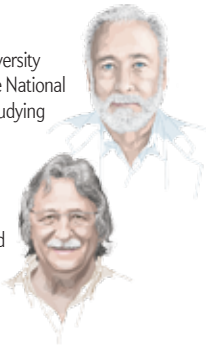
Geoffrey L. Greene is a professor of physics at the University of Tennessee, with a joint appointment at the Oak Ridge National Laboratory's Spallation Neutron Source. He has been studying the properties of the neutron for more than 40 years.

Peter Geltenbort is a staff scientist at the Institut Laue-Langevin in Grenoble, France, where he uses one of the most intense neutron sources in the world to research the fundamental nature of this particle.



LUCKILY FOR LIFE ON EARTH, MOST MATTER IS NOT RADIOACTIVE. WE TAKE THIS FACT FOR granted, but it is actually somewhat surprising because the neutron, one of the two components of atomic nuclei (along with the proton), is prone to radioactive decay. Inside an atomic nucleus, a typical neutron can survive for a very long time and may never decay, but on its own, it will transform into other particles within 15 minutes, more or less. The words "more or less" cover a disturbing gap in physicists' understanding of this particle. Try as we might, we have not been able to accurately measure the neutron lifetime.

This "neutron lifetime puzzle" is not just embarrassing for us experimentalists; resolving it is vital for understanding the nature of the universe. The neutron decay process is one of the simplest examples of the nuclear "weak" interaction—one of nature's four fundamental forces. To truly understand the weak force, we must know how long neutrons live. Furthermore, the survival time of the neutron determined how the lightest chemical elements first formed after the big bang. Cosmologists would like to calculate the expected abundances of the elements and compare them with astrophysical measurements: agreement would confirm our theoretical picture, and discrepancy could indicate that undiscovered phenomena affected the process. To make such a comparison, however, we need to know the neutron lifetime.

More than 10 years ago two experimental groups, one a Russian-led team in France and the other a team in the U.S., attempted separately to precisely measure the lifetime. One of us (Geltenbort) was a member of the first team, and the other (Greene) was a member of the second. Along with our colleagues, we were surprised and somewhat disturbed to find that our results disagreed considerably. Some theoreticians suggested that the difference arose from exotic physics—that some neutrons in the experiments might have transformed into particles never before detected, which would have affected the different experiments in divergent ways. We, however, suspected a more mundane reason—perhaps one of our groups, or even both, had simply made a mistake or, more likely, had overestimated the accuracy of its experiment. The U.S. team recently completed a long, painstaking project to study the most dominant source of uncertainty in its experiment in hopes of resolving the discrepancy. Rather than clearing up the situation, that effort confirmed our earlier result. Similarly, other researchers later confirmed the findings of Geltenbort's team. This discrepancy has left us even more perplexed. But we are not giving up—both groups and others continue to seek answers.

TIMING NEUTRONS

IN THEORY, measuring the neutron lifetime should be straightforward. The physics of nuclear decay are well understood, and we have sophisticated techniques for studying the process. We know, for instance, that if a particle has the possibility of transforming into a lower-mass particle or particles while conserving such characteristics as charge and spin angular momentum, it will. Free neutrons display this instability. In a process called beta decay, a neutron breaks up into a proton, an electron and an antineutrino (the antimatter counterpart of the neutrino), which collectively sum to a slightly lower mass but the same total charge, spin angular momentum and other conserved properties. These conserved properties include "mass-energy," meaning that the daughter particles carry the difference in mass in the form of kinetic energy, the energy of motion.

We cannot predict exactly when a particular neutron will decay because the process is a fundamentally random quantum phenomenon—we can say only how long neutrons live on average. Thus, we must measure the average neutron lifetime by studying the decay of many neutrons.

Investigators have employed two experimental methods—one called the "bottle" technique and the other the "beam" approach. Bottle experiments confine neutrons in a container and count how many are left after a given time. The beam method, in contrast, looks not for the disappearance of neutrons but rather for the appearance of the particles into which they decay.

The bottle approach is particularly challenging because neutrons can pass easily through matter and thus through the walls of most containers. Following a suggestion first explicitly made by Russian physicist Yuri Zel'dovich, experimentalists who use the bottle approach—as Geltenbort and his colleagues in France do—get around the problem by trapping extremely cold neutrons (that is, those with a very low kinetic energy) within a container of very smooth walls [*see box on page 40*]. If the neutrons are slow enough and the bottle smooth enough, they reflect from the walls and hence remain in the bottle. To achieve this effect, the neutrons must move at speeds on the order of just a few meters per second, as opposed to the roughly 10 million meters per second neutrons travel when emitted during nuclear fission, for instance. These "ultracold" neutrons are so slow that you could "outrun"

© 2016 Scientific American

them. The most accurate bottle experiment to date took place at the Institut Laue-Langevin (ILL) in Grenoble, France.

