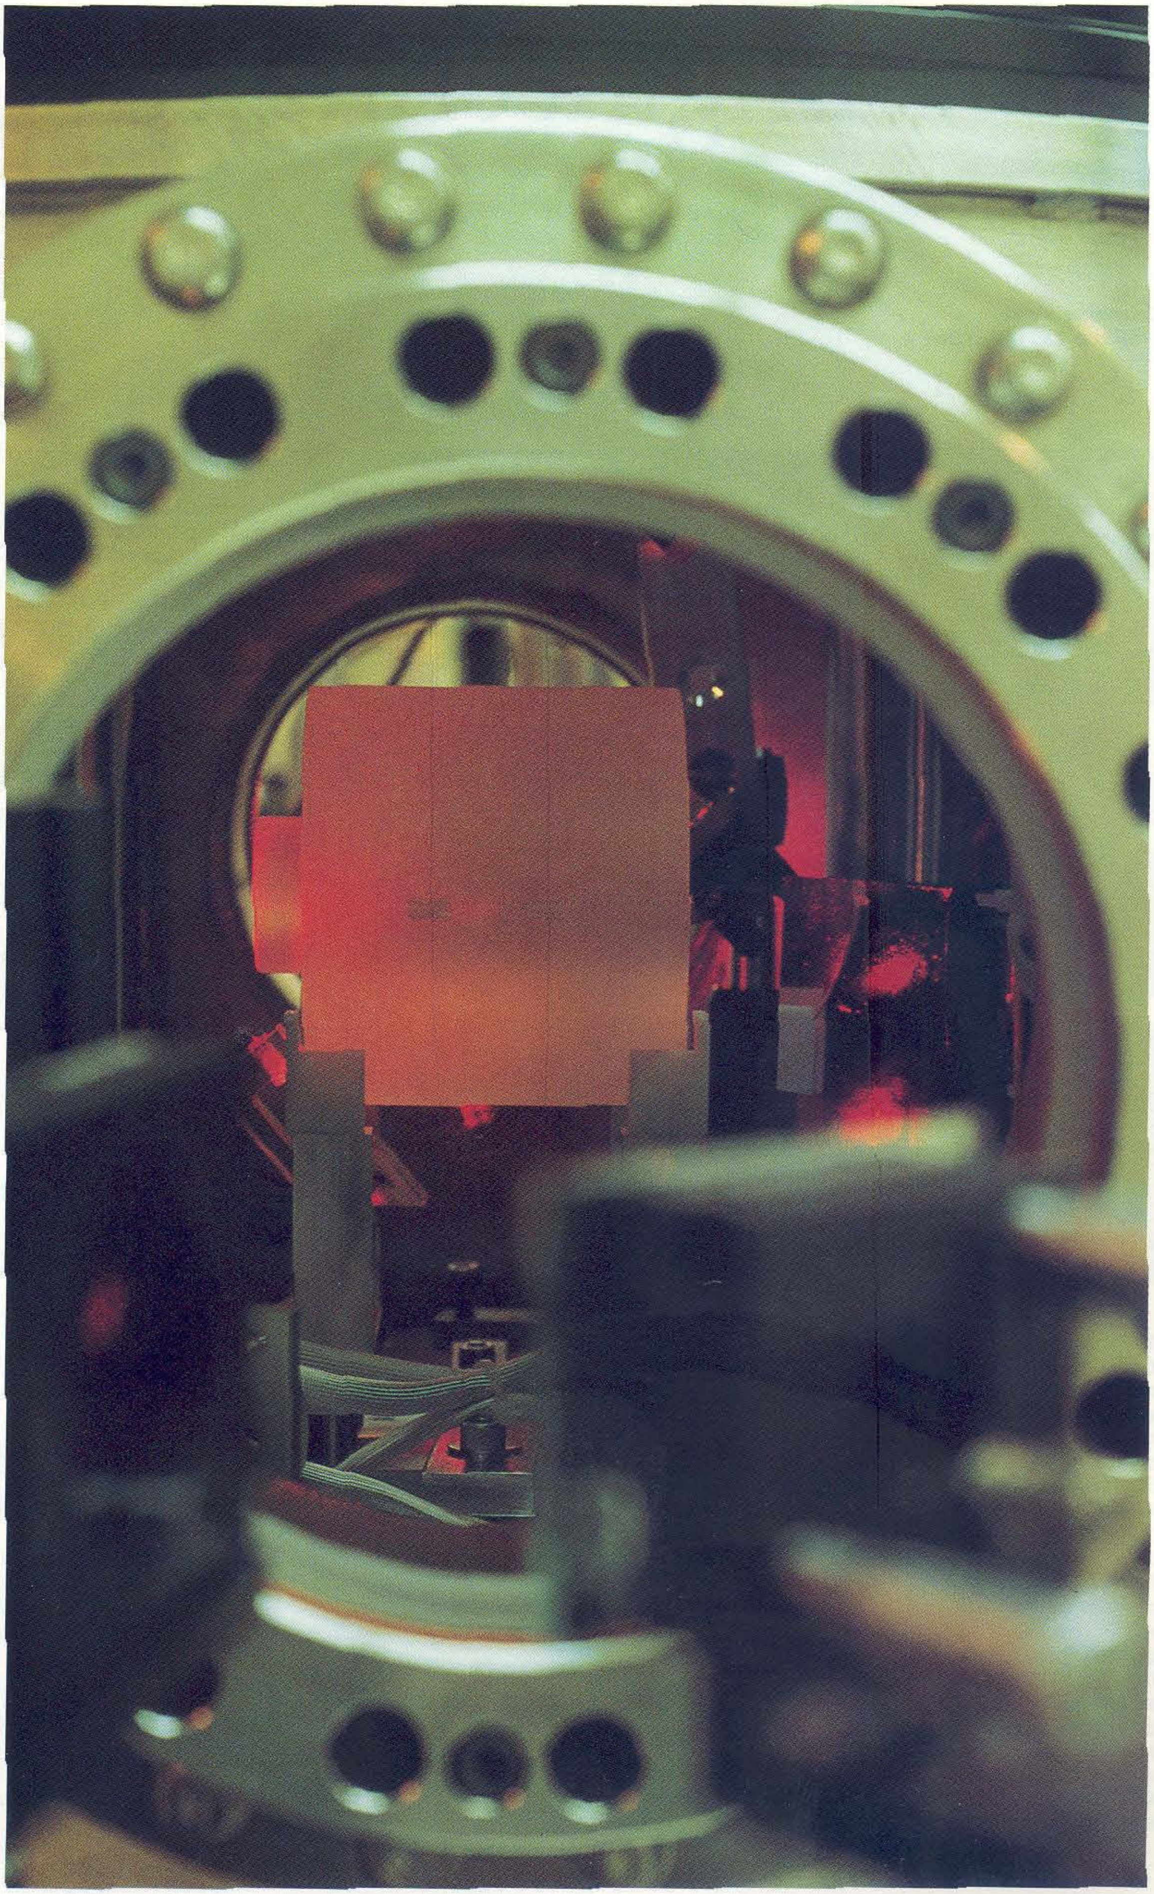

Lasers, Mirrors, and Gravitational Waves

by Frederick J. Raab

These waves, which bear tidings of cataclysms like the deaths of stars and the collisions of black holes, have so far escaped our view.

A three-pound test mass of fused silica—glass—hangs in a vacuum chamber from wires three-thousandths of an inch thick. The wires can be seen as fine vertical lines that divide the mass into thirds. The mass is a four-inch-diameter horizontal cylinder; the mirror is visible on its left side. The ruby suffusion bathing the mass comes from a secondary laser system, part of a servomechanism that keeps the mass in alignment. The mass hangs above, but does not touch, the legs visible in front of it. These legs prevent it from swinging wildly in the event of an earthquake or other sudden jolt that could damage the system.

We know of four fundamental forces in the universe. Two of them, called the "strong" and the "weak" interactions, are the stuff of particle physics, and are very short-range forces. There are also two long-range forces, electromagnetism and gravity. The electromagnetic force can propagate as waves in an electromagnetic field, which we perceive as photons of light or as radio waves. Einstein's theory of general relativity predicts analogous gravitational waves in the gravitational field. In the language of relativity, the fabric of space and time will be distorted in the vicinity of a massive object. Picture Earth as floating on the surface of a pond that represents space-time, and putting a little dimple—its gravitational field—on the water's surface. The dimple, in turn, "attracts" nearby objects whose motion is affected by the dimple's shape. If Earth made violent motions, it would create ripples that would propagate to the far edges of the pond. In this metaphor, the ripples are gravitational waves.

Both electromagnetic and gravitational waves carry information about the physical processes that created them. Astronomy has principally been concerned with interpreting the electromagnetic spectrum, which is imprinted by things that happen in the atmospheres of stars, in gas clouds, and in other places where photons are created or scattered. But gravitational waves are created by the bulk motions of matter. These waves, which bear tidings of cataclysms like the deaths of stars and the collisions of black holes, have so far escaped our view. Astrophysicists and other people interested in such violent events would like to be able to detect these waves and read the stories written on them.

