ments tended to suggest solid–liquid coexistence in small temperature windows of "maybe 5 K or 10 K." Those observations fit with the predictions of numerical models that account for finite-size effects that become significant at nanometer length scales.[2] But the models can't explain the 600 K coexistence range observed for the Ga nanodroplets.

Losurdo and her colleagues suspect that the unusual behavior is related to the droplet's interactions with the underlying sapphire support. Because sapphire's crystalline lattice nearly matches that of γ-Ga in size and shape, the partial crystallization of Ga at the sapphire surface would relieve interfacial stress and reduce the droplet's energy—even at temperatures ordinarily too high for the solid to exist. At the droplet's outer surface, the high energetic cost of forming a gas–solid interface would prevent the liquid shell from completely freezing.

When collaborating theorist Kurt Hingerl (Johannes Kepler University Linz) modified standard energy-balance equations to include those surface effects, he got results consistent with the group's working hypothesis. Follow-up experiments provided yet more supporting evidence: Solid cores failed to materialize in Ga droplets deposited on amorphous glass instead of sapphire. Still, Hingerl acknowledges, the team's theoretical model remains incomplete. It doesn't explain, for instance, why only a small fraction of the Ga along the sapphire boundary crystallizes, despite energy balances that indicate an overwhelming preference for the solid–solid interface.

As Gaston sees it, however, the TEM images alone are reason enough for excitement. Interfaces buried in the interiors of metal nanoparticles are both notoriously difficult and important to study. Understanding how they form and evolve is crucial to deploying the particles in phase-change memories, nanoplasmonics, and other applications. (For more on metal nanoparticles, see PHYSICS TODAY, June 2007, page 26.) Losurdo and her colleagues aren't the first group to use TEM to tackle that problem, but they were able to achieve exceptional resolution—high enough, for example, to confirm theoretical predictions that Ga's solid–liquid interfaces are not atomically abrupt but roughly four atoms thick.

"The quality of the experimental images is fantastic," comments Gaston. "Soon we'll be able to combine these kinds of experiments with theory to start sketching out phase diagrams at the nanoscale."

Ashley G. Smart

References

1. M. Losurdo et al., *Nat. Mater.* **15**, 995 (2016).
2. A. Aguado, M. F. Jarrold, *Annu. Rev. Phys. Chem.* **62**, 151 (2011).

Solid-state NMR resolves protein structures—no deuteration required

Spinning the samples at extraordinarily high frequency is key to untangling the biomolecules' complicated spectra.

A murder-mystery villain, if she were resourceful and patient enough, could bring about her victim's demise by replacing all his drinking water with heavy, or deuterated, water. That's because D_2O, although superficially similar to H_2O, behaves differently enough in enzymatic reactions that in large quantities it's detrimental to living tissues. Rats die within a week when given nothing but D_2O to drink.

Escherichia coli bacteria, on the other hand, can survive in D_2O. That's fortunate for protein researchers, who genetically engineer the bacteria to manufacture deuterated proteins to use in nuclear magnetic resonance spectroscopy: Replacing most of a protein's abundant magnetic hydrogen-1 atoms with deuterium, which is invisible to NMR tuned to ^{1}H, helps to simplify the spectrum. To ensure uniform isotopic substitution, researchers first use bacteria to create entirely deuterated proteins, then back-substitute some