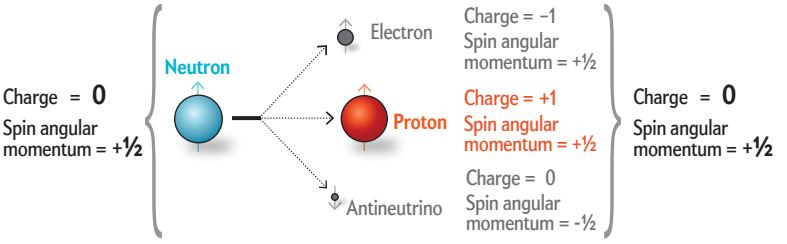
Unfortunately, no bottle is ever perfect. If neutrons occasionally leak out of the bottle, we will attribute this loss to beta decay and will get the wrong lifetime. We must therefore be sure to correct our calculations so as to count only those particles that actually undergo beta decay.

To make that correction, we use a clever technique. The number of neutrons lost through the walls of the bottle depends on the rate at which neutrons bounce against the walls. If the neutrons are slower or the bottle is bigger, the bounce rate, and thus the loss rate, will go down. By varying both the size of the bottle and the energy (velocity) of the neutrons in successive trials, we can extrapolate to a hypothetical bottle in which there are no collisions and thus no wall losses. Of course, this extrapolation is not perfect, but we do our best to account for any error this calculation introduces.

BASICS

How Neutrons Decay

Despite decades of trying, scientists have not been able to definitively measure how long neutrons live outside of atomic nuclei—the best experiments in the world produce conflicting results. Although the length of the neutron lifetime is undetermined, the cause of neutron decay is well known. Through a process called beta decay, a neutron transforms into a proton and releases an electron and an antineutrino, the antimatter counterpart to the neutrino particle. The decay ensures that the final particles' charge and spin angular momentum tally to equal those of the original particle.



In the beam method—used by Greene and others at the National Institute of Standards and Technology (NIST) Center for Neutron Research—we send a stream of cold neutrons through a magnetic field and a ring of high-voltage electrodes that traps positively charged particles [*see box on page 41*]. Because neutrons are electrically neutral, they pass right through the trap. If, however, a neutron decays within the trap, the resulting positively charged proton gets "stuck." Periodically we "open" the trap and expel and count the protons. In principle, the proton trapping and detection are nearly perfect, and we must make only very small corrections for the possibility that we missed decays.

WHERE COULD WE GO WRONG?

TO BE USEFUL, a measurement must be accompanied by a reliable estimate of its accuracy. A measurement of a person's height that has an uncertainty of one meter, for example, is much less meaningful than a measurement that has an uncertainty of one millimeter. For this reason, when we make precision measurements we always report an experimental uncertainty; an uncertainty of one second, for instance, would mean our measurement had a high probability of being no more than a second shorter or a second longer than the true value.

Any measurement has, in general, two sources of uncertainty. Statistical error arises because an experiment can measure only a finite sample—in our case, a finite number of particle decays. The larger the sample, the more reliable the measurement and the lower the statistical error.

The second source of uncertainty—systematic error—is much more difficult to estimate because it arises through imperfections in the measurement process. These flaws may be something simple, like a poorly calibrated meter stick used to measure a person's height. Or they can be more subtle, like a sampling bias—in a telephone poll, for example, one might overly rely on calls to land lines rather than to cell phones and thus fail to capture a truly representative population sample. Experimentalists go to great lengths to reduce these systematic errors, but they are impossible to eradicate completely. The best we can do is carry out a detailed study of all imaginable sources of error and then estimate the lingering effect each might have on the final result. We then add this systematic error to the statistical error to give a best estimate of the overall reliability of the measurement. In other words, we put great effort into estimating the "known unknowns."

Of course, our great fear is that we have overlooked an "unknown unknown"—a systematic effect that we do not even know we do not know—hidden within the experimental procedure. While we go to extreme pains to explore all possible uncertainties, the only way to overcome this type of additional error with real confidence is to perform another, completely independent measurement using a totally different experimental method that does not share the same systematic effects. If two such measurements agree within their quoted uncertainties, we have confidence in the results. If, on the other hand, they disagree, we have a problem.

