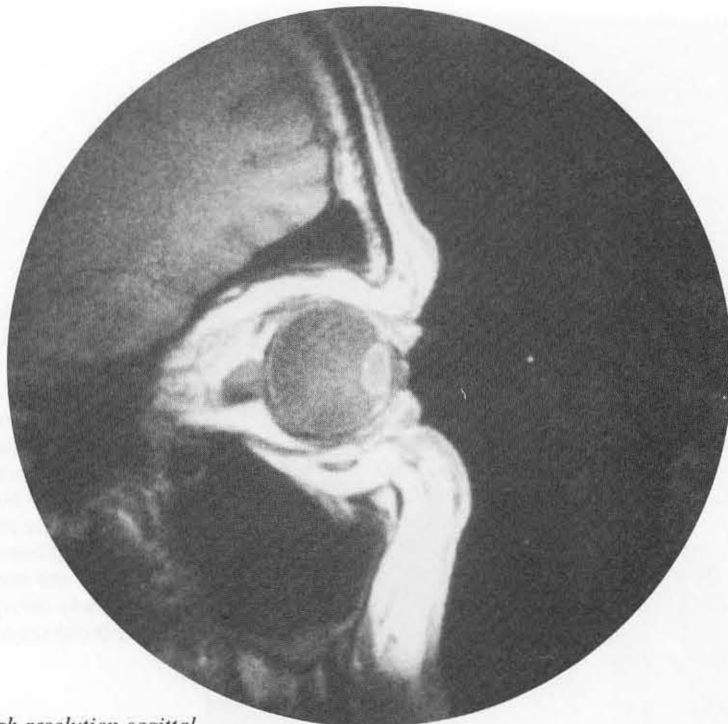

Very high resolution sagittal section of a human eye, made by proton NMR and a surface coil as receiver to improve the signal-to-noise ratio (GE).

Biomedical Applications of NMR

by John D. Roberts

The NMR images used in this article were kindly supplied by the Medical Systems Group of the General Electric Company or the NMR Imaging Laboratory of Huntington Medical Research Institutes and identified, respectively, by GE or HMRI in the figure captions.

NUCLEAR MAGNETIC RESONANCE (NMR), a laboratory technique that has been around for 40 years, has in the past decade been applied to medical imaging. Its potential for studying blood flow and metabolic processes in the heart and brain, as well as abnormal tissue, is enormous, and NMR offers advantages in competition with other noninvasive imaging techniques, particularly in not using ionizing radiation.

Nuclear magnetic resonance (NMR) spectroscopy was first demonstrated independently by physicists Edward Purcell (Harvard), Felix Bloch (Stanford) and their associates, work that was recognized by the award of a Nobel Prize in 1952. Caltech got into the field almost immediately through the perspicacity of Chemistry Professor Don M. Yost, who carried out some truly pioneering studies. Two outstanding Caltech graduate students of that period, John Waugh, now at MIT, and James Shoolery of Varian Associates, have contributed enormously to the continued development of NMR. My own work, starting 30 years ago, involved early applications of NMR to organic and biochemistry and helped to demonstrate the usefulness of NMR for studying molecular structure and molecular dynamics. There is no method as generally powerful as NMR for the study of liquids or substances in solution.

The history of NMR has been punctuated with quantum leaps associated with new technical developments: higher field magnets, more stable frequency oscillators, more sensitive detectors, and particularly advances in signal processing. In 1973, Paul C. Lauterbur (now at the University of Illinois) showed the potential of NMR for making images, and a technical explosion in that application followed. The NMR Imaging Laboratory at the Huntington Medical Research Institutes (HMRI) in Pasadena, starting in 1982 under the direction of Dr. William G. Bradley (Caltech BS 1970), has been a leader in both clinical and basic research in NMR imaging. Currently the whole-body imager at HMRI has had a greater throughput of patients than any other imager in the world.

What is NMR and how can it be applied to imaging? On the one hand, we have to be able to achieve spatial resolution, that is, measure distances between objects. On the other, we have to be able to establish a contrast scale that will differentiate materials of, say, the eye's lens from the vitreous humor, from bone, or from the brain.

Let's start with the N — for "nuclear." The nuclei of all atoms are positively charged and very much smaller than the negatively charged electron clouds surrounding them. Many nuclei have a property that corresponds to spin. Although the nuclei of ordinary oxygen and carbon don't have spin, the nuclei of hydrogen atoms (protons), the H in H_2O, do have this property. As the nuclei spin, their charges circulate and generate magnetic fields in the same way that charges moving in a circular loop of wire generate a magnetic field. And that's the M — the "magnetic."

Such magnetic nuclei, which have north and south magnetic poles, ordinarily have no preferred orientation in space. But if we put them in a magnetic field, they tend to line up with the field in a manner analogous to the way that a compass needle lines up in the earth's magnetic field. This, then, is the magnetic property that we use in NMR — the tendency of the nuclei to align in an applied electromagnetic field and produce a more favorable energy state.

