

Slow-motion spectroscopy paves the way for a nuclear clock

For the first time, the energy level of an atomic nucleus has been directly manipulated with a laser.

Jules Janssen was a 19th-century umbraphile: He traveled the world to witness solar eclipses. During one such event in 1868 in India, he saw something that would help to shape the understanding of both the Sun's inner workings and the periodic table. When he sent the highly attenuated sunlight through his spectrometer, he found a yellow line that no one had ever seen before. It turned out to be an atomic emission from an as-yet undiscovered element: helium.

The discrete colors of light absorbed and emitted by atoms have long been integral to our understanding of the physical world. A few decades after Janssen's discovery, the pioneers of quantum mechanics looked to atomic energy levels, deduced from their spectra, to work out the laws that govern the dynamics of electrons in atoms. And modern atomic clocks, the best of which waver by less than a second over the age of the universe, are the basis for GPS navigation, tests of general relativity, and the very definitions of the fundamental units of measure. (See the article by David Newell, *PHYSICS TODAY*, July 2014, page 35, and the Quick Study by Emily Caldwell and Laura Sinclair on page 54 of this issue.)

Atomic nucleons, like electrons, organize into energy shells, and nuclei, like atoms, can be promoted between discrete energy states. But there's no nuclear optical spectroscopy, and there aren't yet any nuclear clocks. The problem is that almost all nuclear transitions lie at impractically high energies: in the realm of gamma rays, not visible light, and far beyond the reach of tabletop lasers.

In fact, out of all the known nuclides, only one—the rare, radioactive isotope thorium-229—hosts an excited state, denoted $^{229\text{m}}\text{Th}$, that lies within a handful of electron volts of the ground state. Laser

spectroscopy of nuclei will never span the periodic table the way atomic spectroscopy does. But for the purpose of building a nuclear analogue of an atomic clock, one transition of the right energy is all that's needed. And ^{229}Th will do the job.

Now a team of researchers led by Thorsten Schumm (Technical University of Vienna) and Ekkehard Peik (PTB, the National Metrology Institute of Germany) has succeeded in exciting ^{229}Th with a laser.¹ Meanwhile, at UCLA, Eric Hudson and colleagues have achieved similar results.² Laser excitation is just one step on the path to a nuclear clock, but it's an important one. And the hurdles encountered along the way—including finding the transition energy in the first place—highlight the differences between nuclear and atomic physics.

The nuclear option

The heart of an atomic clock is a spectroscopic frequency measurement: The clock's ticks are the electromagnetic cycles associated with an atomic energy transition. If precision timekeeping is the goal, it's necessary to tamp down all sources of uncertainty in the measurement, such as Doppler broadening, which can make the transition look bluer or redder depending on whether the atom is moving toward or away from the observer.

For those uncertainty-mitigation efforts to do any good, it's necessary to start with a transition with an inherently narrow linewidth. By the Heisenberg uncertainty principle, that means the excited state must have a long radiative lifetime. Most atomic transitions won't work: Their excited states relax back to the ground state in a fraction of a second.

In contrast, the $^{229\text{m}}\text{Th}$ state's radiative lifetime is comparatively long—minutes to hours, depending on the environment—but not extraordinarily so: Plenty of atomic states have even longer lifetimes. An excited state in ionic ytterbium, used in some state-of-the-art clocks, has a lifetime of more than a year. And the excited state of the microwave-frequency cesium transition that's the basis for the definition of the second has a lifetime so long that it defies measurement.

But the metrological appeal of the ^{229}Th transition—remarked on in 1996 by Eugene Tkalya and colleagues³ and developed into a full-fledged nuclear clock proposal in 2003 by Peik and Christian Tamm⁴—goes beyond its narrow inherent linewidth. For example, nuclear energy levels are less susceptible than electronic states to distortion by stray electromagnetic fields.

Furthermore, the nuclei used in a nuclear clock can be embedded in a solid. That's not possible for atomic clocks: The chemical bonds of a crystal lattice would hopelessly alter the energy of the electronic transition. All atomic clocks, therefore, use atoms held station-




FIGURE 1. THORIUM-229 ATOMS, shown in light blue in this artist's impression, are held in the lattice of a calcium fluoride crystal. The isotope is prized for its unusually low-lying nuclear excited state, the transition to which can now be driven with a laser. (Image by Oliver Diekmann, Technical University of Vienna.)

ary in the gas phase, such as in electromagnetic or optical traps. The option to use a crystal instead makes it easier to hold large numbers of nuclei in a small space, and it could make nuclear clocks more portable.