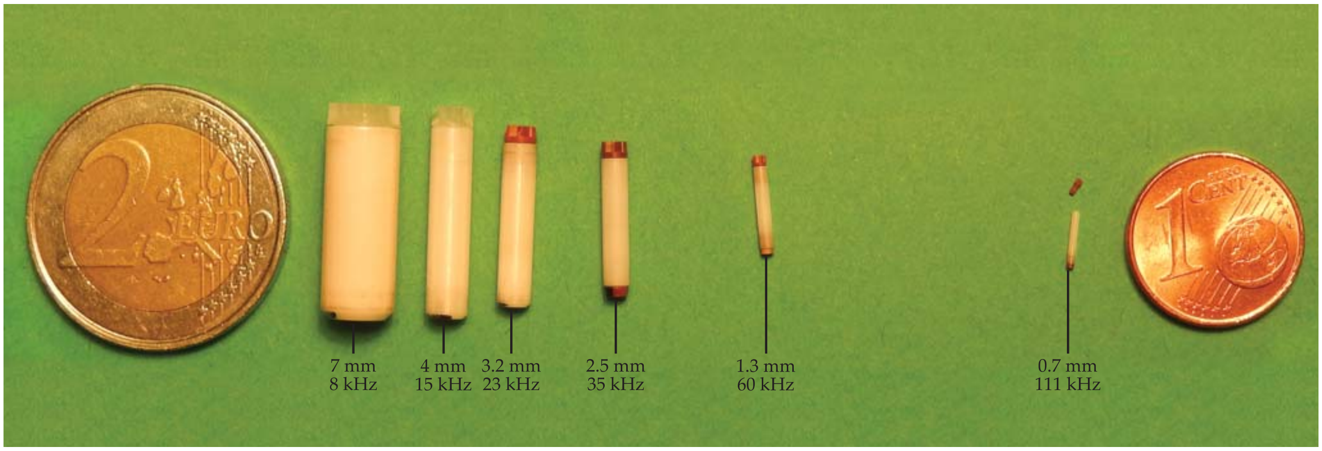

FIGURE 1. IN MAGIC-ANGLE SPINNING NMR, a sample is loaded into a cylindrical rotor that is spun rapidly to mimic the effects of molecules tumbling in a fluid. As shown here, the maximum spinning frequency depends on the rotor diameter. The smallest rotor, 0.7 mm in diameter, was recently debuted. (Courtesy of Guido Pintacuda.)

of the D atoms with ^{1}H at predetermined chemical locations.

Although D_2O doesn't kill the bacteria, it slows their metabolism, and they produce protein only in minute amounts. Furthermore, the back-protonation step isn't easy. As a result, partially deuterated samples are difficult and expensive to make—and for many proteins, not even possible.

Now Guido Pintacuda and colleagues at the European Center for High-Field NMR in Lyon, France, have used ^{1}H NMR to find the structures of two nondeuterated proteins, including a virus coat protein whose structure was previously unknown.[1] To narrow and separate the spectral peaks, the researchers were aided by some of the most advanced equipment in the world, including an NMR spectrometer with a magnetic field of more than 23 T. More importantly, they took a decades-old technique called magic-angle spinning (MAS) and boosted its effectiveness by rotating their samples at speeds of more than 100 kHz. Such high spinning frequencies became possible only a few years ago.

Magnetic dynamics

NMR is widely used to elucidate the structures of molecules both large and small. When placed in a magnetic field, a molecule's magnetic nuclei precess at a field-dependent frequency. But the local field felt by each nucleus isn't exactly equal to the applied field: It's slightly modified, for example, by the swirling cloud of electrons surrounding the nucleus and by proximity to other magnetic nuclei. The spectrum of those deviations, called chemical shifts and measured in parts per million, provides information about the specific environment of each nucleus and thus about the molecule's structure.

For a small, simple molecule, a one-dimensional spectrum of chemical shifts may contain enough information to determine the structure. But for larger molecules, researchers typically require multidimensional spectra: Applying specially designed sequences of RF pulses, they probe spin excitations that coherently hop from one atom to another, with each chemical shift recorded on a separate, orthogonal axis. Such a spectrum offers invaluable clues about which of a molecule's atoms—or which of a protein's amino acids—are spatially close enough to transfer a spin excitation between them.