To this end, we hope to start building LIGO—the Laser Interferometer Gravitational Wave Observatory—in the next few years. LIGO funding has been proposed in President Bush's budget for 1992, which is pending before Congress. LIGO is a Caltech–MIT joint project, directed by Rochus Vogt, Avery Distinguished Service Professor and professor of physics; and including Stan Whitcomb as deputy director; Ronald Drever, professor of physics; Kip Thorne, Kenan Professor and professor of theoretical physics; me; and Rainer Weiss, professor of physics at MIT, plus an excellent technical staff of engineers, physicists, and students at both institutions. A LIGO facility—there will be two of them—consists of a central building containing lasers and detectors, plus two four-foot-diameter vacuum pipes that stretch approximately two and a half miles (four kilometers) from the central building to form an "L." These pipes carry laser beams that sense the separation between test masses—20-pound masses at first, to be replaced by one-ton masses later on—hanging in the central building and in buildings at the far end of each arm. We hope that by monitoring their separation very carefully, we will see them move infinitesimally as gravitational waves pass.

There would be two such detectors, operating in concert, on opposite sides of the United States. We need at least two well-separated stations to know when we have seen a gravitational wave, as opposed to some local disturbance. A real wave would trigger both detectors within one-sixtieth of a second of each other. We can use this differ-

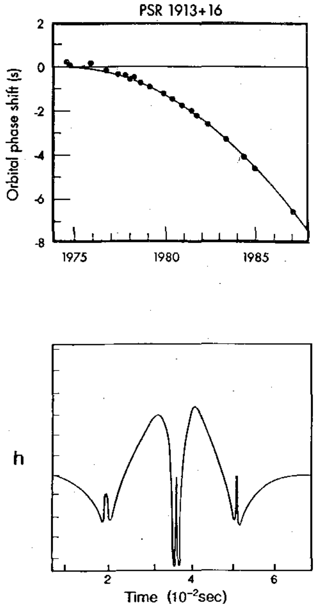

Top: PSR 1913+16's gradually decaying orbit has offered the first indirect evidence for the existence of gravitational waves. [Taylor and Weisberg, 1989]
Bottom: A collapsing supernova's gravitational wave might look like this ("h" is the fractional change in separation between test masses the wave causes). The complex shape arises because different parts of the core collapse and rebound sequentially. The collapse begins along the star's rotational axis, where material isn't supported by centrifugal force. [Saenz and Shapiro, 1978]

ence in arrival times to determine where in the sky the source is. Unfortunately, having two detectors merely specifies the source as being within a ring of sky. Three detectors would cover most of the sky unambiguously, and four would cover the whole sky. The Europeans, Australians, and Japanese are considering building similar detectors, so eventually there should be an international network, of which LIGO will be a part.

Sanduleak was a very interesting star, although no one knew it then, because it had died 160,000 years ago. That night, the signal from its death reached Earth.