The next thing we do is to change the orientation of the nuclei in the field — turn them over and make them point the other way. Because this is a less favorable state, we have to put energy into the system. This energy can be obtained from a radio transmitter with the proper frequency. The phenomenon of turning over the nuclei by a transmitter putting in just the right amount of energy is called "resonance" and provides the R — the last of our string of NMR letters. The transmitter frequency, ω, to obtain resonance is equal to the magnetic-field strength, H, times a nuclear constant, γ, ($\omega = \gamma H$). This relationship holds very precisely, which is important to us. In the magnetic fields that we use for imaging at HMRI (about 3,500 gauss), the resonance frequency for protons in water is 15 MHz.

There are many ways to detect NMR signals. For most spectroscopy and imaging, we turn the magnetized nuclei from the favorable state to the unfavorable state by a short burst of energy at the resonance frequency and analyze the change in magnetization as the nuclei return to the favorable state — a process that can take from milliseconds to minutes depending on the kind of magnetic nuclei and the nature of the sample.

The magnets used for whole-body medical imaging must be large enough to accommodate reasonably hefty adults, and it is difficult to fabricate a high-field electromagnet with electrical coils surrounding an iron yoke with the necessary large gap between the pole faces. Superconducting solenoids are particularly useful where we need a large magnet with a high magnetic field. Superconducting magnets most often use coils of niobium-tin wire cooled with liquid helium to 4 K (-450° F). At this temperature, the electrical resistance of the coils drops to zero, and currents induced in the coils will circulate and generate magnetic fields indefinitely, as long as the coils are kept cold.

But how can NMR be used for imaging? You can't do it by the ordinary ways in which you see things as the result of diffraction, reflection, and so on, because the radio

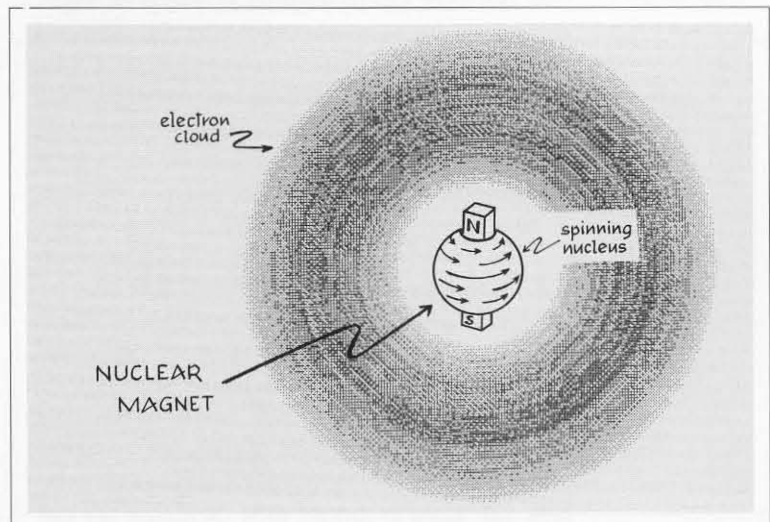

Schematic representation of an atomic nucleus surrounded by its electron cloud. Nuclear spin results in generation of a magnetic dipole that can become aligned in an external magnetic field. The size of the nucleus is here greatly exaggerated with respect to the size of the surrounding electron cloud.

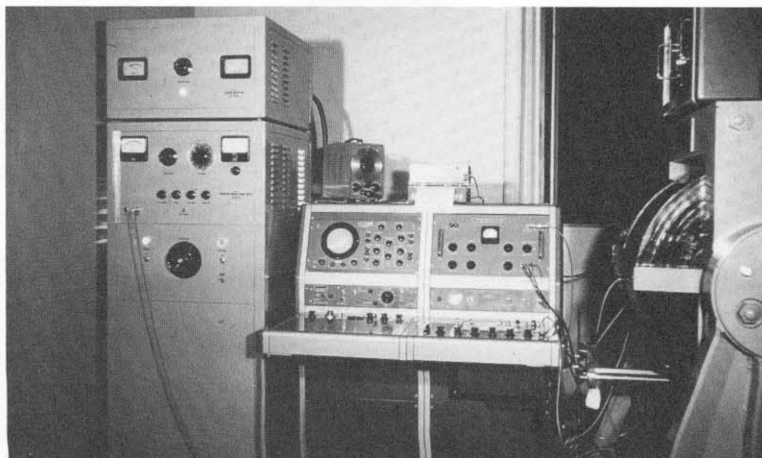

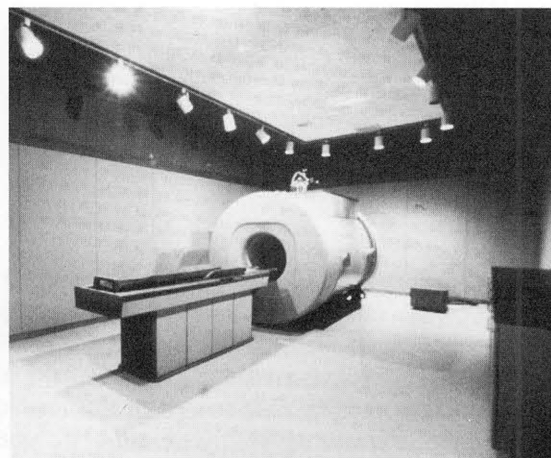

Above: Caltech's first Varian 40-MHz NMR spectrometer installed in 1955.

