, the lid dangles down on top of the gadolinium ion, keeping the water molecules away. But if a stream of calcium ions passes by, the lid swings up to grab one, exposing the gadolinium ion and lighting up the water molecules. This technique can in principle be expanded to include a variety of other secondary messengers.

Ultimately, these new MRI agents used in a traditional MRI machine may replace Positron Emission Tomography, or PET, which is how brain activity is currently mapped. PET uses a radioactive tracer, which MRI does not, and has a lower spatial and temporal resolution. (PET scans also cost several thousand dollars a pop.) Furthermore, PET needs to be done in conjunction with MRI anyway, because PET doesn't give much anatomical detail. One of the trickiest things in PET imaging is making sure that the PET and MRI images are precisely in register—otherwise you can't be sure what anatomical structure is responsible for the brain activity you see. But if all you need is the MRI scan, this problem disappears.

Our agents may have potential in the early and accurate clinical diagnosis of Alzheimer's disease, by differentiating Alzheimer's sufferers from manic-depressive individuals. The psychological manifestations of the early stages of Alzheimer's are very similar to those exhibited by manic-depressives, and it's hard to tell one disorder

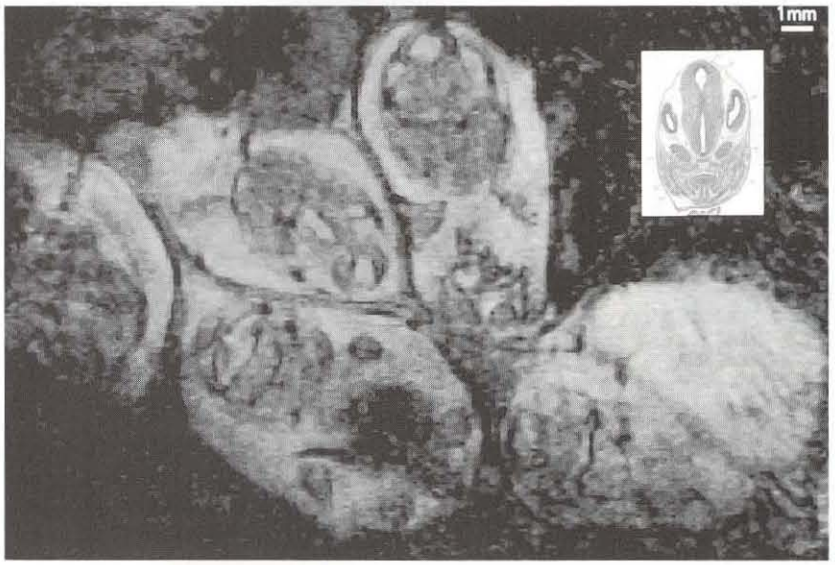

Above: The five pups in utero, as seen by MRI. Some parts of the pups are behind the image plane (the paws and belly of the pup at bottom center, for example, and the head of the one above it), while the parts in the plane show up in cross section (such as the head of the pup at top center). The inset is an anatomical map of the mouse brain, for comparison. Right: The mouse mom and the radio-frequency coil she was in, half an hour after the MRI scan was made.

from the other in a clinical interview. (The physiological changes don't become apparent until much later.) The pathology of the brains of those suffering from these diseases are quite different, however. If we were to map the calcium distributions in the brains of several known Alzheimer's patients and an equal number of manic-depressives, we might find enough differences between the two groups to form the basis for a diagnostic method. And, of course, researchers who study brain function to figure out how we perceive the world around us, how we learn, and so forth, would benefit greatly as well.

While my lab has been working with MRI contrast agents to bring out anatomical and functional details, another lab in the Biological Imaging Center has been working to push the level of detail we can see. Clinical MRI systems can see structures as small as a few millimeters, but Scott Fraser, the Rosen Professor of Biology, and Russell Jacobs, member of the Beckman Institute, have built a system that can see detail as fine as 10–15 microns—roughly the size of an individual cell. The system is purely experimental, because the magnet is too small to fit a person into. An MRI system's resolution is proportional to the strength of the magnetic field(s) used to generate the image data, and building a people-sized magnet this powerful is beyond the capabilities of current technology. But whole-body systems have recently been built that can get about 100 microns' resolution—10 times better than the clinical instruments and only 10 times worse than Scott and Russ's. (In fact, the above experiments were done using their MRI system—the capillary tubes in the pictures are a mere two millimeters in diameter.) At left is another example of their work—it's an MRI slice through the abdomen of a pregnant mouse, revealing five pups each no more than three or four millimeters long. You can see a wealth of anatomical detail in the pups' brains. And I want to emphasize that months after this MRI was taken, mom is still alive and perfectly happy, and her offspring are too.