Lost in transition

The prospect of a thorium nuclear clock faced a barrier that would be unheard of in the world of atomic clocks: Nobody knew the energy of the ^{229}Th nuclear transition, not even approximately.

It's been known for decades that the $^{229\text{m}}\text{Th}$ state exists. When ^{229}Th is produced from the radioactive decay of other elements—such as the alpha decay of uranium-233—the pattern of the gamma rays that are released makes it clear that the isotope is being produced in two states of nearly equal energy. But finding the energy difference between the two states would have meant measuring the difference between two large gamma-ray energies, and nuclear experiments aren't capable of measuring things so finely.

Theory, likewise, wasn't equipped to grapple with such small energies. So in the 1976 paper⁵ that first recognized the low-energy state, the best the authors could say was that the energy difference was something less than 100 eV—that is, something less than the energy of a 12 nm soft x ray. That's not much help.

Experimenters would need the energy information to have any hope of exciting the nuclear transition directly. The usual spectroscopic approach of exciting atoms with broadband light and looking for photons absorbed or emitted wasn't going to work on ^{229}Th . The long radiative lifetime of the $^{229\text{m}}\text{Th}$ state means that only a trickle of photons is being emitted at any given time.

Meanwhile, radiative decay is competing with radioactive decay. The ^{229}Th isotope's 8000-year radioactive half-life dwarfs the $^{229\text{m}}\text{Th}$ state's radiative lifetime—but each radioactive decay produces thousands of photons, not just one. Only with a tightly concentrated laser frequency would researchers have a hope of exciting enough nuclei for the radiative signal to outcompete the radioactive background. Scanning the whole 100 eV range would take forever.

Happily, theory and experiment improved, and researchers started to home in on the ^{229}Th transition energy. Unhap-




FIGURE 2. LASER EXCITATION of the thorium-229 nucleus looks different from different directions. To efficiently scan the range of possible wavelengths (only a small portion of which is shown here), Thorsten Schumm, Ekkehard Peik, and colleagues recorded these spectra at a rate of five minutes per data point. But the $^{229\text{m}}\text{Th}$ excited state's lifetime is close to 10 minutes, so the spectral line is skewed. (Adapted from ref. 1.)

pily, for a while at least, they were homing in on the wrong answer. In the late 1990s and early 2000s, when Tkalya, Peik, and others were formulating ideas for a nuclear clock, the best guess was that the transition was somewhere around 3.5 eV, or 354 nm. If true, that would have been fortunate: Light of that wavelength propagates not only in air but also in plenty of crystals, including some pure thorium compounds, such as thorium oxide. But all searches for the transition in that vicinity turned up empty—because it wasn't there.

The actual energy, it would turn out, was more than twice as high: not in the near UV but in the vacuum UV. Thorium oxide is opaque in that range, as are most crystals. But a few wide-bandgap materials, such as calcium fluoride, are transparent, and they can be doped with thorium, as illustrated in figure 1.

Moreover, vacuum-UV light, as the name suggests, doesn't propagate in air. The whole experiment, from laser to sample to detector, would need to be placed under vacuum—a change that would prompt a rethinking of every aspect of the experimental design. "It becomes harder and harder to build a good laser when the wavelength becomes shorter," says Peik. "And you can't just reach in and tweak things with your hands."

Great catch

A breakthrough came last year, when the ISOLDE group at CERN published the first observations of photons emitted from the thorium radiative transition.⁶

The ISOLDE researchers did it by peppering a calcium fluoride crystal with radioactive precursors to ^{229}Th . Some of the precursors would decay into the $^{229\text{m}}\text{Th}$ state, which then relaxed to the ground state. That approach had been tried before, but always using ^{233}U , which decays to ^{229}Th by emitting alpha particles that carry significant momentum and could have been damaging the crystal. The ISOLDE researchers instead created ^{229}Th through the beta decay of actinium-299, a gentler process that allowed them to finally observe $^{229\text{m}}\text{Th}$'s radiative decay.

The ISOLDE observation narrowed the transition energy down to 8.338 ± 0.024 eV, or a wavelength somewhere between 148.3 and 149.1 nm. At last, researchers could optimize their lasers and begin their search in earnest. But scanning even that narrow range was a painstaking process. Both Peik and Schumm's group and Hudson's used the same basic approach: Illuminate a ^{229}Th -doped crystal with a laser for a while, then switch off the laser and look for any photons emitted by nuclei that had been excited by the laser. Repeat the process with different laser wavelengths until the entire energy range is covered.

