Diamond defects enable nanoscale nuclear magnetic resonance

Detecting the tiny magnetic field from a few thousand atomic nuclei is a first step toward imaging complex molecular structures directly.

Understanding a protein's function is difficult if you don't know its structure. The standard tools of the trade, x-ray crystallography and bulk nuclear magnetic resonance (NMR), work well for molecules that can be produced in macroscopic quantities and arrayed in a crystal or dissolved in solution. But for many proteins and other biomolecules, those requirements have proven prohibitive.

If shrunk down to the nanoscale, magnetic resonance imaging (MRI), the common medical imaging technique, could offer a powerful tool for probing molecules that don't crystallize or dissolve. Inspired by that approach, Daniel Rugar and colleagues at IBM's Almaden Research Center in San Jose, California, developed magnetic resonance force microscopy (see PHYSICS TODAY, September 2004, page 21). With it, one can image single virus particles with 10-nm resolution.[1] But it requires cryogenic temperatures that rule out any possibility of observing proteins in their natural environment.

Now two teams have taken a first step toward room-temperature nanoscale MRI by making NMR measurements of nanometer-sized volumes. One team was led by a collaboration[2] of Rugar and David Awschalom (University of California, Santa Barbara). The other was led by Friedemann Reinhard and Jörg Wrachtrup of the University of Stuttgart in Germany.[3] Both teams used nitrogen–vacancy (NV) centers—point defects in diamond best known for their promise as qubits, the potential building blocks of an eventual quantum computer—as sensitive magnetometers.

Magnetic resonance

Bulk NMR and MRI are two different ways of gaining information about a sample from the resonance frequencies of magnetic nuclei (hydrogen-1, carbon-13, and others) in a magnetic field. That resonance frequency, in the RF, corresponds to the energy splitting between nuclear spin states, which depends on the nuclear gyromagnetic ratio and the local magnetic field felt by a nucleus. In bulk NMR, the goal is to pick up the subtle internal magnetic fields that combine with an applied field and affect the total field. Those internal fields, from surrounding electrons and nearby nuclei, give information about the chemical environment of the nucleus and thus the structure of the molecule.

On the other hand, MRI uses magnetic resonance signals to determine the spatial distribution of nuclei with more or less identical chemical environments—^{1}H nuclei in water, for example. In an inhomogeneous applied magnetic field, nuclei in different locations show up at different frequencies, so a frequency spectrum provides spatial information.

Simply by shrinking the conventional MRI apparatus, researchers can resolve features as small as a few microns, but so far no smaller. To image the subnanometer features inside single molecules, new tools are needed.

Enter the NV center, a point defect in diamond that comprises a nitrogen atom (in place of a carbon atom) adjacent to a lattice vacancy. NV centers are well studied for their spin coherence properties: The defect's spin-1 electronic ground state can be prepared in a specific superposition of its $|0\rangle$, $|1\rangle$, and $|-1\rangle$ sublevels and left there for many hundreds of microseconds, uninfluenced by any noise from its environment. When excited with an optical-frequency laser, the $|0\rangle$ state is 30% more likely to fluoresce than the $|1\rangle$ or $|-1\rangle$ state. That offers an easy way to read the NV center's state, at least in an experiment that can be repeated enough times to accumulate sufficient statistics.

Those characteristics make NV centers good qubits. They also make them sensitive and tiny magnetometers. When an NV center initially in the $|0\rangle$ state is hit with a so-called $\pi/2$ microwave pulse, it is transferred into an equal superposition of the $|0\rangle$ and $|1\rangle$ states. The relative phase of those two components evolves at a rate proportional to the local magnetic field. When the superposition is "undone" with an opposite $\pi/2$ pulse, the probability that the NV center returns to the $|0\rangle$ state depends on the total accumulated phase. In a suitably designed experiment, that phase can be made to depend on the quantity of interest: the abundance of nearby atomic nuclei with a given resonance frequency.