For the measurement of the neutron lifetime we have two such independent methods: the beam and the bottle. The most recent result from the beam experiment at NIST gave a value for the neutron lifetime of 887.7 seconds. We determined the statistical uncertainty in our estimate to be 1.2 seconds and the systematic uncertainty 1.9 seconds. Combining those errors statistically gives a total uncertainty of 2.2 seconds, which means that we believe the true value of the neutron lifetime has a 68 percent probability of being within 2.2 seconds of the measured value.

The bottle experiment at ILL, on the other hand, measured a neutron lifetime of 878.5 seconds with a statistical uncertainty of 0.7 second, a systematic uncertainty of 0.3 second and a total uncertainty of 0.8 second.

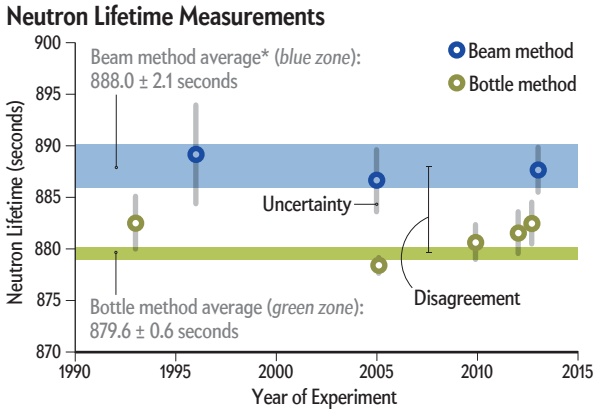
These are the two most precise neutron lifetime experiments of each type in the world, and their measurements differ by approximately nine seconds. Such a time span may not sound like a lot, but it is significantly larger than the calculated uncertainties for both experiments—the probability of obtaining a

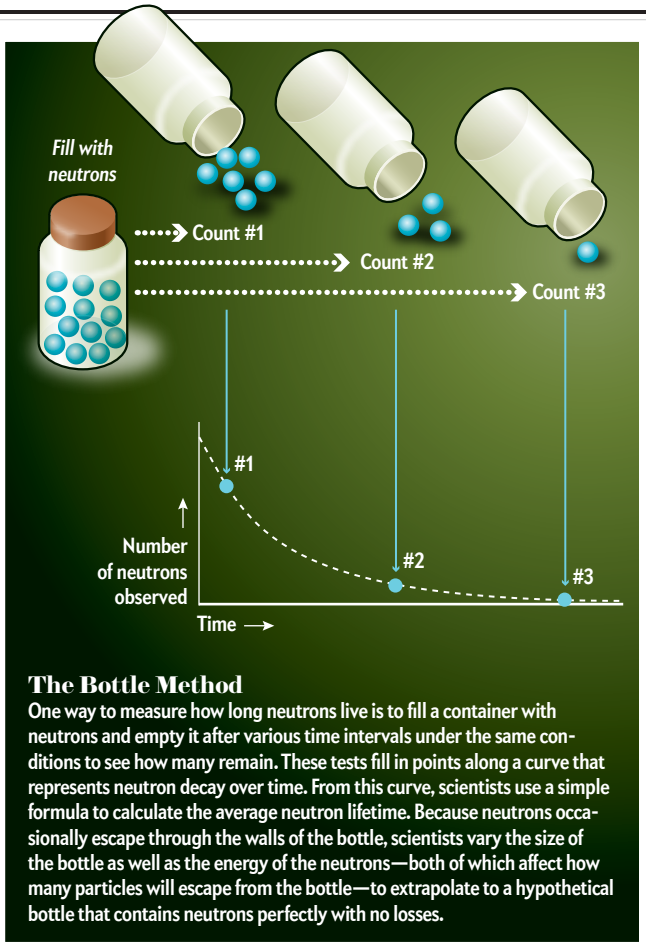
Graphics by Nigel Hawtin

© 2016 Scientific American

Different Techniques, Different Results

Scientists have tried two main techniques to measure the average neutron lifetime: the "bottle" and the "beam" methods. The various bottle measurements over the years tend to agree with one another within their calculated error bars, as do the beam measurements. The results from the two techniques, however, conflict. The discrepancy, about eight seconds between the bottle and beam averages, may not seem like much, but it is significantly larger than the measurements' uncertainty, which means the divergence represents a real problem. Either the researchers have underestimated the uncertainty of their results, or, more exciting, the difference arises from some unknown physical phenomenon.