But that's not all. The pulse sequences for multidimensional NMR spectra can be tailored to probe and follow changes in the molecule on time scales of picoseconds to milliseconds—not just changes in spin states, but also movements of atoms. That capability gives NMR one of its biggest advantages over other techniques for investigating protein structures. X-ray crystallography and cryo-electron microscopy, for example, can resolve structures with exquisite resolution. (For more on cryo-EM, see PHYSICS TODAY, August 2016, page 13.) But both techniques are limited to molecules in static environments—immobilized in a crystal or frozen in a sheet of ice. NMR can probe not just a protein's structure but also the dynamic processes that are central to its biological function.

Conventional solution-phase NMR relies on molecules' rapid random tumbling to average out anisotropic dipolar couplings between nuclei. But that averaging doesn't work for solids and solid-like samples, such as viscous, sticky agglomerates of bulky biomolecules. In those cases, each spectral line is broadened into a thicket of unresolved lines arising from molecules with different orientations. Fortunately, researchers can recover the benefits of rapid tumbling by spinning the sample about an axis oriented at a particular "magic" angle, 54.7°, with respect to the field. At that angle, nuclear couplings fall to zero. (See the article by Clare Grey and Robert Tycko, PHYSICS TODAY, September 2009, page 44.)

The higher the spinning frequency, the more effective MAS is at narrowing and resolving spectral lines, but frequency is limited by the size of the cylindrical rotor that houses the sample. Spinning a too-large rotor too quickly can destroy the sample through frictional heating, destabilize the gas-lubricated rotation, or even disintegrate the rotor itself. Figure 1 shows some examples of rotors and their maximum spinning speeds. The smallest one, at 0.7 mm in diameter, was designed for the new experiments by Pintacuda's collaborators at Bruker Biospin in Germany.

The march to higher MAS frequencies has been a decades-long technical challenge. Building and operating a smaller rotor requires miniaturizing everything from the cooling system to the seals that hold the container closed. It also entails inserting a gooey protein sample into a container the size of a mechanical pencil lead. A further challenge: Shrinking the sample itself, all else being equal, lowers the measurement sensitivity.

Isotope exchange

Although the spin precession of any magnetic isotope can be the basis for an NMR spectrum, solution-phase experiments are most often tuned to detect ^{1}H. Its large gyromagnetic ratio leads to strong signals, and its high natural abundance means that samples can be prepared without any special isotopic labeling. But in solid-state NMR, dipole–dipole coupling between ^{1}H nuclei is a major source of anisotropic line broadening that MAS at modest frequencies can't adequately reverse.

As a result, nearly all solid-state NMR experiments were, until recently, based on isotopes other than ^{1}H; for proteins and other biomolecules, that usually meant carbon-13 and nitrogen-15. Because of the low natural abundances of those isotopes, samples had to be isotopically enriched. The process is similar to deuteration: Insert the gene for the protein into bacteria that are fed with isotopically enriched nutrients. Fortunately, bacterial metabolism isn't slowed by ^{13}C and ^{15}N the way it is by D_2O.

Both ^{13}C and ^{15}N have low gyromagnetic ratios. To obtain reasonable signals from those nuclei, researchers needed relatively large samples, which limited MAS frequencies. For example, in 2002 when Hartmut Oschkinat and colleagues obtained the first complete structure of a protein[2] from ^{13}C and ^{15}N MAS NMR, they used several protein samples of 6–10 mg each, packed into 4-mm-diameter rotors spun at just 8–13 kHz. Although faster spinning is possible with a rotor that size, the need for many-milligram samples was thought to place an effective cap on the frequencies that would ever be feasible.

Around the same time, several groups were experimenting with ^{1}H NMR on heavily deuterated samples; in 2007 Chad Rienstra and colleagues used the method to resolve the structure of GB1, a model protein often used to test solid-state NMR techniques.[3] The switch from ^{13}C and ^{15}N to ^{1}H boosted the signal by almost an order of magnitude and motivated the development of smaller rotors and higher MAS frequencies. Much of the progress was spearheaded by Ago Samoson, of the Tallinn University of Technology in Estonia, and his colleagues, who crossed the 100 kHz threshold with a protein sample in 2012.