Above right: Diasonics whole-body proton imaging apparatus at Huntington Medical Research Institutes, operating at 15 MHz with a superconducting magnet (HMRI).

waves we're using in NMR have enormously long wavelengths relative to the 0.5 mm resolution that NMR imagers now routinely achieve. Expecting that radiation with a 20-meter wavelength (corresponding to 15 MHz) could image the brain would be like expecting a toy boat to measurably perturb ocean waves. So the technique must be different from the usual imaging techniques based on light waves, x-rays, or electrons. The key to spatial resolution in NMR imaging is the fact that the resonance frequency is equal to the nuclear constant times the magnetic field. Thus we can control the resonance frequency precisely by controlling the magnetic field.

The use of this relation can be illustrated by resonance signals from two identical water samples. First we put them in a uniform magnetic field. The resonance frequency of both samples will be exactly the same, and so if we do our nuclear resonance experiment, we will get a single composite signal from them. Now, suppose we change the system by having a gradient in the magnetic field such that the field at one sample is different from that at the other. The two samples will then come into resonance at different radio frequencies. If the field gradient is linear and we know how fast it changes with distance, the distance between the samples can be calculated simply by determining the difference in resonance frequencies. If there is a 50-Hz difference in frequency for a particular field gradient with two samples 5 cm apart, then we know that a 100-Hz difference corresponds to samples 10 cm apart.

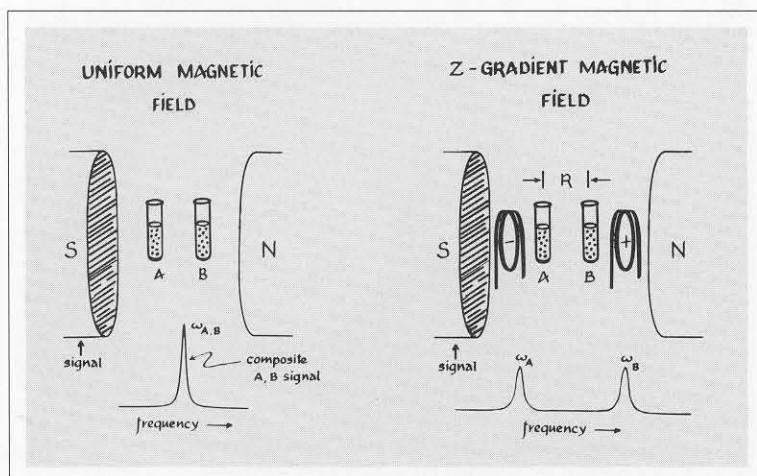

Below: NMR signals expected from two otherwise identical samples in a uniform magnetic field (left) and a field with a gradient (right). The relation $\omega = \gamma H$ tells us that the signals from the sample will superimpose in the uniform field and will have different frequencies in the gradient field, provided that the gradient is not zero along the line of separation, R, of the samples.

We can actually measure the *x*, *y*, and *z* coordinates of particular volume elements in, say, the brain by using *x*, *y*, and *z* gradients to localize the frequency that comes from that volume element. So the magnetic-field gradients are the keys to getting the spatial information. The way we normally do this in practice is to set up the gradient along one axis and activate the nuclei at the proper frequency for a particular segment of that gradient and then analyze the signals while using additional *x* and *y* gradients during the return of the nuclei to the favorable orientation. This can give us an image of a slice through the subject. We can routinely obtain images for 20 such slices in 4 to 17 minutes. Longer times give better resolution.

Unlike other imaging techniques, NMR allows us to change the orientation of the slice without moving the patient. By simply applying the gradient in the *x* direction, we can excite the nuclei in slices parallel to the *y-z* plane to get sagittal images (representing vertical slices parallel to a plane through the nose and the back of the head). And by changing the gradient to analyze slices parallel to the *x-z* plane, we can get a coronal image (representing vertical slices through the head parallel to the ear-to-ear plane).

But how do we differentiate between the various kinds of tissue we have after getting

distance information? This is vital information for images. We must be able to distinguish between blood, fat, muscle, and so on. This is a more difficult problem, whose solution can be illustrated as follows.

