

Frequency-dependent squeezing makes LIGO even more sensitive

Researchers at the gravitational-wave observatory were already using nonclassical states of light to boost their measurement precision. Now they've unveiled a still-subtler trick.

It's been called the world's most precise machine, and it's one of the most ambitious. From an audacious dream in the 1960s to the securing of funding in 1992 to the first direct detection of gravitational waves on 14 September 2015, the Laser Interferometer Gravitational-Wave Observatory (LIGO) is the product of decades of theoretical and experimental work, much of it carried out with no guarantee of any reward.

But LIGO's work is never done. After the 2015 upgrade to Advanced LIGO, the observatory could detect signals a fraction the size of a proton, but that was still only barely sensitive enough to observe gravitational waves from one of the most violent events in the universe—the merging of two black holes some 30 times the mass of the Sun. (See *Physics Today*, April 2016, page 14, and December 2017, page 16.) To fully realize the potential of gravitational-wave astronomy, LIGO researchers have been working tirelessly to reduce noise and boost the observatory's sensitivity even more. They quickly ran up against the limits of quantum mechanics itself.

The Heisenberg uncertainty principle can drive a hard bargain, but it's willing to make a deal: It allows lowering the quantum uncertainty in one quantity at the expense of increasing it in another. Since LIGO's third observing run began in 2019, the observatory has been taking advantage of that offer by using so-called squeezed states of light. But it soon ran into another trade-off. Beyond a certain point, squeezing light increases the sensitivity to high-frequency gravitational



FIGURE 1. AT LIGO'S TWIN DETECTORS—this one in Livingston, Louisiana, and the other in Hanford, Washington—powerful laser light circulates in perpendicular 4-km-long arms. Through interferometry of the light from the two arms, researchers can detect changes in the arm lengths of less than 10^{-18} m. (Photo courtesy of Caltech/MIT/LIGO Lab.)

waves, but it makes the noise for low-frequency signals worse.

Now, for their fourth observing run, which began in May, LIGO researchers have broken through even that apparent barrier.¹ Using a technique called frequency-dependent squeezing, they've achieved the best of both worlds: reducing the noise across a broad range of frequencies while not increasing it anywhere. Thanks to the effort, LIGO can now detect gravitational-wave events over a larger volume of the universe than ever before—and moreover, a fundamental obstacle to pushing LIGO's quantum noise even lower has at last been removed.

Squeezed vacuum

LIGO, as the name suggests, works through laser interferometry. Light beams travel out and back along the two long arms of an L-shaped interferometer, and they recombine at their source. The system is tuned so that usually the beams interfere destructively: No gravitational

wave means that (almost) no light is detected. When a gravitational wave does pass through, it alternately stretches each arm while compressing the other. The length changes disrupt the interference and create an optical signal.

The stretches and compressions are tiny. Even the powerful gravitational wave from a black hole merger, by the time it gets to Earth, creates fractional length changes on the order of just 10^{-21} . To have any hope of seeing anything at all, LIGO researchers take every opportunity to boost the signal and suppress noise. The heavy mirrors that reflect the light are hung from sophisticated pendulums to protect them from vibrational noise. The facility uses not one interferometer but two—and a growing network of partner facilities around the world—to bolster the case that any wave they simultaneously detect is not a fluke. And the interferometer arms are 4 km long, as shown in figure 1, and the circulating laser power is in the hundreds of kilowatts, so even a small




FIGURE 2. SQUEEZING LIGHT shifts the quantum uncertainty from evenly spread along the waveform (brown) to concentrated in just one part (purple). Electromagnetic vacuum states, too, can be either unsqueezed (green) or squeezed (blue). (Adapted from R. Schnabel, *Phys. Rep.* **684**, 1, 2017.)

fractional length change can leak a detectable amount of light out of the interferometer.

But measurements of a light wave's intensity, like those of any other physical quantity, are subject to quantum uncertainty—and that's true even when there's no light wave present. Even in the electromagnetic vacuum state, there's always a chance that some photons will appear. There's no way for an interferometer output to ever truly be zero, and the quantum background can easily mask the feeble signal of a gravitational wave.

So what is squeezed light, and how can it help? Left to its own devices, quantum uncertainty tends to spread out uniformly along a waveform, as shown in brown in figure 2, but that's not the only option. With nonlinear optics, you can squeeze the uncertainty out of one part of the waveform and concentrate it in another. For example, the wave shown in purple has reduced uncertainty in its amplitude and increased uncertainty in its phase. If you're looking to measure the amplitude, and you don't care about the phase, the squeezed state offers a big improvement.