The same advantages that have made MRI the technique of choice in medical imaging make it an ideal tool for biological experiments, so Scott, Russ, and I are using MRI to take moving pictures of embryos as they develop. We're looking at a range of organisms, from African tadpoles to small primates. This effort will open up whole new vistas in developmental biology. For example, many researchers focus on a single cell in an embryo and say, "Cell, what do you want to be when you grow up? When was that decision made, and who made it?" Up to now,

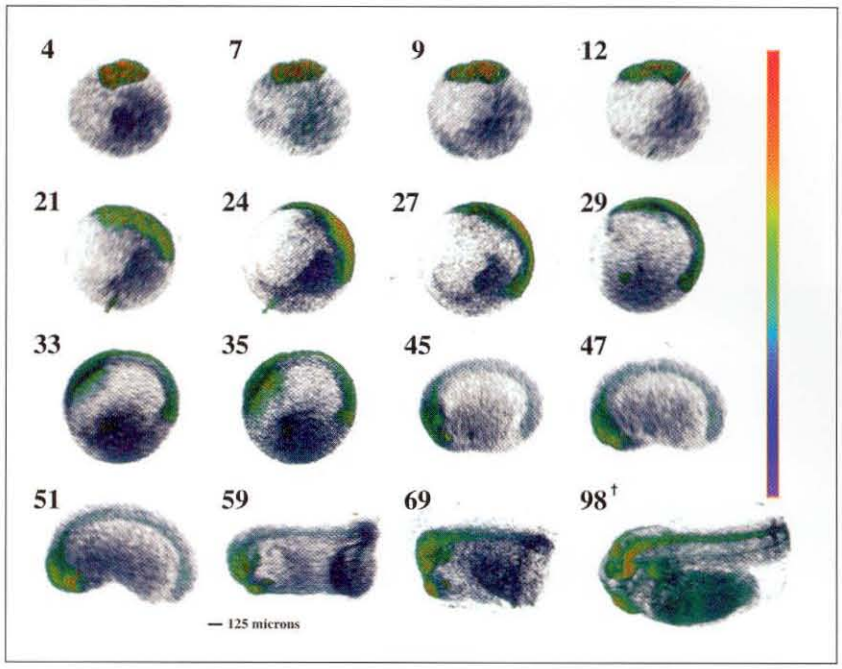

This series of images shows a *Xenopus laevis* embryo as it develops. The number next to each image denotes the elapsed time in hours since the egg was fertilized. When the egg had divided to the point where the embryo contained 32 cells (about two hours after fertilization), one cell was injected with an MRI contrast agent. By F+4 hours, that cell had become eight cells. In later frames, these descendants marched off two or three abreast to become the spinal cord until, at F+33 hours, one end of the spinal cord begins thickening to become the brain. A small platoon of progeny, visible at F+29 hours as a green spot near the center of the embryo's lower left quadrant, became the heart. The color bar along the right side of the panel is keyed to the contrast agent's concentration, with red being the highest and purple the lowest.

the way to find out was to inject a contrast agent into that cell and follow its descendants as the animal developed. But now, we can not only watch the migrations of a certain family of nerve cells, say, as the embryo's brain wires itself up, but we can also see when those cells decided to become nerve cells in the first place. We could inject the original cell with our really smart contrast agent, which we had programmed to light up when some enzyme specific to the nerve cell becomes active. Or, we could add our moderately smart agent to the solution in which the embryo is swimming, and wait for the cells to sprout nerve-cell-type receptors.

But what really makes our contrast agents so powerful is that MRI data are three-dimensional. We can build a 3-D model of our organism and rotate, tilt, and slice through it any way we please. We can extract innumerable images from a single 3-D scan without ever pulling a scalpel out of the drawer. And when we make a succession of 3-D scans over time, to follow an organism's growth or a disease's spread, it gives us amazing flexibility in the questions we can ask of the data. Here's a vivid example—the pictures above are frames from a video by Russ Jacobs' group. They have labeled one of the 32 cells in a frog embryo with an MRI contrast agent. That cell is going to split many times as the embryo grows, and we can track the great-great-great-great-granddaughter cells in 3-D. Some of the offspring will form what will become the spinal cord, others will become the heart, and most of them will eventually wind up in the brain. (Had we started with another cell, we might have gotten the intestinal tract.) So we can cut an embryo in slices any way we please while it's still growing, and we can make out changes in very fine three-dimensional detail in a single animal over a long period of time.