Already, NV-center NMR has been used to detect ^{13}C nuclei elsewhere in the diamond itself. To be useful for applications, though, nanoscale MRI needs to be sensitive to nuclei in arbitrary samples outside the diamond. Toward that end, both teams used an experimental setup like the one in figure 1: An NV center a few nanometers beneath the diamond surface could respond to the ^{1}H spins in an organic sample just above the surface. The dimensions of the detection volume roughly equal the NV center depth: 20 nm for the California team's experiment, 5 nm for the Stuttgart group's. In all cases, when the sample was removed, the ^{1}H signal went away.

Frequency and time

The California team's setup included a microwire, lithographically added to the diamond surface, that could generate both RF pulses to manipulate the

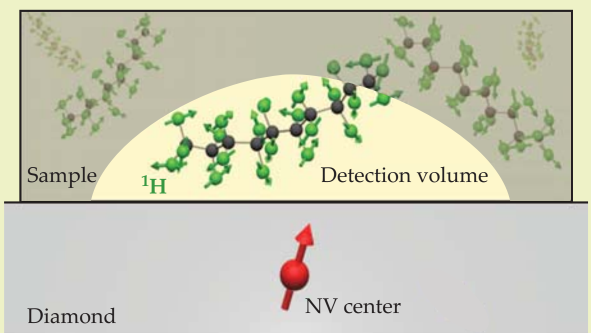

Figure 1. A nitrogen–vacancy (NV) center a few nanometers beneath a diamond surface can sense the tiny magnetic fields of the hydrogen-1 nuclei in a nearby organic sample. (Adapted from ref. 3.)

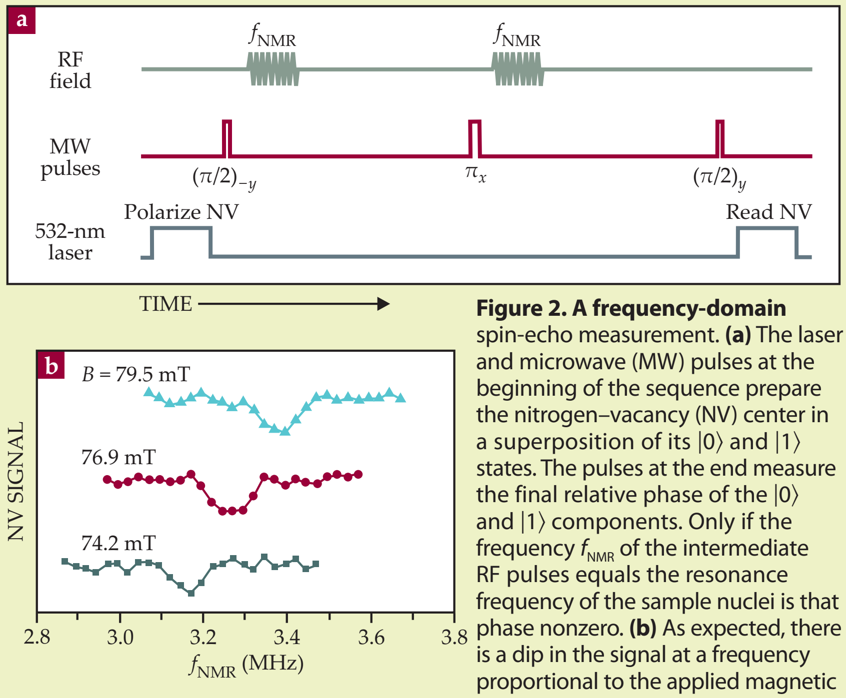

Figure 2. A frequency-domain spin-echo measurement. (a) The laser and microwave (MW) pulses at the beginning of the sequence prepare the nitrogen–vacancy (NV) center in a superposition of its $|0\rangle$ and $|1\rangle$ states. The pulses at the end measure the final relative phase of the $|0\rangle$ and $|1\rangle$ components. Only if the frequency $f_{\rm NMR}$ of the intermediate RF pulses equals the resonance frequency of the sample nuclei is that phase nonzero. **(b)** As expected, there is a dip in the signal at a frequency proportional to the applied magnetic field *B*. (Adapted from ref. 2.)

sample's nuclear spins and microwave pulses to manipulate the NV center's electronic spin. That allowed them to do a spin-echo measurement, as shown in figure 2a, reminiscent of conventional NMR. Midway between the $\pi/2$ pulses, they applied a microwave π pulse (with a different polarization) that essentially reversed the sign of the phase accumulated during the first half of the sequence. If the field felt by the NV center is the same throughout the sequence, the total accumulated phase at the end should therefore be zero, so an opposite $\pi/2$ pulse should reliably return the NV center to the $|0\rangle$ state.