The Bottle Method
One way to measure how long neutrons live is to fill a container with neutrons and empty it after various time intervals under the same conditions to see how many remain. These tests fill in points along a curve that represents neutron decay over time. From this curve, scientists use a simple formula to calculate the average neutron lifetime. Because neutrons occasionally escape through the walls of the bottle, scientists vary the size of the bottle as well as the energy of the neutrons—both of which affect how many particles will escape from the bottle—to extrapolate to a hypothetical bottle that contains neutrons perfectly with no losses.

*The beam method average does not include the 2005 measurement, which was superseded by the 2013 beam study.

difference of this size by chance alone is less than one part in 10,000. We must therefore seriously consider the possibility that the discrepancy results from an unknown unknown—we have missed something important.

EXOTIC PHYSICS

AN EXCITING explanation for the difference could be that it actually reflects some exotic physical phenomenon not yet discovered. A reason to think such a phenomenon might exist is that although the bottle and beam methods disagree, other beam studies show good agreement among themselves, as do other bottle studies.

Imagine, for example, that in addition to the regular beta decay, neutrons decayed via some previously unknown process that does not create the protons sought in beam experiments. The bottle experiments, which count the total number of "lost" neutrons, would count both the neutrons that disappeared via beta decay as well as those that underwent this second process. We would therefore conclude that the neutron lifetime was shorter than that from "normal" beta decay alone. Meanwhile the beam experiments would dutifully record only beta decays that produce protons and would thus result in a larger value for the lifetime. So far, as we have seen, the beam experiments do measure a slightly longer lifetime than the bottles.

A few theorists have taken this notion seriously. Zurab Berezhiani of the University of L'Aquila in Italy and his colleagues have suggested such a secondary process: a free neutron, they propose, might sometimes transform into a hypothesized "mirror neutron" that no longer interacts with normal matter and would thus seem to disappear. Such mirror matter could contribute to the total amount of dark matter in the universe. Although this idea is quite stimulating, it remains highly speculative. More definitive confirmation of the divergence between the bottle and beam methods of measuring the neutron lifetime is necessary before most physicists would accept a concept as radical as mirror matter.

Much more likely, we think, is that one (or perhaps even both) of the experiments has underestimated or overlooked a systematic effect. Such a possibility is always present when working with delicate and sensitive experimental setups.

WHY THE NEUTRON LIFETIME MATTERS

FIGURING OUT WHAT WE MISSED will of course give us experimentalists peace of mind. But even more important, if we can get to the bottom of this puzzle and precisely measure the neutron lifetime, we may be able to tackle a number of long-standing, fundamental questions about our universe.

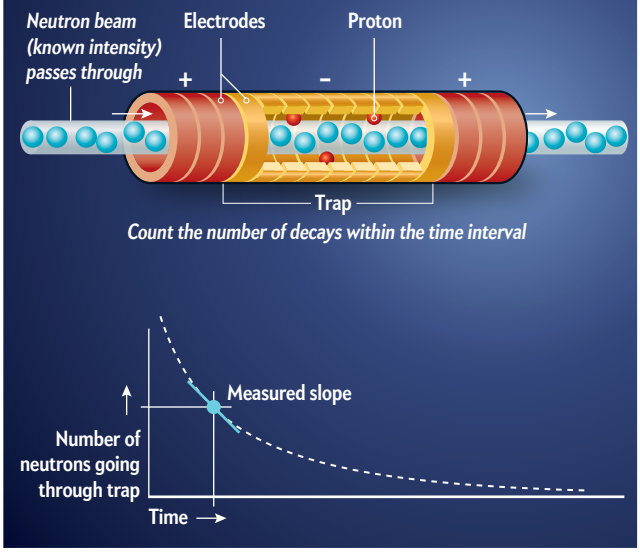
First of all, an accurate assessment of the timescale of neutron decay will teach us about how the weak force works on other particles. The weak force is responsible for nearly all radioactive decays and is the reason, for instance, that nuclear fusion occurs within the sun. Neutron beta decay is one of the simplest and most pure

SCIENTIFIC AMERICAN ONLINE See a video about neutron beta decay at ScientificAmerican.com/apr2016/neutron-lifetime

© 2016 Scientific American

The Beam Method

In contrast to the bottle method, the beam technique looks not for neutrons but for one of their decay products, protons. Scientists direct a stream of neutrons through an electromagnetic "trap" made of a magnetic field and ring-shaped high-voltage electrodes. The neutral neutrons pass right through, but if one decays inside the trap, the resulting positively charged protons will get stuck. The researchers know how many neutrons were in the beam, and they know how long they spent passing through the trap, so by counting the protons in the trap they can measure the number of neutrons that decayed in that span of time. This measurement is the decay rate, which is the slope of the decay curve at a given point in time and which allows the scientists to calculate the average neutron lifetime.



examples of a weak force interaction. To calculate the details of other, more complex nuclear processes involving the weak force, we must first fully understand how it operates in neutron decay.