You may be asking yourself, "It's pretty easy to believe in electromagnetic waves, what with TV and all, but should I buy any of this stuff about gravitational waves? Do they really exist?" Let me show you a system whose gravitational radiation can be calculated rather accurately. . . a very simple system consisting of two masses—two stars orbiting each other. If Einstein is right, gravitational radiation would carry energy away from the system, so that the orbits would gradually decay and the stars would spiral in toward each other. Around 1975, science got lucky—although luck is always due to very good work—and Joseph Taylor and his colleagues discovered a very interesting object, called PSR 1913+16, that consists of two neutron stars orbiting each other. (Neutron stars are several-mile-diameter balls of essentially pure neutrons. They're the compressed remnants of stars that have long since died.) One of the pair is a pulsar that sends out a narrow beam of radio waves. As the pulsar spins on its axis, the beam sweeps across Earth like a lighthouse beacon, and one can measure the interval between flashes with exquisite precision. From these data, one can derive the orbit's phase shift. That is, one can measure the interval between successive times when the two stars have a given alignment as seen from Earth, and see if there is any drift in that interval. In the topmost graph above, the horizontal line at zero seconds is what would be observed if there were no energy loss from the orbiting pair, and certainly the data points don't agree with it. The curved line is what we would expect if the energy in the orbit was being lost—as general relativity predicts—by the radiation of gravitational waves. In fact, Taylor and his colleagues have measured a drift of about eight seconds over the years. You can see how well the data fit that curve, so it's a fairly safe bet that we've already seen the result of gravitational radiation in the orbital decay of this particular pulsar.

What kinds of events might one study using gravitational waves? One promising event is the supernova—the death throes of a large star. For instance, on February 23, 1987, astronomers Robert McNaught and Ian Shelton were separately photographing a region of sky that happened to contain a star called Sanduleak. Sanduleak was a very interesting star, although no one knew it then, because it had died 160,000 years ago. That night, the signal from its death reached Earth and was recorded on film—the famous Supernova 1987a—the brightest one seen from Earth in 300–400 years.

Now, normally a star burns nuclear matter by fusion—the star's tremendous gravity squeezes atoms together until they fuse—gradually turning the star's hydrogen and helium into the elements that we know and love and are made of. But at some point the fuel runs out, and the star's

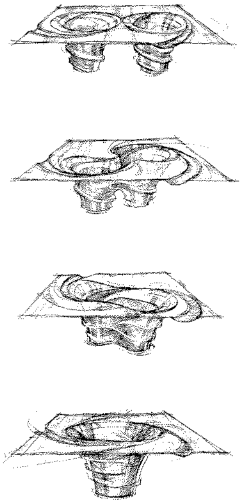

Above: Two black holes orbiting each other eventually coalesce into one.

Below: In their final moments as separate entities, just before each one's immense gravity starts to disrupt the other, the merging black holes should produce this gravitational-wave signature.

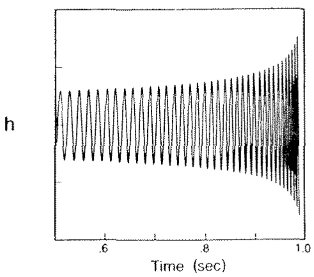

nuclear engine can no longer provide the outward pressure needed to keep the star from collapsing under its own enormous weight. A star the size of our sun will just fizzle out. Much more massive stars undergo a gravitational collapse that is one of the truly spectacular events in the universe. The star's iron core, which is about the mass of our sun, collapses to a few miles in diameter in less than a second. The collapse releases subatomic particles called neutrinos. These interact fairly weakly with matter, but enough so that it takes them a few seconds to boil out through the star, and in the process they deposit enough energy in the star's outer layers to blow them apart. The supernova that appears on photographic plates is a record of the photons released from this explosion. This visible display begins to appear hours later, when the photons finally escape. The material gets dispersed through the heavens, and it later forms planets and people. This is where the stuff that we're all made of comes from.

One can watch the supernova's light show with all the tools for detecting electromagnetic radiation—optical telescopes, radio telescopes, gamma-ray detectors, and so on. (In the case of Supernova 1987a, the neutrinos were detected as well.) One then practices forensic medicine, examining the corpse of the star for clues about the manner of its demise.

But if a gravitational-wave detector catches a supernova in the act, we can glimpse the inner workings of a collapsing star. The waves released as the core collapses and rebounds pass right through the star's outer layers as if they weren't there. Even after the collapse's waves are long gone, the object left behind may still emit gravitational waves. This object may be a neutron star or a black hole. It can remain hidden from our electromagnetic view for years by the veil of exploded stellar material, or it may fail to emit electromagnetic waves.

Another thing we can do with a supernova is determine whether gravitational waves travel at the speed of light, or merely very close to it. This question is related to whether the graviton—the hypothetical object that "carries" gravitation the way photons carry light—has mass or not. The fact that gravity is a long-range force is interpreted to mean that the graviton is massless, like the photon. But we don't know absolutely for sure that the graviton has exactly zero mass. It could have a very small mass, which would mean that the range of gravity isn't infinite after all, but merely very long. We could answer this question by racing gravitons against something whose mass is known, like photons, because the maximum propagation velocity of an object is related to its mass.

This has already been done with the neutrinos and photons emitted from Supernova 1987a, which had to travel 160,000 years to get here because Sanduleak was 160,000 light-years from Earth. It's like watching sprinters. Carl Lewis will always beat me in a foot race, even if I jump the gun a bit, as long as my head start is small compared to the length of the racecourse. Similarly, photons will always beat massive neutrinos if the race is long enough. By timing when the photons and neutrinos from a supernova arrive at Earth, we get an upper limit on the speed and mass of the neutrino relative to the photon, even though the neutrinos get a head start. This has confirmed that the neutrino's mass is less than about 20 electron volts. Once gravitational waves have been detected, we can use a similar technique to actually measure their speed and thus test the theoretical underpinnings of relativity.

We can also look for binary systems made up of two black holes. (Black holes are regions of gravity so strong that even light can't escape.) We hope to see the waves emitted when two black holes capture each other. Again, their orbits gradually decay through the emission of gravitational radiation until the black holes become so close that they start to disrupt each other tidally. We can predict what the waveforms will look like before the tidal distortions start. Life becomes much more complicated for theorists after that, but, with the help of supercomputers, they will be able to calculate the details of the death spiral. These waveforms could be used as a signature—their detection would be proof that black holes actually exist.

We could look for all kinds of waveforms, in fact. Deciphering a waveform's details should tell us what created it. A supernova collapse and a black-hole binary coalescence should look very different. There are other things that we know should make gravitational waves, and there are probably just as many other things we can't even dream of yet. When detecting these waves becomes routine, we will have a powerful new tool for astronomy, comparable, perhaps, to the advent of radio astronomy, which made possible the discovery of objects like pulsars and neutron stars in the first place.

