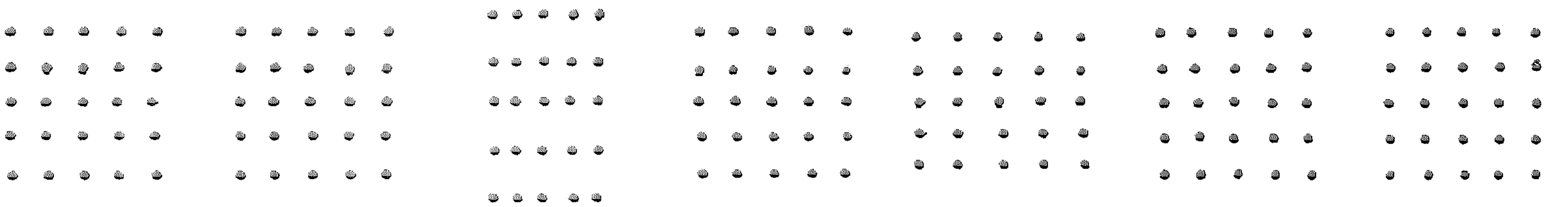

The Search for Gravitational Waves

by Ronald W. P. Drever

Illustrated above is a possible apparent effect of a gravitational wave pulse, traveling in a direction into the page, on a square array of light test particles floating freely in space. Seven successive snapshots taken at intervals of one-thousandth of a second might reveal the kind of relative particle motions indicated. The motions shown are highly exaggerated – by a factor of around a billion billion.

CALTECH's first prototype gravity wave detector, four years in planning and construction, is approaching completion in its initial form and is nearly ready for preliminary experiments. The sensitive instrument, with 40-meter-long arms, is housed in its own temperature-controlled and vibration-isolated "building" attached to the north and west sides of the Central Engineering Services facility. It may seem strange to be putting so much experimental effort into gravitational waves when it is not at all clear that anyone has ever even observed them yet. But if this project — or any of the others around the world — do actually find gravitational waves, it could lead to a new window on the universe. Right now we don't even know all the things we might see through that window; an entirely new area of astronomy would be opened up.

Gravitational waves are a form of radiation predicted by Einstein's relativity theory. Einstein showed that when masses move, the resulting changes in the gravitational field propagate out at a speed equal to the speed of light. An oscillating acceleration of masses would give a periodic effect, and it is quite reasonable to call this propagating strain in space a gravitational wave. Analogies can be made with electromagnetic waves, which are produced by accelerating charges — such as the electrons carrying the currents in the antenna of a radio broadcasting station. When thinking about gravitation, we can in some ways regard mass as analogous to electrical charge in electromagnetism, and it might not be too surprising that accelerating masses might produce gravitational waves.

But there are some basic differences between gravitational waves and radio waves. In electromagnetism, for example, we have both positive and negative charges; but in the case of gravitation, as far as we know, there is only one sign of "gravitational charge." Although in principle I might produce gravitational waves by shaking a heavy object up and down, when I make it accelerate in one direction, my body and the ground I am standing on would tend to accelerate in the opposite direction. And as my body and the ground would have effectively the same relative gravitational charge as the object being shaken, they would give an effect which would tend to cancel the waves from that object. In fact, generation of gravitational waves depends on the overall change in distribution of mass, or flow of mass, in a system. In principle the simplest gravitational wave generator might consist of a pair of massive objects made to oscillate with respect to one another, or rotate about one another.

How could we observe the waves from such a generator? In the radio wave analog we might set up an antenna and look for the relative motions of electrons and positive charges within it, detecting the resulting currents. In the gravitational wave case, with all masses equivalent, we have to do something different; and the simplest thing we might consider looking for could be relative motion of a pair of free test masses placed some distance from one another. And indeed this is a way of detecting gravitational waves, at least in principle.

So, to see a gravitational wave, we might set up two test masses, floating in space or suspended somehow so that they are free to move, and look for relative motions. It turns out that gravitational waves, like radio waves, are transverse in nature, and the maximum effect is obtained when the line joining the two particles is perpendicular to the direction of propagation of the wave, and is suitably oriented relative to its polarization. If the distance between the masses is small compared with the wavelength, then the change in separation is proportional to the separation, and the fractional change in separation may be taken as a measure of the amplitude of the wave. Another pair of test masses placed to reveal transverse motions in a direction orthogonal to that tested by the first pair will give relative motions of the same magnitude but opposite sign; and the usual definition of wave amplitude is in fact twice the fractional effect seen in either direction alone.

How big are these effects? Very, very small in most cases — and that is what makes experiments hard. If we tried to generate gravitational waves in the lab by making masses move, the radiated waves would be so weak that there would be no hope of detecting them by any method we can think of just now. There are, however, some violent astronomical processes that might generate detectable gravitational waves, such as, for example, a supernova explosion — the collapse of a normal star to give a neutron star.

When a star burns up its nuclear fuel, the central part can become unstable and start to collapse inward upon itself. As the parts get closer together, they're pulled more and more strongly by their own gravity, and the collapse becomes very rapid in its last stage. Sometimes a neutron star is formed — a star just a few miles in diameter with the extremely high density of nuclear matter. In the final stage of the formation of the neutron star the system may oscillate in shape from a pancake to a football shape, or bounce in other ways in times as short as a thousandth of a second, and these violent motions can produce a significant flux of gravitational waves.

If the mass of the original star is rather larger than the mass of the sun, the collapse can proceed further to form a black hole — which has a density so high that the gravitational attraction prevents anything getting out of it — and this process may produce gravitational waves quite well too. Of course not even gravitational waves can escape from inside a black hole, but as pieces of matter cross over the boundary into the black hole, they're accelerating rapidly and could throw out gravitational waves just before they vanish.