Discerning the exact rate of neutron decay would also help test the big bang theory for the early evolution of the cosmos. According to the theory, when the universe was about one second old, it consisted of a hot, dense mixture of particles: protons, neutrons, electrons, and others. At this time, the temperature of the universe was roughly 10 billion degrees—so hot that these particles were too energetic to bind together into nuclei or atoms. After about three minutes, the universe expanded and cooled to a temperature where protons and neutrons could stick together to make the simplest atomic nucleus, deuterium (the heavy isotope of hydrogen). From here other simple nuclei were able to form—deuterium could capture a proton to make an isotope of helium, two deuterium nuclei could join together to create heavier helium, and small numbers of larger nuclei formed, up to the element lithium (all the heavier elements are thought to have been produced in stars many millions of years later).

This process is known as big bang nucleosynthesis. If, while the universe was losing heat, neutrons had decayed at a rate that was much faster than the universe cooled, there would have been no neutrons left when the universe reached the right temperature to form nuclei—only the protons would have remained, and we would have a cosmos made almost entirely of hydrogen. On the other hand, if the neutron lifetime were much longer than the time required to cool sufficiently for big bang nucleosynthesis, the universe would have an overabundance of helium, which in turn would have affected the formation of the heavier elements involved in the evolution of stars and ultimately life. Thus, the balance between the universal cooling rate and the neutron lifetime was quite critical for the creation of the elements that make up our planet and everything on it.

From astronomical data we can measure the cosmic ratio of helium to hydrogen, as well as the amounts of deuterium and other light elements that exist throughout the universe. We would like to see if these measurements agree with the numbers predicted by big bang theory. The theoretical prediction, however, depends on the precise value of the neutron lifetime. Without a reliable value for it, our ability to make this comparison is limited. Once the neutron lifetime is known more precisely, we can compare the observed ratio from astrophysical experiments with the predicted value from theory. If they agree, we gain further confidence in our standard big bang scenario for how the universe evolved. Of course, if they disagree, this model might have to be altered. For instance, certain discrepancies might indicate the existence of new exotic particles in the universe such as an extra type of neutrino, which could have interfered in the process of nucleosynthesis.

One way to resolve the difference between the beam and bottle results is to conduct more experiments using methods of comparable accuracy that are not prone to the same, potentially confounding systematic errors. In addition to continuing the beam and bottle projects, scientists in several other groups worldwide are working on alternative methods of measuring the neutron lifetime. A group at the Japan Proton Accelerator Research Complex (J-PARC) in Tokai is developing a new beam experiment that will detect the electrons rather than protons produced when neutrons decay. In another very exciting development, groups at ILL, the Petersburg Nuclear Physics Institute in Russia, Los Alamos National Laboratory, the Technical University of Munich and the Johannes Gutenberg University Mainz in Germany plan to use neutron bottles that confine ultracold neutrons with magnetic fields rather than material walls. This is possible because the neutron, though electrically neutral, behaves as though it is a small magnet. The number of neutrons accidentally lost through the sides of such bottles should be quite different from that of previous measurements and thus should produce quite different systematic uncertainties. We fervently hope that, together, continuing bottle and beam experiments and this next generation of measurements will finally solve the neutron lifetime puzzle. SA

MORE TO EXPLORE

Measurement of the Neutron Lifetime Using a Gravitational Trap and a Low-Temperature Fomblin Coating. A. Serebrov et al. in *Physics Letters B*, Vol. 605, Nos. 1–2, pages 72–78; January 6, 2005.

The Neutron Lifetime. Fred E. Wietfeldt and Geoffrey L. Greene in *Reviews of Modern Physics*, Vol. 83, No. 4, Article No. 1173; October–December 2011.

Improved Determination of the Neutron Lifetime. A. T. Yue et al. in *Physical Review Letters*, Vol. 111, No. 22, Article No. 222501; November 27, 2013.

FROM OUR ARCHIVES

Ultracold Neutrons. R. Golub, W. Mampe, J. M. Pendlebury and P. Ageron; June 1979.

The Proton Radius Problem. Jan C. Bernauer and Randolf Pohl; February 2014.

scientificamerican.com/magazine/sa

© 2016 Scientific American