Clearly there's a lot to be gained from making the routine detection of gravitational waves a reality. The hard part is building detectors with the requisite sensitivity. The basic LIGO detector is a system of suspended masses that are free to move horizontally; when a gravitational wave comes by, it perturbs the distance between them

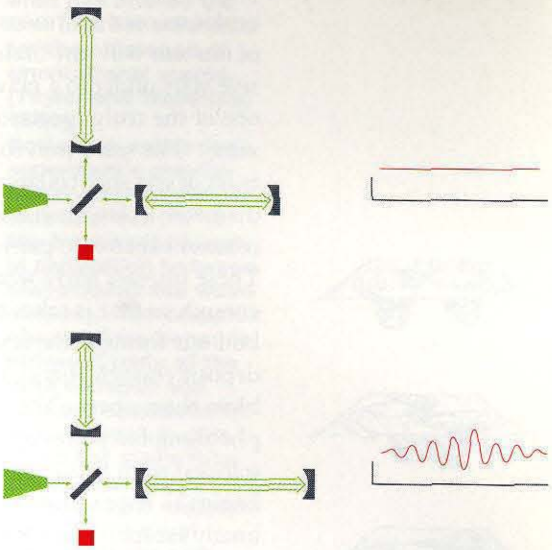

Right: How an interferometer works. When no gravitational waves are passing by (top), the distances between the two sets of mirrors are the same, the reflected beams cancel each other, and no light reaches the photodetector. When a wave comes by (bottom), it nudges one set of mirrors slightly apart and the other set slightly together. The beams no longer cancel, and the photodetector records a signal whose shape follows the wave. Below: The 40-meter prototype. The optical setup in the foreground filters and stabilizes the laser beam, which enters the interferometer via the right-hand end of the horizontal pipe. Of the three mesh cages, the middle one houses the beam splitter and the other two house test masses. Note the 6' 3"-Raab standing at the arm's far end.

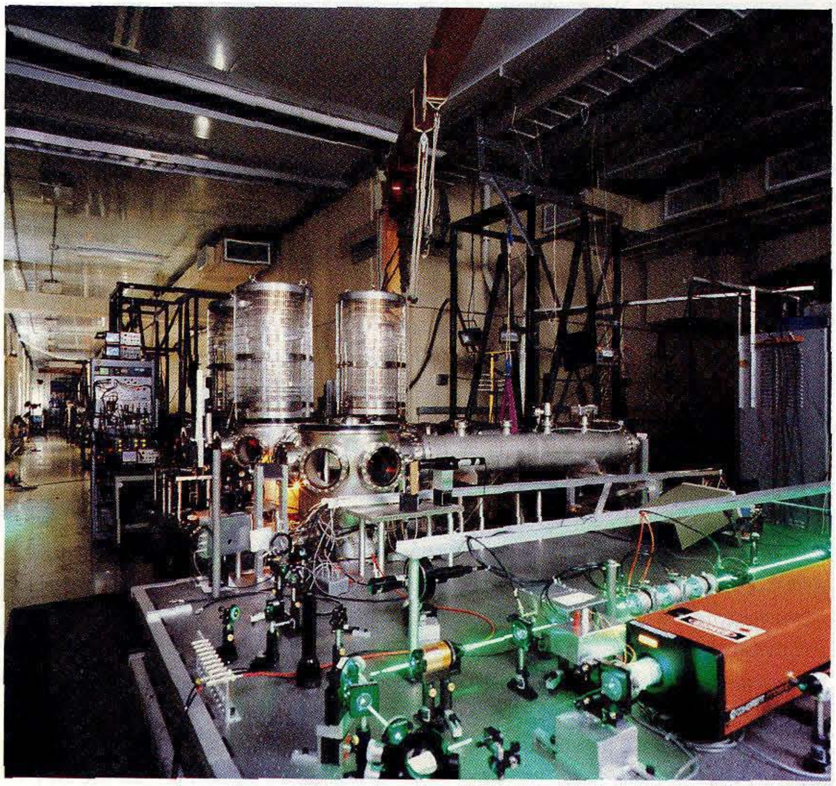

as measured by a laser beam. Imagine sending the beam out from one mass, reflecting it off the other one, and measuring how long the trip took. Since we know the speed of light, we can figure out the distance. LIGO has two pairs of masses in perpendicular directions, forming the two arms of an "L." A passing gravitational wave will change the distance between the masses by some fraction of that distance. Think of the masses as being glued to a rubber sheet. If the sheet stretches in one direction, moving the masses apart, it will shrink in the other, moving the masses together. Thus we need only compare the lengths of the arms to detect the waves.

This comparison can be done by actually splitting the laser beam, sending half of it down each arm of the L, and then recombining the reflected beams by passing them through a device called a beam splitter. This setup makes use of something called interference—hence the word *interferometer* in the name LIGO. Light waves are just oscillations in the amplitude of an electric field, and the beam splitter adds the two recombining fields together. Normally, the crests of the field from one arm line up with the troughs of the field from the other arm—or if they don't, the masses can be moved slightly so that they do—and the beam splitter adds a plus field and a minus field to get zero. This is called destructive interference. No light reaches the photodetector. But if a gravitational wave makes one arm a little longer and the other a little shorter, then crests start to line up with crests, and troughs with troughs. Now when the beam splitter adds the waves, it gets a field with some amplitude. This