But along with the π pulse, they applied an RF pulse. If that pulse is out of resonance with the ^{1}H nuclei in the sample, it has no effect. But if it's in resonance, it flips the spins of the sample nuclei, so the field felt by the NV center is not the same during the two halves of the sequence. (To guard against the possibility that the RF pulse has some spurious effect on the NV state, they applied an identical pulse at the beginning of the sequence.) From the NV final state as a function of the RF pulse frequency, they should thus be able to detect the sample protons' resonance frequency.

Sure enough, as shown in figure 2b, they saw a dip in the spin-echo signal, the position of which was proportional to the applied magnetic field. As expected, the constant of proportionality was the gyromagnetic ratio of the ^{1}H nucleus.

The Stuttgart group's setup lacked the microwire on the diamond surface. They used an external resonator to create the microwave pulses to drive the NV center, but they aren't yet able to manipulate the sample spins. Instead, they designed a sequence of microwave pulses that, placed between the two $\pi/2$ pulses, created an accumulated phase that's sensitive to the spontaneous precession of the randomly oriented spins. The sequence comprises about 100 π pulses, equally spaced by a delay τ. By varying τ and then transforming their data from the time domain into the frequency domain, they found a signal peak corresponding to the sample nuclei's resonance frequency.

The California team could do time-domain measurements as well, by breaking their RF pulses into two equal halves separated by a delay τ. The half pulses interfered either constructively or destructively, depending on whether τ was a multiple or half multiple of the resonant period (see the article by Norman F. Ramsey in Physics Today, July 1980, page 25; reprinted in January 2013, page 36). That approach gave them an even sharper resonant peak.

What lies ahead

Detecting a single ^{1}H peak in a nanometer-sized sample is just the first step, and nanoscale MRI is still a long way off. To produce images, the researchers need a way to take measurements at different points in the sample. They could apply an inhomogeneous magnetic field, as in conventional MRI. Or they could attach the NV center to a scanning probe tip. The Stuttgart researchers have already built the setup for the latter approach, Reinhard says, "so I hope the first proof of principle won't take too long."

But the bigger challenge, says Rugar, lies in improving the signal-to-noise ratio enough to detect small numbers of nuclei. At present, the California team's detection volume contains about a million nuclei; the Stuttgart group's, about 10 000. Moving the NV center even closer to the diamond surface would shrink the detection volume, although it would also reduce the defect's coherence time. More efficient detection of the fluorescent photons that reveal the NV center's state would also help. The California researchers estimate that with realistic improvements in both those areas, imaging a molecule atom by atom is a goal within reach.

Johanna Miller

References

1. C. L. Degen et al., *Proc. Natl. Acad. Sci. USA* **106**, 1313 (2009).
2. H. J. Mamin et al., *Science* **339**, 557 (2013).
3. T. Staudacher et al., *Science* **339**, 561 (2013).

Gamma-ray spectra show that supernova remnants create cosmic-ray protons

Spectral signatures of pion decay have long been sought as direct evidence of proton acceleration in supernova shock fronts.

The interstellar cosmic-ray flux is dominated by high-energy protons presumably accelerated by sources within the Milky Way. For decades the best guess as to those intragalactic acceleration sites has been the shock fronts of supernova remnants (SNRs). They're thought to be capable of accelerating ambient protons to energies as high as 10^6 GeV, a million times the proton mass. But the evidence has remained frustratingly indirect.