Both research groups used lasers with bandwidths of around 10 GHz, or 40 μeV , so they needed some 1000 data points to span ISOLDE's energy range. How long does that take? It depends on how patient you are. In a crystal, $^{229\text{m}}\text{Th}$ has a radiative lifetime of around 10 minutes, but

Get **nervous**
before a job
interview?

You're **not**
alone.



Physics Today Jobs has several resources on our website for job seekers, including recorded webinars led by experts to help you navigate the job interview process with ease.

Find your future at
physicstoday.org/jobs

PHYSICS TODAY

that's the half-life—only half of the excited nuclei decay in that time. To reset all the nuclei back to the ground state, you'd need to wait for several half-lives.

Hudson and colleagues took that approach, doing their laser scan at a rate of an hour per data point, or 1000 hours of data-taking to cover the whole ISOLDE range. Schumm, Peik, and colleagues took a more expedient approach—although still sluggish by spectroscopic standards—of five minutes per data point. As a result, they were the first group to find the excitation. But because the signal they observed at each step included nuclei that had been excited in previous steps, the spectral line they observed was skewed, as shown in figure 2. To compensate, they scanned the range in both directions.

Schumm and Peik's team measured the resonance wavelength to be 148.3821 ± 0.0005 nm; Hudson and colleagues found it to be 148.3822 ± 0.0002 nm. The agreement lends credence to both measurements, especially because the two experiments used different crystal materials: calcium fluoride in Schumm and Peik's experiment, and lithium strontium aluminum fluoride in Hudson's. There was an outside chance that the spectroscopic feature could have been an artifact, perhaps the result of a crystal defect created by thorium's radioactivity. But the odds of a spurious signal at exactly the same wavelength in two different materials are slim to none.

Fine structure

There's still a long way to go on the path to a nuclear clock. To achieve precision of a part in 10^{18} , as the best atomic clocks do, researchers would need a wavelength measurement that's accurate to the 15th decimal place, not the 4th. The next step, therefore, is to repeat the excitation measurement using narrower-band lasers. There's plenty of room for improvement. The inherent linewidth of the ^{229}Th transition is thought to be smaller than 1 mHz, more than 13 orders of magnitude narrower than the 10 GHz laser bandwidth used in the initial search.

As researchers zero in on the transition wavelength, a problem that's plagued the search so far—background due to radioactive decay—should naturally resolve itself. With a 10 GHz bandwidth, the vast majority of the

laser power is always wasted: Even when the laser is centered on the transition wavelength, most of the light is too far off resonance to excite any nuclei. With a narrower bandwidth, more light can contribute to exciting the nuclei. "For the same laser power, we get a much stronger signal," says Schumm. "Then we can reduce the doping concentration to reduce the radioactive background."

More than two decades after they were initially proposed, nuclear clocks are looking like an achievable goal. "But atomic clocks are a moving target," notes Peik. "Researchers have made so much progress in recent years, and they've solved a lot of the problems that the nuclear clock was designed to circumvent." From a purely timekeeping perspective, it remains to be seen whether nuclear clocks will ever get the upper hand.

But there may be more subtle benefits to a clock based on nuclear rather than atomic physics. For example, one way that researchers test whether the fundamental constants have changed over time is to compare two clocks based on different atomic transitions. If the fine-structure constant, say, isn't actually constant in time, the clocks' relative tick rates would slowly drift.

So far, no such drift has been observed—but that could be either because the fine-structure constant isn't changing or because its change is too small to be observed. Because of the overall larger energy scale of nuclear physics, a nuclear clock should be orders of magnitude more responsive than an atomic clock to any changes in the fundamental constants. So researchers could have a chance at spotting a drift in the fine-structure constant that they otherwise would never observe in a million years.

Johanna Miller

References

1. J. Tiedau et al., *Phys. Rev. Lett.* **132**, 182501 (2024).
2. R. Elwell et al., <https://arxiv.org/abs/2404.12311>.
3. E. V. Tkalya et al., *Phys. Scripta* **53**, 296 (1996).
4. E. Peik, C. Tamm, *Europhys. Lett.* **61**, 181 (2003).
5. L. A. Kroger, C. W. Reich, *Nucl. Phys. A* **259**, 29 (1976).
6. S. Kraemer et al., *Nature* **617**, 706 (2023).