We'll never make a handheld diagnostic MRI system, because the magnet still has to be big enough to fit a person into, but we could perhaps make a system that would fit into an ambulance or minivan. (Today's "mobile" systems only qualify in the broadest sense of the world—they live in those semitrailers you sometimes see parked behind medical centers. You have to hitch them up to a diesel cab to take them anywhere.) In contrast, our system for DNA analysis truly is portable—a simple one could be built into a device the size of a garage-door opener.

DNA is our genetic material—and not just human genetic material, but that of every living thing. Knowing what kind of DNA one has in a sample and who that DNA belongs to, whether searching for a disease or a hardier strain of vegetable, requires a lot of technology and is rather cumbersome. You have to collect a sample, extract the DNA, and then purify it. There are all sorts of ways that contamination can creep in, and innumerable opportunities for mistakes to be made. The procedures require a lot of complicated equipment and all sorts of chemicals, so you can't take the lab to the samples; you have to bring the samples to the lab. The current methods are very powerful, but they're also labor intensive, time consuming, and extremely expensive.

If, however, there existed a handheld tricorder

that could run a whole bunch of DNA tests simultaneously, a lot of applications would open up that are impractical today. We could, for example, rapidly check the entire U.S. blood supply for contamination by all known strains of the AIDS virus. We could monitor our air and water supplies for infectious agents. This could be anything from tracing the progress of a particularly virulent strain of the flu, to dealing with a situation like in the movie *Outbreak*, where an airborne Ebola-like virus got loose in the United States. More prosaically, the food industry is very concerned about bacterial contamination, especially by *E. coli*, which can cause food poisoning. (You may remember the fast-food scare of a couple of years ago.) There are forensic applications—finding out who the murderer is from bloodstains left behind. In addition, there's agricultural monitoring. Disease-resistant genes can be inserted into plants to improve crop yields, but, again, there's currently no way to tell if the gene has "taken," short of waiting for the seedlings to sprout and then screening them for whatever resistance the gene was supposed to impart. Being able to test the DNA of the seed itself, before it ever leaves the lab, would assure that only the disease-resistant seeds get planted.

We can extract innumerable images from a single 3-D scan without ever pulling a scalpel out of the drawer.

Our method uses electron-transfer reactions to analyze DNA. Electron-transfer mechanisms have been a popular subject here at Caltech for many years—Rudy Marcus, the Noyes Professor of Chemistry, won the Nobel Prize in 1992 for electron-transfer work. And Harry Gray, the Beckman Professor of Chemistry, has been studying electron transfer in proteins for more than a decade. His group discovered that electrons could travel across large distances in complex biological molecules such as electron-transfer proteins. Such processes are not unique to proteins; for example, Professor of Chemistry Jackie Barton's group has data indicating that electrons can go through DNA very rapidly. In a nutshell, you do these experiments by putting an electron donor on one end of a piece of DNA, and an electron acceptor on the other end. Then you launch an electron from the donor and measure how long it takes to arrive at the acceptor.

It would seem reasonable to assume that how fast the electron goes might depend on the exact nature of the DNA it's traveling along, which led us to wonder if we could use electron-transfer rates to identify a DNA molecule—specifically to tell us whether it matches a reference piece of DNA whose identity is already known. DNA is a long, linear molecule, and normally two strands of it interlock like the two halves of a zipped zipper. It's the way that the four chemical "letters"

In a piece of double-stranded DNA, the chemical "letters" (seen here edge-on), which carry the genetic information, recognize one another and pair up like rungs on a ladder. The surrounding tracery of yellow and red is the phosphate backbone on which the letters are strung,and whose natural twist imparts to the molecule its classic shape. The magenta and yellow spheres are the electron donor and acceptor added for electron-transfer experiments.