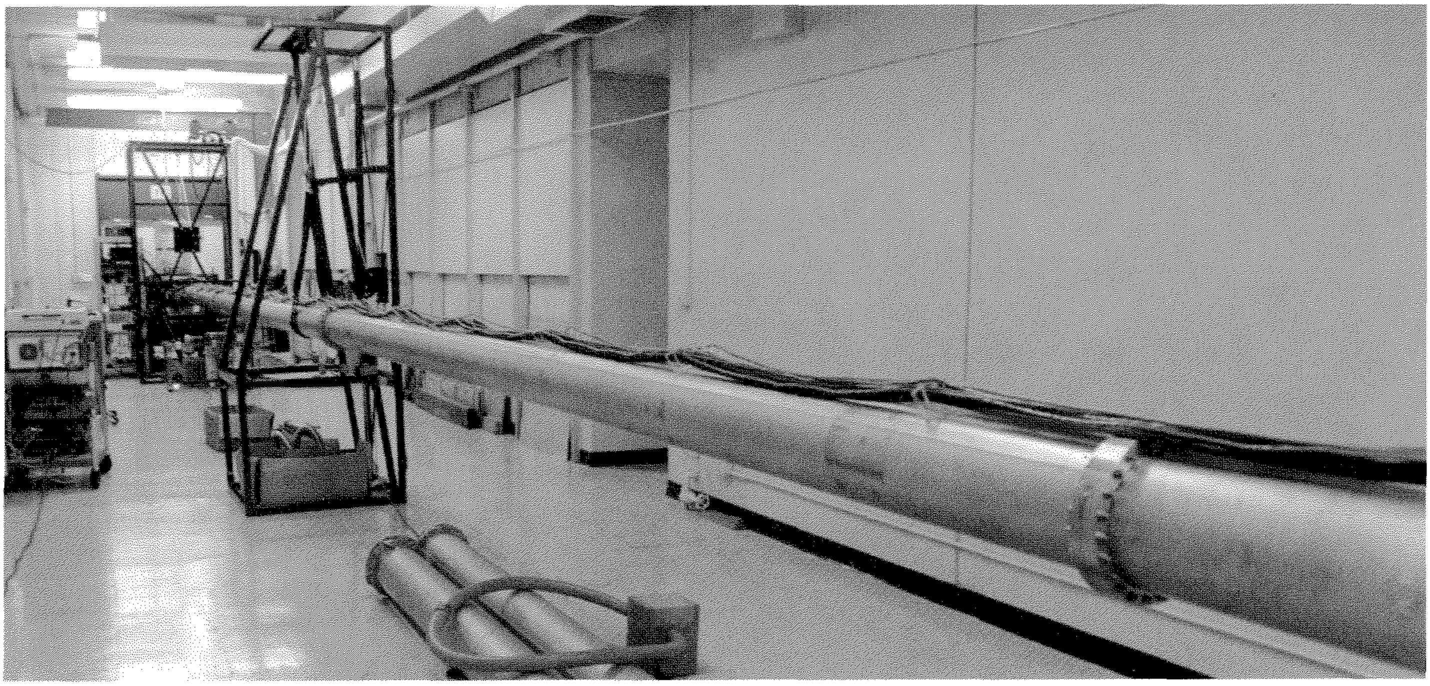

The view down the hall: looking along one of the 40-meter interferometer's arms toward the test-mass chamber at its far end. The photo was shot from about the arm's midpoint.

field hits the photodetector, which typically detects the square of the field, and we get a curve that follows the shape of the gravitational wave.

Our group has built a working prototype gravitational-wave detector. Its arms are 40 meters long, and it's housed in a prefab building wrapped around two sides of Caltech's Central Engineering Services building. In fact, we've prototyped almost all of the techniques I'll describe, in either this 40-meter system or in some other apparatus. When Ford builds a prototype vehicle, they don't do it because the CEO wants to drive to the store for a loaf of bread; they do it to find out how to build a better car. Similarly, our prototypes are not likely to detect gravitational waves; we use them to develop and test the technology that will eventually be incorporated into LIGO's detectors. This technology works on paper, but we want to see it work in the real world, and learn how to operate it. Our 40-meter prototype has masses suspended very much as they would be in LIGO. A number of optical schemes—I'll describe them presently—have also been prototyped. Since we don't need the suspended masses to test purely optical properties—just the detector's lasers, mirrors, and such—we build these kinds of prototypes on an ordinary optical bench.

I have one big problem. Remember the supernova, and the star's outer layers that the gravitational waves went right through? We're all made from that material, and so are any gravitational-wave detectors we might build. If the waves zipped through several times the mass of the sun with impunity, they sure won't move any test masses much. With LIGO as it is presently designed, and considering the sources we want to see, we expect fractional changes in the distance between the masses of about 3×10^{-22}. That is, the distance change between masses two and a half miles apart will be about one-thousandth of the diameter of an atomic nucleus. This could give you pause. It would give *me* pause, except that we're operating a prototype where we're close to detecting that small a change now. We've gotten down to 10^{-18}—about one-thirtieth of a nuclear diameter over a distance of 40 meters for an event that lasts for a few thousandths of a second. So if we make the distance between masses 100 times greater so that they would move 100 times farther in response to a passing wave (as planned for LIGO), and if we can improve the precision of our measurement of small displacements by a factor of 30 (which we think we can do), we can make this scheme work. It's a challenging program, but the technology to do it is within reach.

We would like to separate the masses as far as possible, because the masses' motion is proportional to their separation, but this isn't as easy as it sounds. With LIGO's four-kilometer arms, we're already running into problems with Earth's curvature, not to mention acquiring the real estate. Instead, we put mirrors at both ends of each arm to bounce the light back and forth many times, so that we get more signal for a given mass motion. The setup we use is called a Fabry–Perot cavity.

In each arm the front mirror—the mirror closest to the laser—is only partially reflective,