Now, at last, the collaboration that runs the Large Area Telescope (LAT) aboard the *Fermi Gamma-Ray Space Telescope* has found convincingly direct evidence of proton acceleration to cosmic-ray (CR) energies in two 10 000-year-old SNRs a few thousand light-years away.[1] After four years of data taking and analysis, the gamma-ray spectra of the two SNRs revealed a long-sought signature of abundant pion production, which can't be happening unless the remnants are indeed accelerating protons to CR energies.

Shock-front acceleration

Cosmic-ray protons represent a mean energy density throughout the galaxy roughly equal to that of all the starlight. In 1964 Vitaly Ginzburg argued that the only plausible intragalactic sources abundant and powerful enough to generate so much kinetic energy over time are supernovae. The mechanism by which a supernova might eventually convert a significant fraction of its explosive energy into CR protons had been introduced 15 years earlier by Enrico Fermi. He had conjectured that randomly moving magnetic fields in the interstellar medium could stochastically accelerate protons to CR energies in small increments by bouncing them back and forth for centuries.

The expanding shock fronts of supernova remnants, where the exploding star's ejecta impinge upon the undisturbed interstellar medium, provide strong, turbulent magnetic fields. The Milky Way witnesses a few supernovae per century. Ginzburg pointed out that most of the galaxy's CR flux could be accounted for if something like 10% of a supernova explosion's energy goes to Fermi proton acceleration over perhaps 20 000 years, after which its shock front becomes too attenuated and slow.

Theoretical elaboration of the shock-acceleration scenario in recent decades has strengthened the prediction that many SNRs should be profusely making CR protons. And there is much observational evidence, albeit indirect. For example, the theory predicts that SNR shock fronts should also be accelerating electrons. And indeed, x-ray observations of SNRs show telltale synchrotron radiation from those high-energy electrons.

But electrons represent less than 3% of the galactic CR flux. And because they're so much lighter than protons, they produce gammas much more profusely when they scatter off ambient

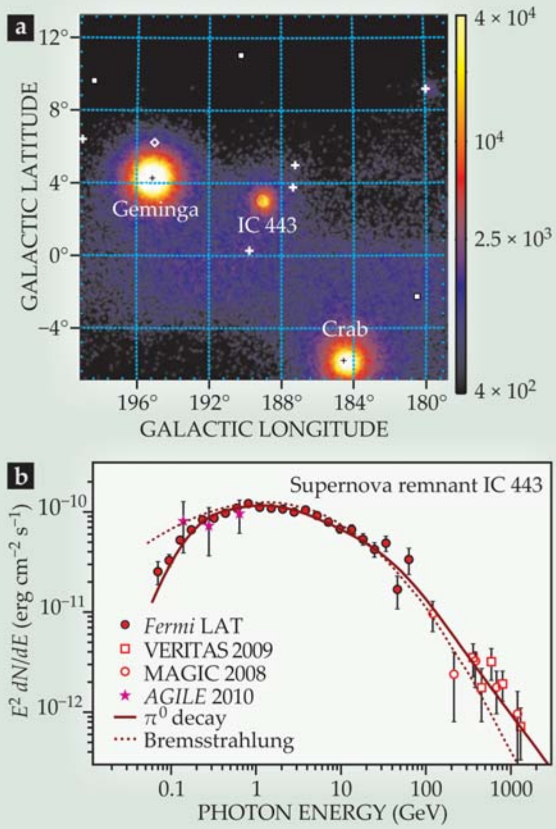

Figure 1. Supernova remnant IC 443, a bright source of gamma radiation. **(a)** Sky map of IC 443's vicinity from observations by the *Fermi Gamma-Ray Space Telescope* is color coded by total gamma-photon counts per $(0.1°)^2$ recorded over four years by *Fermi*'s Large Area Telescope (LAT). Nearby are two other bright sources dominated, unlike IC 443, by pulsar radiation. Marked in white are faint point sources. **(b)** Energy spectrum of IC 443's gamma radiation measured by LAT and other gamma-ray telescopes. The photon-count spectrum dN/dE is multiplied by E^2. Curves show best fits to the LAT data for models assuming only bremsstrahlung from accelerated electrons or only two-gamma decay of neutral pions created by accelerated protons. (Adapted from ref. 1.)