In 2014 Samoson collaborated with Matthias Ernst, Beat Meier (both at ETH Zürich), and Anja Böckmann (at the University of Lyon) to resolve a protein structure with 100-kHz MAS NMR from two samples of just 0.5 mg mass, each placed in a 0.8 mm rotor.[4] Pintacuda's new work, in comparison, used a single 0.5 mg sample of each protein.

The expense of growing bacteria in D_2O is less daunting when all that's needed is a milligram of protein, but the challenge of back-protonation remains a limitation. Each of a protein's amino acids contains one H site—the so-called amide H, bound to the N atom that puts the "amino" in "amino acid"—that's more easily displaced than the rest. When a deuterated protein comes in contact with H_2O, amide D atoms readily change places with H atoms in the water, while the rest of the protein's D atoms remain in place.

But that process can swap *all* the amide D atoms only if the protein is not yet folded, or if it's prone to quick, stochastic folding and unfolding. For proteins that fold tightly and remain folded, the only amide D atoms that are exchanged are the ones near the surface of the folded structure. Large parts of the protein remain invisible to the NMR study. And the class of tightly folded proteins includes many of particular medical interest, including membrane proteins (which are often the targets of drugs) and virus capsids. Misfolded proteins known as amyloid fibrils, which are associated with Alzheimer's and Parkinson's diseases (see the article by Tuomas Knowles, Michele Vendruscolo, and Christopher Dobson, PHYSICS TODAY, March 2015, page 36), are also best addressed by solid-state NMR in fully protonated form.

Virus frontiers

Pintacuda and colleagues' initial plan was for a proof-of-principle study of the fully protonated form of GB1, the model protein used by Rienstra and colleagues. But on the strength of those results, the researchers decided also to look at the virus protein AP205CP (short for *Acinetobacter* phage 205 coat protein), whose structure had yet to be resolved by any technique. The protein is a dimer consisting of two copies of a sequence of 130 amino acids; 90 copies of the dimer, in turn, make up the virus capsid. And it's an illustrative example of how back-protonation doesn't always work.

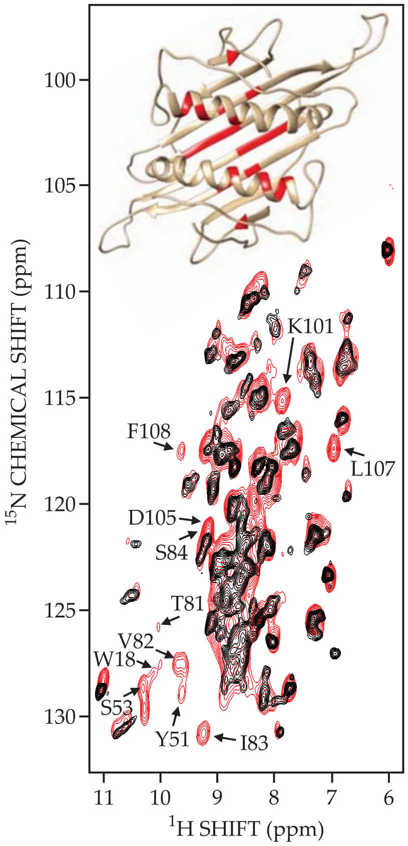

FIGURE 2. TWO-DIMENSIONAL NMR SPECTRA of the virus coat protein AP205CP, partially deuterated (black) and fully protonated (red). The labeled amino-acid spectral peaks in the fully protonated spectrum are missing from the partially deuterated one; for example, "V82" indicates the 82nd amino acid in the protein, which is valine. The missing amino acids' positions are shaded in red in the structural diagram at the top. (Adapted from ref. 1.)