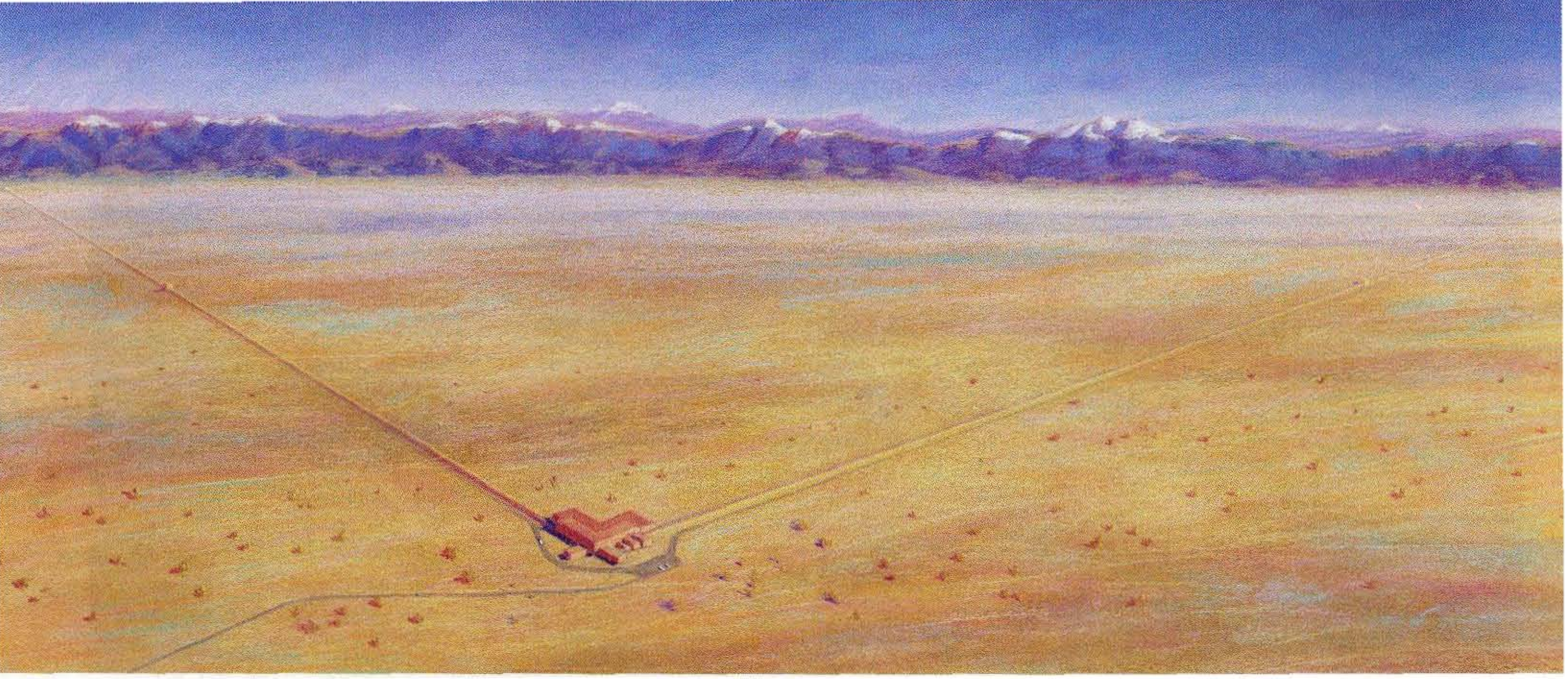

Above: How a real LIGO installation might look. The L-shaped building in the foreground houses the lasers, control equipment, offices, and such. The interferometer arms would be protected by concrete culverts.

to let the light into the arm. However, some of the light bounces back toward the laser. Let's call the electric field associated with the light reflected from the front mirror "F." The light that enters the interferometer arm bounces off the second mirror and rattles around between the mirrors. But with each round trip, a little light leaks out of the cavity through the front mirror. I'll call the leaking field "C." Once again I can invoke interference—in fact, I'm going to invoke it frequently now. If I set the front and back mirrors at the right distance from each other, the crests of wave C and the troughs of wave F line up—destructive interference occurs—and behold, there's less light reflected from the front mirror back to the laser. More light goes into the arm, where we want it, instead. (This is how the anti-reflective coating on your camera lens works. A coating on the front surface of the lens gives an F reflection that cancels the C reflection from the lens itself. That's why you don't see your face very well when you look into your camera.) We typically bounce the light back and forth about 3,000 times in our prototype. This feat requires extremely high-quality mirrors; in the longer LIGO fewer bounces will be required.

There are various so-called noise sources that, depending on how cleverly we deal with them, can determine the smallest gravitational wave we can measure. (See graph, opposite.) These generally fall into two classes: noise sources that prevent the masses from hanging motionless when no waves are present, and sources that affect our ability to detect very small motions of the masses; these latter sources are known as sensing noises.

Below: Schematic diagram of an interferometer made from two Fabry-Perot cavities. "F" is the front field, "C" is the cavity field, and "D" is the photodetector.

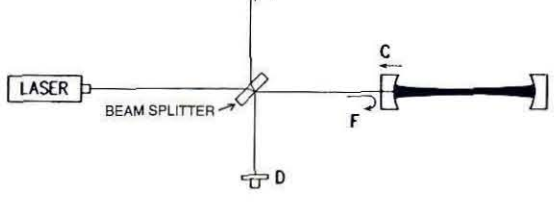

Right: Noise from many sources can hide a gravitational wave. The bars below the graph show the frequency ranges at which various types of sources are expected to emit gravitational waves. The vertical scale shows the smallest wave detectable in the presence of various noise sources. For a one-second-duration wave at a given frequency to be detectable, its amplitude at that frequency must fall on the highest line. (Longer-duration waves can be detected below that line as their signal accumulates.)

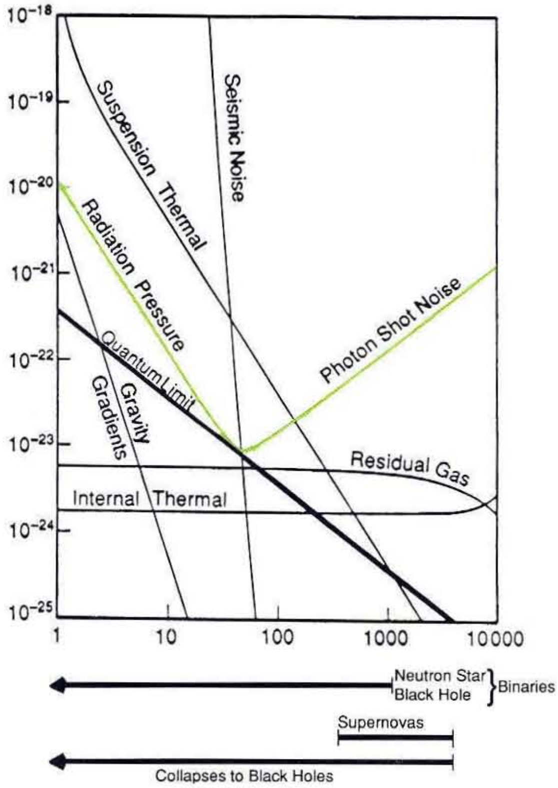

The evolution of a vibration-isolation stack. Clockwise from the top:
1. The first stacks installed in the 40-meter prototype had horseshoe-shaped slabs supported on erasers—a convenient source of uniform pieces of soft rubber.
2. A newer test setup uses pairs of bars stacked crosswise like the logs in a cabin. In the search for elastomers with better vacuum properties, the erasers have given way to silicon-rubber cubes.
3. The eraser and the cube.