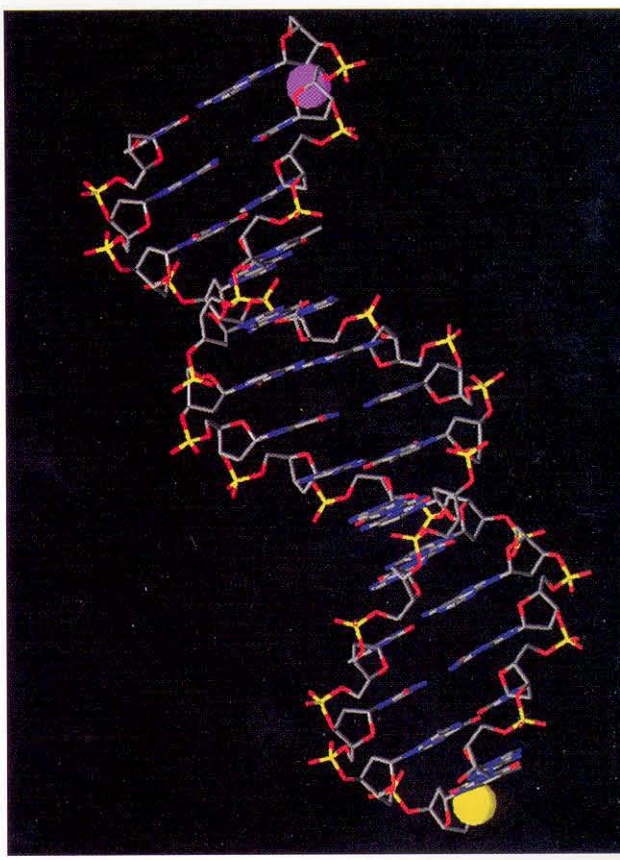

I've often wondered if perhaps the monks stumbled onto the structure of DNA 700 years ago, and the cathedral was the only journal they could find to publish it in.

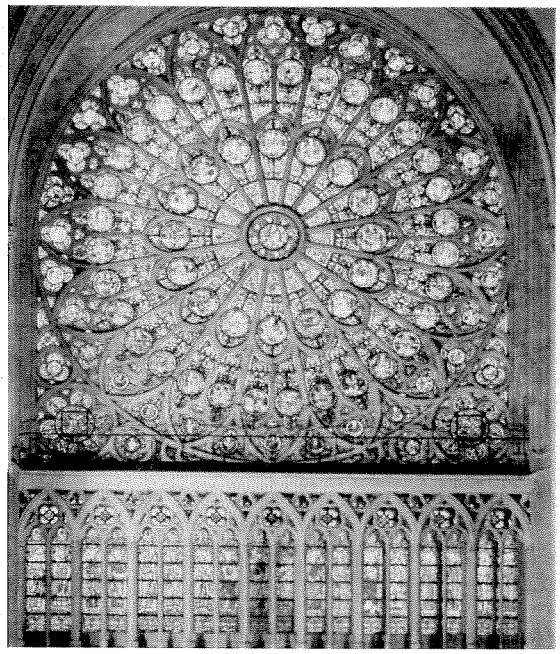

(A, C, G, and T) in the DNA code recognize each other that makes the zipper zip. The code follows two simple rules: an A on one strand only binds to a T on another strand, and a C on one strand only binds with a G on the other. So if you know the sequence of one strand of DNA, you also know the sequence of its complementary strand—the strand that will zip up with it. For every A on one strand, there's a T in that spot on the other, and vice versa; as is the case for C and G. Thus, if we know the sequence of a strand of DNA that's unique to a defective gene, for example, we can make a probe—a single strand of DNA with the complementary sequence—and if that defective gene is present in our sample, the probe will find it and zip it up into a double strand.

This leads to the key question: is the electron-transfer rate along a double strand sufficiently different from the rate along a single strand that we could reliably tell double from single? To find out, we made short pieces of DNA that had an electron donor attached to the terminal letter on one end and an electron acceptor attached to the terminal letter on the other end, taking pains that the donor and acceptor didn't perturb the DNA structure in any way. We synthesized some pieces with the donor and acceptor on the same strand, and some with them on opposite strands. That way we could not only do the single-strand versus double-strand experiment, but we could also find out whether it made a difference whether the electron had to transfer from one strand to the other. Finally, we wanted to be able to track the electron unambiguously, so we gave the donor and acceptor different spectroscopic fingerprints. In fact, we actually had four distinct spectroscopic signals—one from the donor with the traveling electron ready for launch, one from the donor after the electron left, one from the acceptor before the electron arrived, and one from the acceptor after the electron had landed.

It turns out that the electron goes from one end to the other end of a double-stranded piece of DNA very rapidly, much faster than would be predicted from the rate of electron transfer through a single strand. In a single strand, the only way the electron can get from donor to acceptor is by tunneling down the phosphate backbone on which the DNA's letters are strung. But when two strands zip up, the electron can scoot along the letters themselves, which now line up like the steps in a spiral staircase. You can see this alignment better on the inside front cover of this magazine, where there's a nontraditional view of DNA that rather resembles the rose window at Notre Dame shown above. I've often wondered if perhaps the monks stumbled onto the structure of DNA 700 years ago, and the cathedral was the only journal they could find to publish it in.