A small volume element, such as a single drop of water in a magnetic field (as in an imager), will contain a given but very large number of magnetized protons. If we hit these protons with a short burst of energy at the right radio frequency, they respond analogously to identical bells hanging in a row from an essentially weightless beam and all set to ringing at once by a hammer blow on the end of the beam. The bells swing together, ring together, and ring down together. There are two aspects of this that are important to note: One is the rate of loss of the energy imparted to the bells when you hit the end of the beam. The other is the rate of decay of the ringing sound. In our nuclear case, the rate of loss of excess energy and the rate of loss of the radio signal are different forms of what is called NMR relaxation.

Now, if our bells are all identical and they are hit so that they swing in unison, we might expect their ringing to decay away at the same rate as the extra energy is lost. This is an ideal case which may not actually happen with real bells or with real nuclei either. The reason is that otherwise identical bells, or identical nuclei, may not be in exactly the same surroundings. In any case, we can differentiate the characteristic time associated with the rate of loss of excess energy (which we call T_1) and the characteristic time associated with the loss of the ringing sound (which we call T_2). The rate of energy loss may be different from the rate of ringing loss if, for any reason, the bells do not ring with exactly the same frequency. Then their rings can interfere with each other, and the ringing signal will disappear before all the energy is lost. We cannot expect that the ringing signal will ever persist after all the extra excitation is lost by relaxation, so T_2 will always be smaller or equal to T_1.

What difference does that make to the nuclei we use in NMR? When T_1 is long *and* equal to T_2 (what we can call phase-coherent signal decay), and we analyze this time-dependent signal decay for its frequency content (Fourier analysis), we get a sharp line. When T_1 is much greater than T_2 (poor phase-coherent decay), that is, the ringing sound goes away much faster than the excitation energy is lost, we get a broader line when we analyze for frequency content. Differences in line width provide us with one way of differentiating between biological materials.

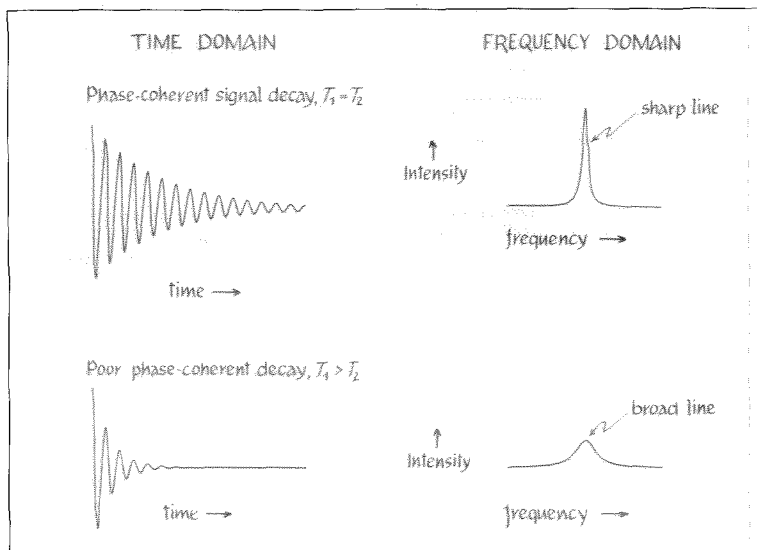

Above: Time domain signal-decay curves in relation to T_1 and T_2. The two decays have the same T_1 but different T_2 values. Analysis of the frequency content (Fourier transform) of the two decays shows a much sharper line in the frequency domain for the longer T_2 value.

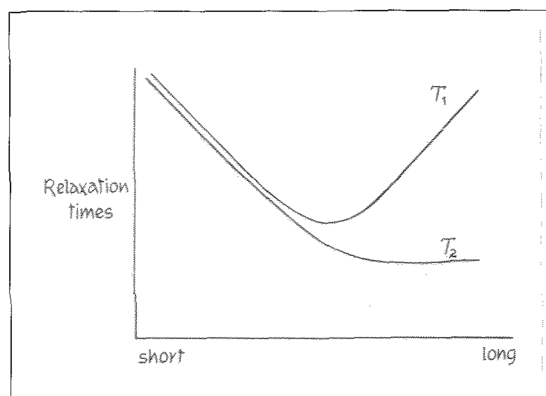

Left: Relationship between NMR relaxation times, T_1 and T_2, of molecules in solution with the average times required for molecular reorientation as the result of collisions. Reorientation of large molecules is generally slow and will be further decreased in media that are viscous.

Researchers have found that T_1 and T_2 tend to be more nearly equal when the molecules that make up the material being investigated move quickly with respect to one another — that is, when the material is non-viscous, like water. When the molecules begin to move more slowly, the material becomes viscous (like honey or maple syrup), and T_1 and T_2 values both become smaller and begin to diverge from one another. Surprisingly, with some very viscous materials present in the body, we find that T_1 goes through a minimum and starts to become larger, while T_2 continues to decrease with increases in viscosity.