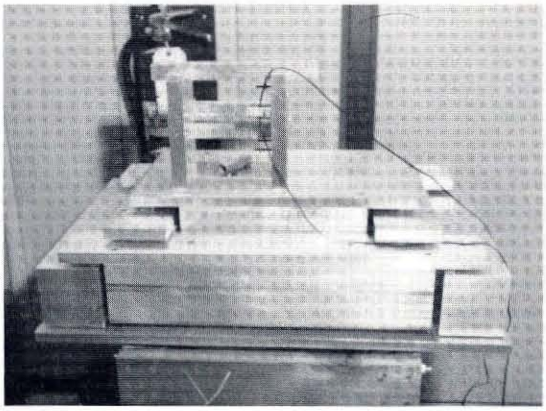

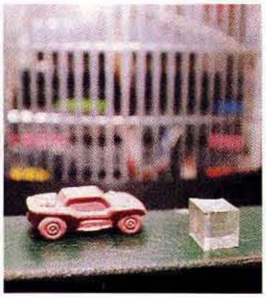

"Seismic noise," which includes the rumble of passing trucks as well as natural ground motions such as earthquakes, falls into the first class of noise sources. We can make the test masses very insensitive to seismic noise by the clever use of vibration-isolation systems. The basic principle is that a mechanical oscillator—a pendulum, or a mass supported by a spring—doesn't move much when its support is pushed by a force at a frequency higher than the oscillator's "fundamental resonance." (The fundamental resonance is the frequency at which the pendulum naturally swings or the mass naturally bounces.) Therefore, we hang each test mass like a pendulum, in slings of fine wire hanging from a frame that sits atop a stack of metal slabs. Each slab is separated from its fellows by springs of a resilient, rubbery elastomer not unlike the caulk around your bathtub. The springs isolate the upper slab and its test-mass cargo from the lower slab and its eventual connection to the outside world, much the way that your car's suspension isolates the chassis from the pounding of the tires on a bumpy road.

Since LIGO detectors will operate at room temperature, the mirrored surfaces of the test masses will move randomly in response to thermal excitation of the wire slings ("suspension thermal" on the graph) and internal vibrational modes of the test masses themselves ("internal thermal"). Seismic noise and thermal noise can be reduced by technological improvements such as the development of better vibration-isolation systems and higher-mechanical-quality materials. Unfortunately, gravity gradients caused by density fluctuations underground and in the atmosphere cause motions in the test masses that can't be shielded, and will limit the low-frequency sensitivity of Earth-based detectors.

A familiar example of a sensing noise is the apparent motion of distant, stationary objects on a hot day, due to density fluctuations in the intervening atmosphere. The space between LIGO's test masses must be evacuated to a pressure low enough that the density fluctuations in the residual gas left in the pipe ("residual gas noise") don't become confused with actual test-mass motion. Once again, this isn't as easy as it seems, because stainless steel absorbs hydrogen gas during the manufacturing process. This gas could leak out of the steel, and into our vacuum system, for years. Other projects—particle accelerators, for example—that use steel piping to hold a high-purity vacuum generally bake their components out at temperatures up to 1600 degrees Fahrenheit to expel the hydrogen. That's impractical for us—the electrical bill alone would be truly astronomical. So we've collaborated with industry to

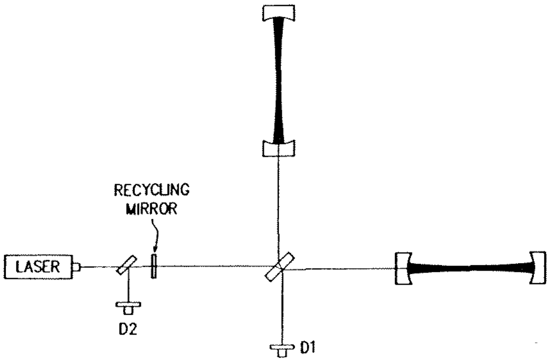

Right: Installing a recycling mirror between the laser and the beam splitter sends the light back to the interferometer's arms. D_1 is the photodetector. D_2 is a secondary detector used to adjust the recycling mirror's position for maximum destructive interference. Placing another recycling mirror between D_1 and the beam splitter would create the dual recycling interferometer described on page 11.
Below: The optical components in the foreground are part of the system that filters and stabilizes the laser light. In the background, undergrad Maggie Taylor (seated), LIGO scientist Seiji Kawamura (kneeling), and Raab confer.

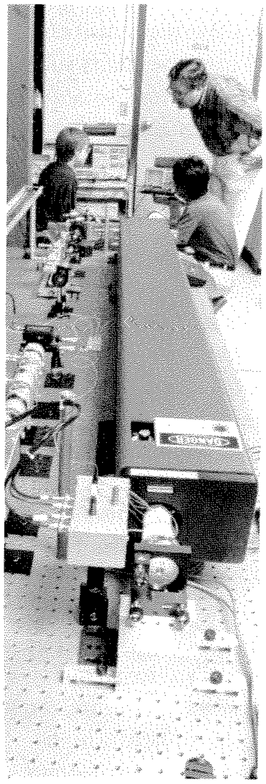

develop a special process to manufacture low-hydrogen stainless steel.

And finally, the green curve shows how the quantum nature of light—the fact that it's made of photons—affects sensitivity. "Photon shot noise" is very much like hearing rain on the roof. We hear a patter because the rain comes as droplets instead of a continuous flow of water. Similarly, with light we may be measuring a million photons per unit of time, on the average, but in any given interval, the million will be either in excess or shy of photons by about a thousand. The fluctuation during any interval generally depends on the square root of the total number. The other arm of the curve is "radiation pressure noise." Every time a photon reflects from a mass, there is a recoil momentum—radiation pressure—given to the mass. If the number of photons striking the mass fluctuates, then the radiation pressure fluctuates. At any given power level, there's a frequency at which a minimum occurs between shot-noise and radiation-pressure fluctuations in the interferometer. On one side of the minimum, we're not sensing accurately enough because we're not using enough photons, and on the other side, we're kicking the masses around because we've got too many photons. If we turn the power up, the shot noise's contribution to sensitivity decreases and the radiation-pressure contribution increases, so that the minimum moves to a higher frequency. The curve labeled "quantum limit" is the line of all possible minima for a 20-pound test mass, and it's set by the Heisenberg uncertainty principle. Unfortunately, the quantum limit, like the speed limit, is not just a good idea, it's the law. And it's a lot more rigidly enforced than any speed limit.