This rate difference is being exploited by building a microelectronic chip, or biosensor, that has the single-strand probe DNA attached to it. The chip would send electrons through the probe and measure its resistance, which is very easy to do. If the probe recognizes its target sequence—say, a viral gene—the resistance would suddenly drop, because electrons would

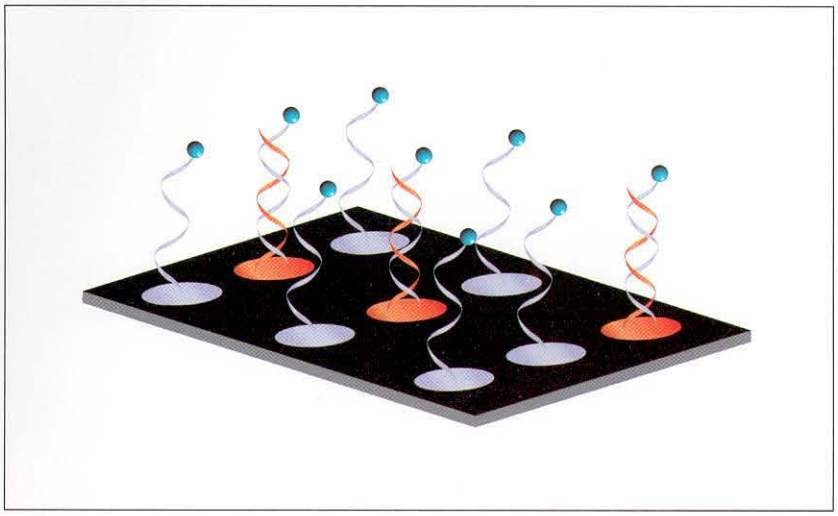

The conceptual basis for a tricorder on a chip. Anchored to the chip is an array of single strands of DNA whose sequences are complementary to, say, an assortment of disease-susceptibility genes. When a piece of the suspect gene drifts by, it binds to its opposite number on the chip, and the decreased resistance through the double-stranded DNA indicates a match.

travel faster through the double-stranded DNA than the single strand. Ultimately, we could put an array of thousands of different probes for all sorts of things on a single chip.

In principle, anything that has DNA in it could be reliably detected on such a chip. We already know the sequences of many pieces of DNA—including the *p53* gene, which has been linked to colon cancer, and all the mutants of the HIV virus, which causes AIDS—and laboratories are working out more sequences all the time. The laws of statistics say that using a 17-letter probe for a sequence that is found only in the piece of DNA we're looking for will virtually guarantee that only that single-stranded piece will match. To further reduce the possibility of false-positive results, a number of different probes for each gene can be placed on a single chip. A collective pattern of decreased resistance from those probes would indicate that a match had been detected.

Today, several companies are integrating DNA into chips with individually addressable circuits. Jon Faiz Kayyem (PhD '92), formerly a postdoc in my group, has cofounded a local company (in cooperation with Caltech's newly formed Office of Technology Transfer) to explore ways of combining this chip-making technology with our rapid-electron-transfer diagnostic methods.

In the next generation, a five-year-old might go to the doctor's office, and the doctor might draw a blood sample, stick it in the tricorder, and say, "Aha. We have decreased resistance through this probe, and this one, and that probe over there, and those two probes there." This pattern might reveal that the child has a mutant *p53* gene and might be at risk, some day, of developing colon cancer. Then the doctor could say to the child, "You need to come back in 45 years, because we'll want to start screening you for colon cancer on a regular basis then." Now, that's what I call early detection—45 years in advance of the potential onset of something is a pretty big lead time.

This brings us full circle to the notion of a diagnostic tool that we began with—a device that's fast, accurate, and makes the right call the first time every time. If Hollywood's vision of the medical technology of the future has any basis in reality, then today we may be laying the groundwork needed to put a tricorder in Dr. McCoy's hands tomorrow. □

Thomas J. Meade earned his undergraduate degree in chemistry from Arizona State in 1980, and proceeded to Ohio State, where he got his master's degree in biochemistry in 1982 and his PhD in inorganic chemistry in 1985. He next worked on magnetic resonance imaging as an NIH postdoctoral fellow at Harvard Medical School, before his first sojourn at Caltech as a research fellow (1987–89), when he studied electron transfer in metalloenzymes with Harry Gray. He joined the Division of Biology and the Beckman Institute in 1991. This article is adapted from a recent Watson lecture.