Several factors determine the strength of NMR resonance signals of protons. Obviously, the more protons you have, the stronger the signal will be. Biological materials such as blood, cerebrospinal fluid, fat, muscle, brain white and gray matter, and bone have different proton contents. Bone has the smallest, with blood, muscle, fat, and

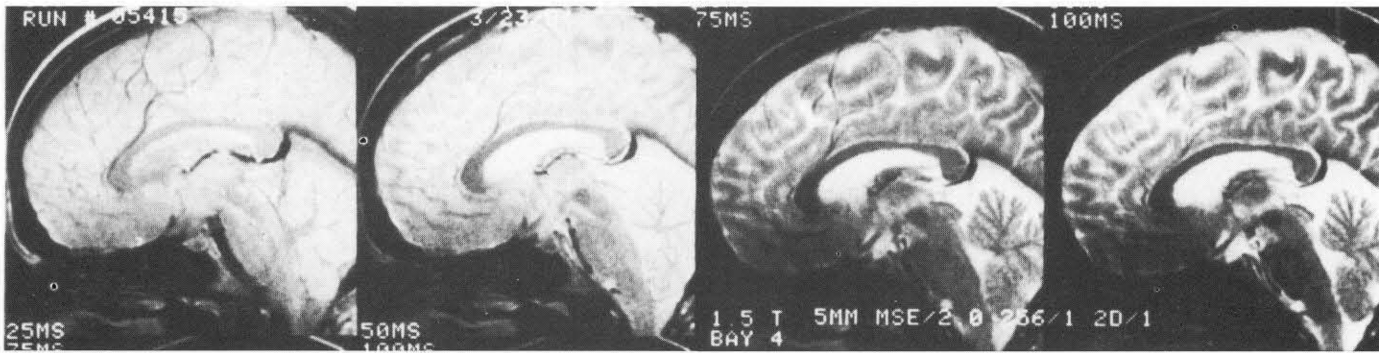

Sagittal brain images as a function of time allowed for relaxation; here, differences arising from T_2 are emphasized over the range 25-100 milliseconds (GE).

cerebrospinal fluid having increased amounts in that order. The T_1 and T_2 values are also very different for these materials, so we can use T_1, T_2, proton content, or combinations of these to establish a contrast scale. For clinical purposes, a combination of proton content and T_2 is most often used.

To illustrate how this works, let's consider bone, which has little water in it, so it's not going to give a very strong proton signal. Further, it has short T_2 values, so the resonance signal will decay away rather quickly. The signal from bone will therefore be weak and fast decaying. So, if we sample the signal from bone at different periods after the initial excitation pulse, it will always be weak. We can color such weak signals as black or very dark gray. Muscle has many more protons per unit volume and longer T_2 values, so it will give a stronger signal that decays more slowly than that of bone. We could assign medium gray to muscle and light gray to blood, which will also have a strong signal with an even slower decay.

We can see in the illustrations above the remarkable difference in contrast in brain images as a function of time. The images correspond to different decay times for sagittal slices ranging from 25 to 100 milliseconds after the original excitation pulse. In the areas that remain light, the protons are losing the "ringing" very slowly. Others (the darker areas) lose the major part of it more rapidly. The overall signals are decreasing in intensity all the while, but each of these pictures has been brought up to the same average intensity by turning up the brightness as one might do on a TV set. At the outset, only a small fraction of the signals has decayed, and there's not much differentiation between the tissues on the basis of the differing T_2 values. However, as time unfolds, large differences develop because the signal from some of the materials decays much more rapidly than others. Rapid decay produces the dark areas;

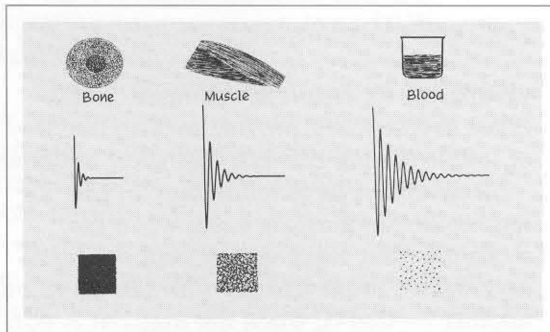

***Right:** NMR signal intensities and relaxations to provide a contrast scale for bone, muscle, and blood.*

***Below:** CT (x-ray) section (left) of patient suffering from multiple sclerosis compared with transaxial NMR images, T_2 (center) and T_1 (right) (HMRI).*