It takes very high light levels inside an interferometer to reach the quantum limit at higher frequencies. But there are some tricks we can play with the interferometer's optics to help reach this sensitivity. We can use our light more efficiently by recycling it. Once again, it's all done with mirrors. Remember that the light from the arms eventually hits the beam splitter, which, in the absence of a gravitational wave, delivers no light to the photodetector. By then, the light has already done its job of measuring the arm's length. But where did the light's energy go? It turns out that, by the law of conservation of energy, when destructive interference happens for light traveling in one direction, constructive interference must occur for light traveling in another direction. In this case, the beam splitter sends the constructively interfered light back toward the laser. It would have been nice if the photons didn't have to leave, because the interferometer loses energy that way. In fact, as my colleague Ron Drever discovered, they don't have to go. We can install a recycling mirror between the beam splitter and the laser and reflect the photons back in. After all, this is the age of conservation.

It works just like the Fabry–Perot cavity. There's an F wave, which is the direct reflection off the recycling mirror. In the absence of a gravitational wave, the beam splitter and the two interferometer arms act like a single mirror, sending the light back toward the laser. Now if we put the recycling mirror in the right place, the light that leaks through it—the C wave—can interfere destructively with the F wave. As a result, there's no light reflected to the laser—it all goes back into the interferometer. Believe it or not, this actually works. We've demonstrated 20-fold recycling factors in benchtop experiments. In principle, we ought to be able to make LIGO's laser appear a thousand times more powerful this way.

This recycling method is useful when we want to capture gravitational waves with high fidelity over a broad frequency spectrum. In many cases we won't know the wave's frequency ahead of time, or we may simply want to hear something go "bang!" But gravitational waves don't just come from burst sources. There are other sources—rapidly rotating neutron stars, for instance—that can broadcast gravitational waves continuously on one frequency. In this case, we might know the frequency, especially if the neutron star is a pulsar, but we won't know the amplitude that the wave should have. To study

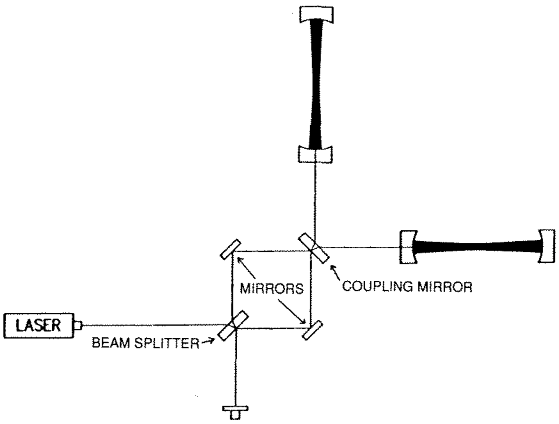

Resonant recycling interferometer.

Like the 200-inch Hale telescope at Palomar, which has stayed on the cutting edge of astronomy over the years as its photographic plates have given way to sophisticated electronics, LIGO is not a single experiment but an experimental facility.

this kind of periodic source, we basically want to make the equivalent of a radio receiver—something that detects a very narrow frequency band, but is very sensitive at that frequency. There's a trick that works for that, and—surprise!—it involves recycling photons.

It's called the resonant recycling interferometer. Here's the trick: fill the arms with light that leaks out in a time equal to half the period of the gravitational wave we want to find, then pass the light leaking out of one arm into the other. For half of its period, the wave is stretching one cavity and shrinking the other. In the next half period, the motion reverses. If we switch the light back and forth between arms in synchrony with their alternating stretching and shrinking, so that one batch of photons is always in the long arm and the other batch is always in the short arm, the signal can build up over a long time. Ultimately, light appears at the photodetector if there actually is a gravitational wave at the chosen frequency. Otherwise, the photons dissipate among the mirrors in the system.

A new trick, called dual recycling, was just demonstrated this March by Brian Meers at Glasgow University. And yes, it's also done with mirrors. It's a modification to the broadband recycling interferometer that gives it the narrowband character of resonant recycling. In a broadband recycling interferometer, a periodic gravitational wave will immediately induce a periodic signal at the photodetector, although it may be very small compared to the background noise. The new mirror recycles this periodic signal, allowing it to build up inside the interferometer, so that each passing crest of the wave gives it a boost before it finally leaks out to the detector. We call this dual recycling, because we recycle both outputs from the beam splitter. We can now set the detector's bandwidth by choosing the reflectivity of the signal-recycling mirror. If the reflectivity is zero—as if the recycling mirror wasn't there—then we're back to the broadband recycling interferometer. If the reflectivity is high, then the bandwidth becomes very, very narrow, approaching the sort of frequency response of the resonant recycling interferometer in the previous case.

Perhaps LIGO's most important feature is that it will evolve. Like the 200-inch Hale telescope at Palomar, which has stayed on the cutting edge of astronomy over the years as its photographic plates have given way to sophisticated electronics, LIGO is not a single experiment but an experimental facility. The graph on page 9 shows the noise levels that we anticipate in LIGO's first generation of interferometers, which will be built from readily available components. For instance, the photon shot noise was calculated for an argon laser producing five watts of green light—the laser seen on page 6, in fact—installed in a broadband recycling interferometer with a 30-fold recycling factor, close to what we've already demonstrated on the benchtop. LIGO, in this incarnation, would have significant scientific potential—it would have been wonderful to have had it to view Supernova 1987a. Finding coalescing neutron-star or black-hole binaries will likely require a more advanced interferometer, one incorporating the improved lasers, optics, and materials that are now being developed, as well as the valuable experience gained by running LIGO's first detectors. If we start now, the first LIGO detectors will be on line in the latter half of the 1990s. This will mark not the closing of a project, but the opening of a new window through which to look at the universe. ☐

Assistant Professor Fred Raab joined Caltech's LIGO project in 1988. He had previously been at the University of Washington, where he worked on a variety of high-precision measurements of fundamental forces. He earned his BS from Manhattan College and his MS and PhD from SUNY, Stony Brook. This article was adapted from his Centennial Seminar Day talk.