Roughly speaking, the uncertainty principle treats a wave's amplitude and phase the same way it treats a particle's position and momentum: The product of the two uncertainties is constrained, but either one can be reduced at the expense of the other. For an interferometer like LIGO's, phase is the more important quantity. It's the timing of the light waves from

the two arms that determines whether they interfere destructively or not.

The idea of using squeezed light for gravitational-wave detection was laid out by Carlton Caves in 1981—decades before LIGO was built and years before anyone had even observed squeezed light in a lab.² Caves anticipated that the way to do it was to squeeze not the state of the laser light itself but rather the state of the electromagnetic vacuum that enters the interferometer where the signal light comes out. Figure 2 shows how vacuum states can be either unsqueezed (green) or squeezed (blue). Although the vacuum lacks either amplitude or phase, those terms can be defined according to its interaction with the interferometer light.

In 2019 LIGO implemented Caves's scheme for using a phase-squeezed vacuum to substantially reduce quantum noise.³ But there was a fly in the ointment: The increased amplitude uncertainty, which transfers to the amplitude of the light inside of the interferometer, is not harmless. When light hits the mirrors at the ends of the interferometer arms, it exerts radiation pressure on them—and because the mirrors are dangling from pendulums, fluctuations in the radiation pressure can set them swinging. The mirrors are heavy and the fluctuations are small, so they don't swing very much. But the signals LIGO seeks to detect are so extraordinarily small that it doesn't take much to obscure them.

Only the low-frequency signals are obscured: The weighty mirrors can't swing fast enough to make any difference in the detection of gravitational waves above about 300 Hz. Low-frequency signals, however, are important. The events LIGO detects—merging pairs of black holes and neutron stars—generate gravitational waves as the massive objects circle one another faster and faster for a few tenths of a second before colliding. If the observatory were to give up on detecting signals until the orbital speed had ramped up to 300 cycles per second, it wouldn't detect much at all.

Phase delay

To avoid the detrimental effect on low-frequency signals, LIGO's 2019 implementation of squeezed light limited its squeezing to three decibels, or about a factor of 2. But the researchers were already working on doing better—and once again, their work built on theoretical foundations that had been laid decades ago.

In a 2001 paper, H. Jeff Kimble and colleagues presented the idea of enhancing gravitational-wave detection by squeezing light differently at different frequencies.⁴ In their analysis, "frequency" refers not to the frequency of the laser light in the interferometer (which is perfectly monochromatic) but to the frequency of the gravitational waves it's trying to detect. The state being squeezed, after all, is the electromagnetic vacuum, which doesn't have an inherent frequency itself but can be thought of as having components of all frequencies.

"Squeezing at every frequency is independent," says Lee McCuller, a LIGO scientist at Caltech, "and it just kind of works out that the way we usually make a squeezed vacuum squeezes the same at every frequency." In LIGO's case, every frequency is phase squeezed. Kimble and colleagues' idea was to instead create a state that varies from phase squeezed at the highest frequencies to amplitude squeezed at the lowest.

Luckily, a phase-squeezed vacuum and an amplitude-squeezed vacuum look exactly the same, and one can be transformed into the other simply by delaying it by a quarter of a wave cycle. So creating frequency-dependent squeezing is just a matter of introducing a frequency-dependent delay—and that can be done by bouncing the




FIGURE 3. NOISE AT LIGO'S LIVINGSTON DETECTOR (black) consists of classical noise (gray) and quantum noise (pink). Using squeezed light can substantially reduce the quantum noise at high frequencies—but the initial approach, frequency-independent squeezing (green), also increases the noise at low frequencies. LIGO researchers have now implemented frequency-dependent squeezing (purple), which reduces the noise across many frequencies without increasing it anywhere. (Adapted from ref. 1.)

phase-squeezed vacuum off a long optical cavity.

“Think of it as how, when you yell into a cave, if your voice is resonant with the cave you hear an echo,” says Victoria Xu, a postdoc in MIT’s LIGO lab. “But if it’s not resonant, you hear nothing.” Similarly, the low-frequency components of the squeezed vacuum enter the cavity and ricochet around for a while before exiting, while the high frequencies ignore the cavity and are reflected straight back.