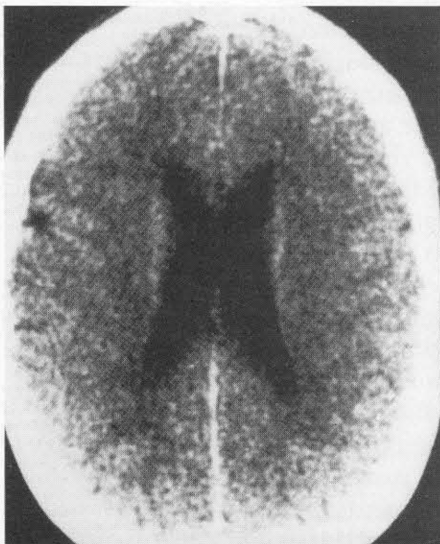

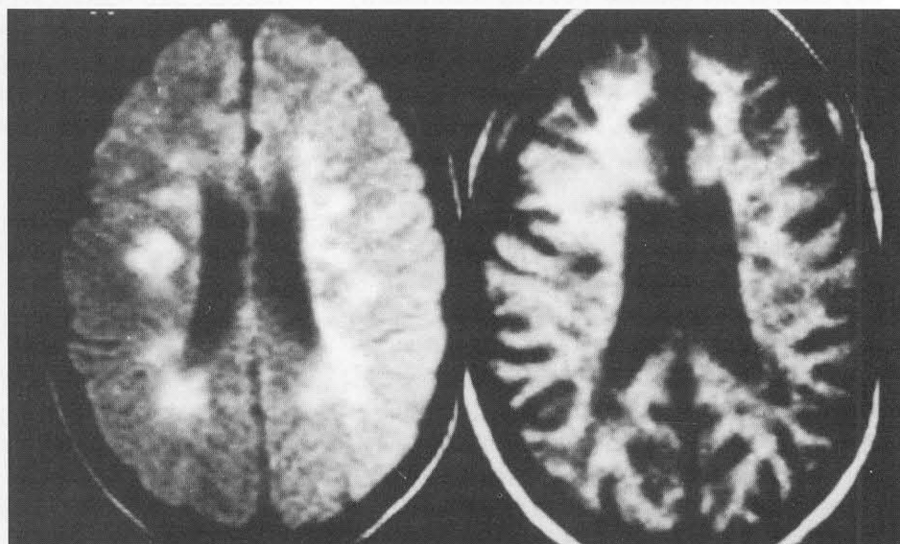

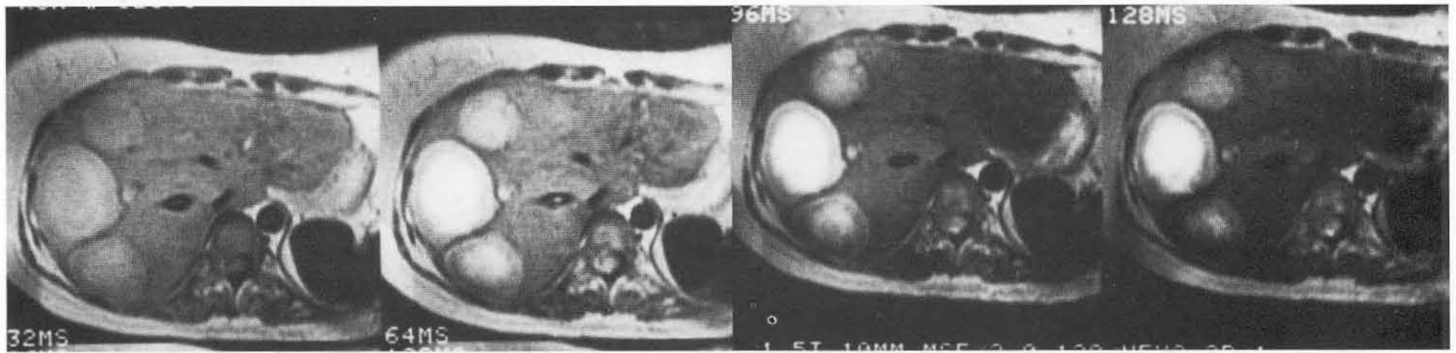

Section through the body, showing how a cancerous growth in the liver becomes prominent in the image as more time is allowed for relaxation. The range is 32-128 milliseconds (GE).

slow decay produces areas that remain bright.

A comparison between CT (x-ray) and NMR scans of a patient with multiple sclerosis shows that the x-ray image provides relatively little differentiation between different areas of the brain. But in the T_2 and T_1 images of the same patient you can see some marked white areas. These are demyelinated areas, where the insulation (myelin) covering the nerves has decomposed with a consequent increase in water content and T_2 values. The way in which the abnormal areas stand out makes the NMR scan favored in this case.

Of particular interest is the fact that the water (proton) content, T_1, and T_2 values of cancerous tissues are, in general, substantially greater than those of normal tissue. By delaying the observation time and allowing the T_2 effect to become more prominent, cancerous tissue in, say, the liver, stands out very brightly, as shown in the illustration above.

In just the last few years, remarkable NMR pictures have been made of the heart. The heart is a very difficult organ to image by either NMR or x-rays, because it pulsates and therefore needs relatively short exposure times to give unblurred pictures. Unfortunately, NMR is not exactly a 1000-speed film; it's more like a relatively slow film, and so we have to repeat our snapshots many times and then add them together to get a useful picture. Thus we need to start the excitation pulse and the analysis period at particular times in the heart cycle. Proceeding in this way, we can make images of the heart at different parts of that cycle.