Figure 2 shows two versions of a 2D spectrum of AP205CP's amide H atoms and their associated N atoms. The spectrum of a deuterated, back-protonated sample is shown in black. The spectrum of the fully protonated sample, shown in red, has several peaks that the black spectrum lacks; those additional peaks correspond to amide sites on the deuterated protein that couldn't be exchanged.

Also shown is the NMR protein structure, with the locations of the unexchangeable amides shaded in red.

Pintacuda and colleagues didn't need to deuterate their sample, but they did use ^{13}C and ^{15}N labeling to get information about the positions of C and N atoms. So their method is still limited to proteins that can be grown in laboratory cell cultures—they can't, for example, harvest a protein sample directly from a human patient. Pintacuda speculates that it could one day be possible to find a solid-state NMR structure of a protein using no isotopic labeling at all, merely the naturally present magnetic nuclei. "But resolving the proton resonances would be much harder," he says, "because we wouldn't have the information from nearby carbons and nitrogens."

Although GB1 and AP205CP are both

PHYSICS UPDATE

These items, with supplementary material, first appeared at www.physicstoday.org.

X-RAY GHOST IMAGING

In its basic form, ghost imaging is a technique that indirectly produces a likeness of an object by combining the information from two light detectors—one that views the object and one that doesn't. By splitting a speckled laser beam so that two beams strike spatially separated locations on a CCD camera screen and by placing an object in one of the beams, as shown here, researchers can calculate the intensity correlations between the two signals as the beams are scanned in synchrony and construct the object's ghost image. Demonstrated in 2002 using visible light, the technique holds promise for remote atmospheric sensing because of its insensitivity to turbulence. Two research groups have now extended ghost imaging to the hard x-ray regime. Daniele Pelliccia (RMIT University in Australia) and his colleagues passed an x-ray beam produced by the European Synchrotron source through a diffracting silicon crystal grating oriented to transmit part of the undiffracted beam at a slight angle from the diffracted beam. An ultrafast camera placed 20 cm downstream of the beamsplitter detected both signals—one of which was partially blocked by the imaging target, a copper wire. Because the system worked with individual synchrotron pulses, retrieving the wire's ghost image took more than just calculating intensity correlations; the group also had to filter out low-frequency artifacts associated with mechanical vibrations induced by the synchrotron pulse rate. By contrast, Hong Yu (Chinese Academy of Sciences) and colleagues periodically moved their test object—a gold film with five slits—into and out of a synchrotron's x-ray beam to create their intensity-correlated signals—no beamsplitter or frequency filtering required. Turbulence is not a problem for x-ray imaging, but radiation exposure is. Pelliccia and colleagues speculate that one could pass a weak beam through fragile tissue but use its intensity correlations with a stronger beam that never touches the tissue to produce a low-noise microscopy image. (D. Pelliccia et al., *Phys. Rev. Lett.*, in press; H. Yu et al., *Phys. Rev. Lett.*, in press.) —RMW