Kimble and colleagues had worked out the theory, but implementing frequency-dependent squeezing to LIGO’s stringent standards posed additional challenges. “The tricky part is to think about what you’re asking for,” says Xu. The low frequencies that LIGO seeks to detect—from tens to a few hundreds of hertz—are extremely low by electromagnetic standards. Creating the requisite phase delay of 3 milliseconds means building a cavity 300 meters long and holding the light inside for a few thousand round trips.

And it all had to be done without losing any photons. “Squeezed light is extremely sensitive to loss,” says Xu. As a nonclassical state of light, it can be thought

of as made up of entangled pairs of photons. “If you lose one photon from a correlated pair, you have nothing,” she says.

By 2020, LIGO researchers had tested frequency-dependent squeezing with laboratory-scale experiments, including one led by McCuller that used a 16 m cavity.⁵ Satisfied that they could make it work, they decided to take the plunge and push to implement the technology for LIGO’s fourth observing run.

“It was amazing that it worked so fast,” says McCuller. “The 16-meter experiment took us four years—but we had just a few researchers and postdocs working on it. The real deal had to come together much faster.”

“Three hundred meters is as far as I can walk in five minutes,” says Xu. “To house a cavity that big, we had to build whole new buildings and new clean rooms. It’s nothing that nobody’s ever done before, but for LIGO it had to be done on a massive scale. And this is the kind of thing that LIGO is really good at.”

A new baseline

Figure 3 shows the resulting noise reduction at LIGO’s Livingston detector.

(Data for the Hanford site are similar.) At high frequencies, the frequency-dependent-squeezing noise (purple) is six decibels lower than what would have been achieved with no squeezing (black), whereas at low frequencies it’s unchanged. And the frequency-independent-squeezing noise (green) matches the purple curve at high frequencies, but at low frequencies it’s much higher.

The black and green curves don’t represent the noise that LIGO achieved during its third observing run (or at any other time), but rather they show what it would have achieved in its fourth run without frequency-dependent squeezing. “We compare the noise not to the previous run but to the best we can do now,” says Lisa Barsotti, a senior research scientist at MIT’s LIGO lab, “and we never make only one improvement from run to run. There’s always a constant effort to keep reducing the classical, technical noise too.”

With that caveat in mind, the researchers estimate that the difference between the black and purple curves means that LIGO can detect events 15–18% farther away—or over a 50–65%

larger volume of the universe—than it otherwise could. But that improvement is only the beginning.

“This is now the baseline for any future upgrade,” says Barsotti. Before, LIGO had to deliberately throttle its light-squeezing efforts to avoid compromising its low-frequency sensitivity, but that’s no longer the case. “The next step

is to improve how much squeezing we can see,” Barsotti explains. “We can squeeze the light as much as we want, and we’re only limited by how well we can get it into the interferometer. This is going to be important not only for LIGO but for all future ground-based gravitational-wave detectors.”

Johanna Miller

References

1. D. Ganapathy et al., *Phys. Rev. X* **13**, 041021 (2023).
2. C. M. Caves, *Phys. Rev. D* **23**, 1693 (1981).
3. M. Tse et al., *Phys. Rev. Lett.* **123**, 231107 (2019).
4. H. J. Kimble et al., *Phys. Rev. D* **65**, 022002 (2001).
5. L. McCuller et al., *Phys. Rev. Lett.* **124**, 171102 (2020).

Uncovering the molten mantle of Mars

A delay in seismic-wave arrival times reveals the presence of an additional layer in the planet’s interior.

Mars is only the second celestial body, after the Moon, that humans have sent seismometers to. A desire to

understand its interior, which holds clues to Mars’s origin and evolution, motivated NASA to develop *InSight* (*Interior Exploration Using Seismic Investigations, Geodesy and Heat Transport*). Before that mission, astronomers’ knowledge of Mars’s interior came primarily from models of solar-system formation, Mar-

tian meteorites, and geophysical observations from satellites orbiting the planet. Those data provided only an incomplete glimpse of Mars’s interior.

The seismometer on the *InSight* lander—known as SEIS, or the Seismic Experiment for Interior Structure—collected data from 2019 to 2022 by listening




FIGURE 1. MARS'S INTERIOR STRUCTURE can be inferred from seismic waves that travel through the planet. Waves go out in all directions from the seismic events; shown here are examples of wave paths that support the existence of a deep liquid silicate mantle layer just above the core. S waves from events near the *InSight* lander reflect off the solid–molten boundary and can be used to determine the inner radius of the solid mantle. An event farther away allows the seismometer to detect P waves that penetrate to deep layers of the planet, revealing the presence of a molten silicate mantle layer. (Adapted by Jason Keisling from ref. 1.)