NMR techniques do very nicely for noninvasive measurements of the diameter of the aorta and for detecting deposits of cholesterol and lipids associated with hardening of the arteries. Dr. Bradley of HMRI has made some important studies on the use of NMR to investigate blood flow. Earlier we indicated that blood should have a contrast of light gray in NMR imaging. But in the image of the aorta on the following page it shows as black. Why would this be? Remember that, when we start at time zero, we use a gradient and are able to pulse the protons in a particular plane or slice of the body a few millimeters thick. So in a midriff cross-section, we pulse only the blood that is in that aorta slice at the particular time. When we analyze the signal from the protons in the same aorta slice some 28 milliseconds later, the blood we magnetized earlier has already moved farther down, wholly out of the slice. It has been replaced by magnetically unexcited blood, which of course gives no NMR signal and hence comes out black in the image. This way of analyzing blood flow can detect turbulence in flow and has very general potential for analysis of the mixing of liquids.

Still another burgeoning application of NMR is the study of metabolism. Dr. Dominique Freeman, a Boswell Postdoctoral Fel-

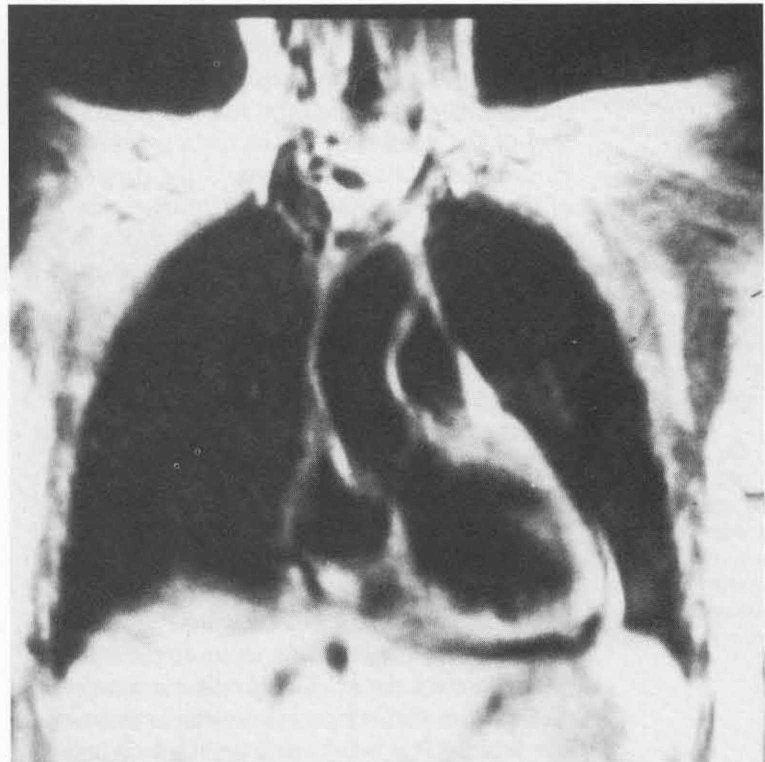

Images made during different parts of the heart cycle by having the NMR excitation pulse synchronized to particular parts of the EKG waveforms (GE).