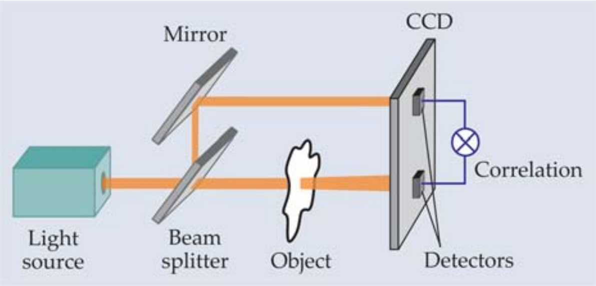

HIGH-VELOCITY CLOUD BLOWS MILKY WAY BUBBLE

Spread through the disk of the Milky Way are about 20 vast bubbles of million-degree plasma wrapped in warm shells of neutral atomic hydrogen. Blowing up a so-called supershell to its typical size of 1000 light-years requires an energy, 10^{45} J, equivalent to at least 30 supernovas. Given that massive short-lived stars cluster in groups of a few dozen, it's indeed possible that supershells are inflated by supernovas. But evidence for an alternative formation scenario has been uncovered recently by Geumsook Park of Seoul National University in Korea and her collaborators. The team used a seven-band receiver at the Arecibo radio telescope in Puerto Rico to resolve the spatial and velocity structure of a supershell known as GS040. At the very center of GS040 they found a high-velocity cloud (HVC). Such clouds contain a million or so solar masses of warm gas and fly through space at speeds 50 km/s higher than is typical for stars and other galactic components. According to previous calculations, when an HVC slams into the galactic disk, it launches a supersonic expansion of hot plasma that is energetic enough to form a supershell. That appears to be what led to the formation of GS040, which, given the shell's current size, took place 5 million years ago. Where did the HVC come from, and how did it acquire its high speed? Park and her colleagues speculate that the HVC is either a fragment of a nearby galaxy that was tidally disrupted by ours or a leftover part of the intergalactic medium from which our galaxy and its neighbors formed. (G. Park et al., *Astrophys. J. Lett.* **827**, L27, 2016.) —CD

PLASMA DISCHARGE FOR FOOD STERILIZATION

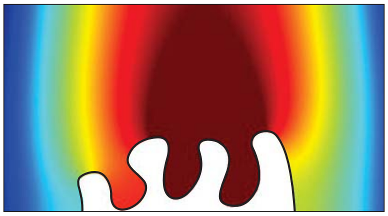

Nonequilibrium plasmas, generated at atmospheric pressure through electric discharge, are in wide commercial use, especially for cleaning surfaces and improving surface adhesion. By producing a range of reactive chemical species, both charged and neutral, they can also inactivate or kill bacteria and other pathogens. When bacteria colonize the surface of fresh fruits and leafy vegetables, they encase themselves in a protective layer known as a biofilm. Plasmas' ability to breach biofilms has caught the attention of the food industry. Yet for plasmas to be deployed as food decontamination tools, their ability to do the job must be understood, controllable, and reproducible. In one step toward that goal, Dawei (David) Liu, Xinpei Lu, and colleagues at China's Huazhong University of Science and Technology and at Shanghai Jiao Tong University have now used a two-dimensional discharge model to analyze the interactions between the plasma and a biofilm on the surface of an apple. In the simulations, which spanned a region 3 mm wide by 5 mm high, the apple served as a floating electrode, and a series of small proteins, high-frequency MAS NMR can potentially be used on much larger molecules. The main challenge is measurement sensitivity: The larger the protein, the fewer copies of it there are in a 0.5 mg sample, and the weaker the resulting signals. Doubling the size of the molecule under study would require an experiment four times as long. But each of Pintacuda and colleagues' structures was found with less than two weeks of data collection. There is plenty of room to expand.

Johanna Miller

References

1. L. B. Andreas et al., *Proc. Natl. Acad. Sci. USA* **113**, 9187 (2016).
2. F. Castellani et al., *Nature* **420**, 98 (2002).
3. D. H. Zhou, *Angew. Chem. Int. Ed.* **46**, 8380 (2007).
4. V. Agarwal et al., *Angew. Chem. Int. Ed.* **53**, 12253 (2014).

---

nanosecond pulses of −25 kV applied to a full-width electrode 4 mm above it triggered narrow plasma filaments known as streamers. To calculate the details of the streamers' propagation and interactions with the 400-μm-wide, 100-μm-thick biofilm, shown here in white, the researchers self-consistently accounted for 115 reactions involving 18 different chemical species, photoionization, and biofilm structure and conductivity. They found that localized ionization occurs above the biofilm, and the plasma spreads into the biofilm recesses. Although the biofilm's irregular shape may shield some nooks from reactive chemical species, diffusion between the plasma pulses evens out the distribution to a level that could satisfy decontamination thresholds. (H. Cheng et al., *Phys. Plasmas* **23**, 073517, 2016.) —RJF

DEUTERON JOINS PROTON AS SMALLER THAN EXPECTED

According to the international Committee on Data for Science and Technology (CODATA), the charge radius of the proton is 0.8768(69) fm. Few researchers would give that number much thought if not for measurements in 2010 and 2013 that yielded a radius 4% smaller than and 7.2 standard deviations distant from the CODATA value. Randolf Pohl of the Max Planck Institute of Quantum Optics in Garching, Germany, and colleagues obtained the curiously low radius after analyzing the energy-level shifts of muons orbiting hydrogen nuclei. With a mass 207 times that of the electron, a muon has a tighter orbital that more closely overlaps the nuclear charge distribution, which makes the negatively charged particle a useful tool for probing nuclear dimensions. The discrepancy between the results of muon-based and other experimental investigations has come to be known as the proton radius puzzle.

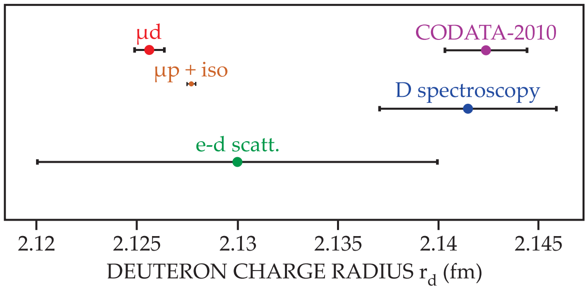

Now Pohl and his colleagues have used the same technique to measure the radius of the deuteron, a nucleus of one proton and one neutron. The researchers shot a beam of muons at a target of D_2 gas. Lasers excited some of the atoms whose electrons were replaced by muons and probed the muons' energy-level transitions. By combining the measurements with theory, the researchers came up with a deuteron charge radius of 2.12562(78) fm. That's 7.5 σ smaller than the CODATA value (see graph; the new result is in red). In addition, both the proton and deuteron sizes are in tension with the values obtained by applying the same technique to atoms with electrons rather than muons.

The new study reinforces the notion that something is amiss in our understanding of particle or atomic physics. The most tantalizing possibility is that the standard model is wrong—perhaps muons interact with other particles differently than electrons do, for example. Pohl considers that explanation unlikely. His group and others are conducting experiments to precisely measure the Rydberg constant, which, if favorably reevaluated, could resolve the discrepancy. (R. Pohl et al., *Science* **353**, 669, 2016.) —AG

UNRAVELING THE JET-LAG ASYMMETRY

Jet lag, the sluggishness we feel after landing in a new time zone, has a directional bias: Studies suggest it takes longer to recover from eastward travel than from westward travel. A new model developed by Michelle Girvan, Edward Ott, Thomas Antonsen, and colleagues at the University of Maryland, College Park, may explain why. The team used tools of nonlinear dynamics to model a region of the brain known as the suprachiasmatic nucleus (SCN), a network of roughly 20 000 time-keeping neurons devoted to maintaining the body's circadian rhythm. Jet lag happens when those neurons fall out of sync with the local cycle of night and day. Modeling the dynamics of the SCN, however, is a tall order. Although the neurons all take cues from the same source—the retina—they don't all respond in the same way. Nor do they share the same natural oscillation periods; absent visual cues, the periods of the average person's SCN neurons would be distributed around an average value slightly longer than a day—about 24.5 hours. In 2008 Ott and Antonsen devised a way to represent such heterogeneous networks of oscillators in terms of just a few key variables. Now they've applied that approach to the SCN. Their model predicts that the slight deviation of the neurons' natural periods from 24 hours can lead to large recovery-time asymmetries. The worst-case scenario turns out to be an eastward trip across nine time zones: An Angeleno arriving in Paris would require six more days to recover than would a Parisian landing in Los Angeles. (Z. Lu et al., *Chaos* **26**, 094811, 2016.) —AGS PT