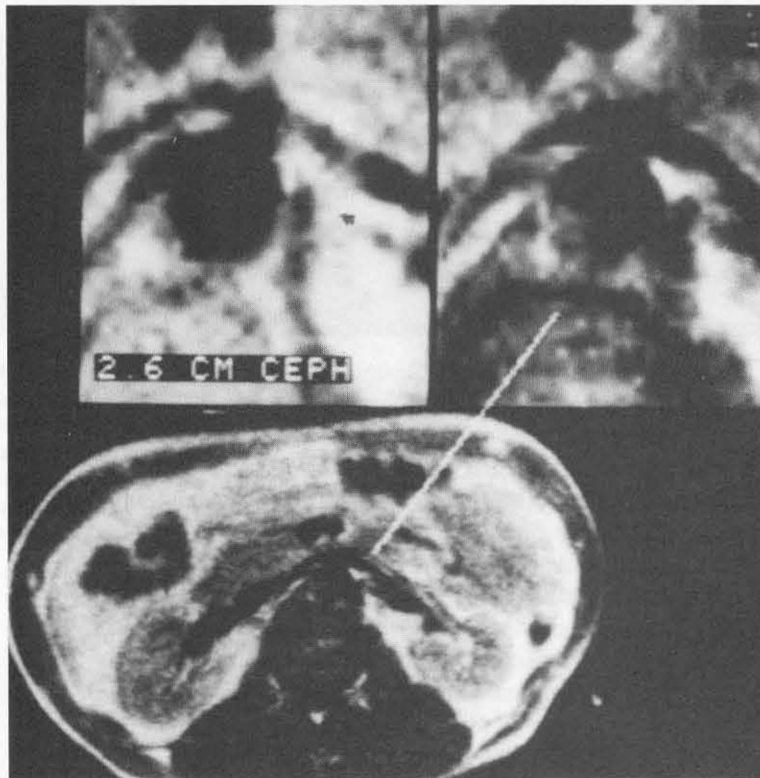

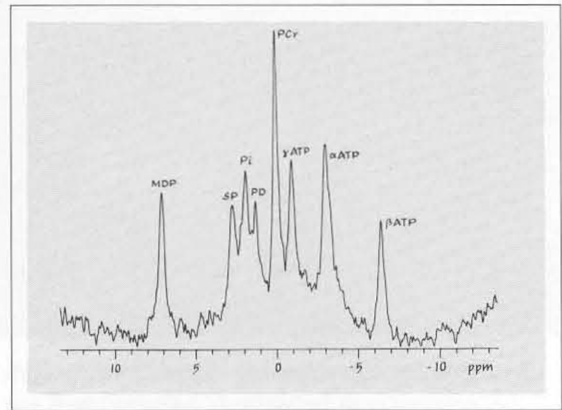

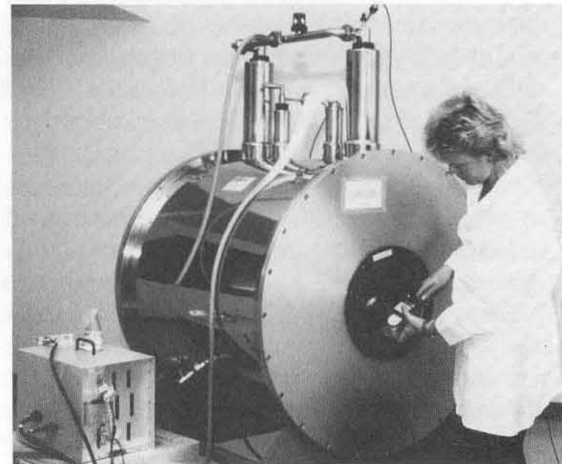

Above: Midriff section showing aorta (at ends of white line, enlargement upper right) with blood appearing black as the result of flow. The light gray area visible in the aorta between 7 and 10 o'clock at the upper right is a deposit of fatty matter characteristic of atherosclerosis. The inset at upper left is taken further down the aorta, and here the deposits are essentially absent (HMRI).

Top right: Phosphorus NMR spectrum of a living rabbit heart. The key resonances for energy metabolism are the three α, β, and γ peaks of adenosine triphosphate (αATP, βATP, γATP) and the single peaks of phosphocreatine (PCr) and inorganic phosphate (P_i).

Bottom right: Dr. Dominique Freeman arranging for a phosphorus NMR spectrum in a high-field, large-bore superconducting magnet system made by General Electric and set up at the Huntington Medical Research Institutes.

low at Caltech, and Harold Mayr from the University of Vienna are working at HMRI in collaboration with Dr. Richard Bing, studying the metabolism of phosphorus, rather than protons, in a living rabbit heart. The phosphorus NMR spectrum (shown above, right) from such a heart has several interesting features. The spectrum is complex, but we can see three different resonance peaks coming from the three different phosphate atoms of the molecule ATP (adenosine triphosphate), which is the fuel for contraction of the muscle in the heart; a peak from P_i, inorganic phosphate, one of the products of the reaction of the muscle contractions produced by the ATP; and a peak from PCr (phosphocreatine), which brings the phosphorus into the heart muscle in the proper form to be easily converted to ATP.

We can study a heart with diseased muscle under stress and see how the concentrations of these metabolites change in comparison with normal heart muscle. Then, when the stress is removed, we can then see how fast the recovery to unstressed conditions occurs. Many people have serious problems in regenerating the proper balance of phosphorus metabolites after periods of stress.

In principle, we could make images of brain slices to determine the distribution of ATP within the brain and also determine how fast this ATP is being used up. This is still beyond the reach of current technology — but not far out of reach. It's very tantalizing. The problem is that phosphorus is 10,000 times more difficult to detect in living tissue than protons. For one thing, the concentration of phosphorus is much lower than the concentration of protons in living tissues. Also, phosphorus has a smaller resonance frequency, and its nuclei are not as favorable to observe as protons.

The uses of NMR in biology and medicine offer a most fertile field of research. Caltech has been very much in the forefront of chemical NMR in the past. Now there's a wonderful opportunity for collaborative work between Caltech and HMRI on these exciting new processes, and we should take full advantage of it. Some years ago, during Harold Brown's presidency, Caltech considered starting a medical school but decided not to proceed for a number of good reasons. Collaborative work with HMRI would give Caltech the opportunity to become involved in various aspects of medical research without the problems of running a medical school. NMR is just one of the several possible attractive modes of present and future collaboration to the benefit of each institution. □