The size of a gravitational wave signal from a supernova depends on how far away the supernova is. If it were in our own galaxy, that is, relatively close, the best estimates are that the amplitude of the wave might be about 10^{-17}. That means, for example, that if our gravitational wave detector had test masses a meter apart, the relative motion we would get from this supernova would be only about 10^{-17} meters. That's less than the size of the nucleus of an atom. If we imagine our detector expanded to the size of the earth itself, the opposite parts would move even then by a distance of only about one angstrom. So gravitational waves might not, at first, look a very promising field for experiment.

But nevertheless, Joseph Weber, of the University of Maryland, had the courage in the late 1960s to start some experiments to look for gravitational waves. It is generally considered that he did not find them, but it was extremely important work because it triggered many other experiments in the field. Weber's technique did not use two separate masses, but made the masses the two ends of a big aluminum bar a couple of meters long. The elasticity of the bar was such that it had a resonance frequency for longitudinal vibrations in the region of a kiloherz, which corresponds roughly to the frequency you would expect from gravitational waves from supernova collapse. As a gravity wave passes, it might set up oscillations in the bar, which would persist for some seconds and might perhaps be detectable.

Weber fitted bars like this with piezoelectric transducers connected to sensitive amplifiers, suspended them in vacuum tanks with isolation from seismic disturbances, and looked for changes in oscillation. He achieved quite good sensitivity, perhaps not far from 10^{-16} in gravitational wave amplitude, and, surprisingly, seemed to see many pulses he thought might be gravitational waves. They were, however, larger than expected, much more frequent than supernovae in our galaxy — which probably occur about once per 30 years on average — and indeed so numerous that they

The whole view (both arms) of the Caltech 40-meter prototype laser interferometer gravitational wave detector pictured on the cover. Test masses are suspended inside the vacuum tank at the center of the photo and in similar tanks at the far ends of the vacuum pipes. Laser beams within the pipes monitor relative motions of the masses.

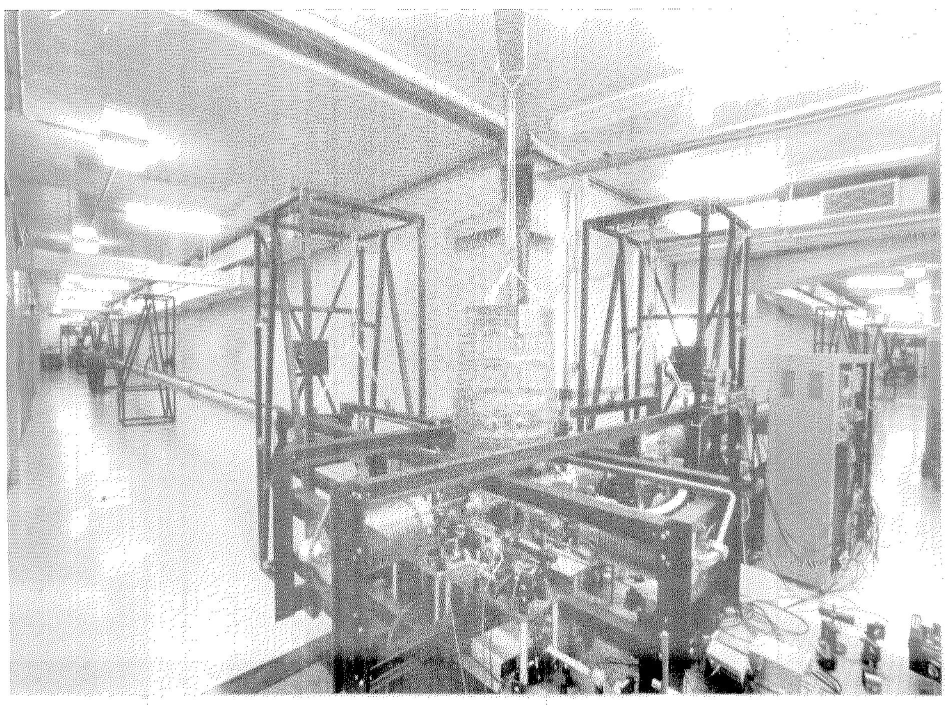

would correspond to a very puzzling loss of energy from the galaxy if they originated there. Of all those who got into the field after Weber nobody else found signals that seemed to correspond to gravitational waves at the level and rates suggested by Weber's results, and the consensus view is that he may have been seeing some other kind of disturbance. An important thing about this work, however, was that it did show that experiments could be carried out at a surprising level of sensitivity, and it brought many other researchers into the field.

Most of the gravitational wave experiments carried out so far have used bar-type detectors similar in principle to Weber's. Currently the most sensitive of these that have gone into operation is at Stanford, in which a large bar is suspended in a liquid helium cryostat to reduce thermal noise, and the vibrations are monitored by a superconducting accelerometer attached to one end of the bar. This instrument has had a sensitivity of order 10^{-18} to date — and achieving this has been a considerable feat, one that should be good enough to see supernovae in our galaxy. But these occur at such a low rate that observation is not very likely. Detection of unpredicted signals is not at all impossible at the level of sensitivity of this detector, however. Really critical experiments have not been possible yet due to the absence so far of another working detector of similar sensitivity with which to correlate data.

In order to do a really promising experiment, however, it would be good if we could aim at getting signals at a rate of around once per month. One way to do this would be to make a more sensitive detector so that we could look further out and cover more supernovae. It looks as though we could expect signals at about once a month if we could detect supernovae events out into the Virgo cluster of galaxies, but the amplitudes of the signals from these distances would probably be of order 10^{-21} or so. That does not look very easy. It was hard enough to do the early experiments at sensitivities of about 10^{-16} or 10^{-17}, and now we are asking to do experiments four orders of magnitude better in amplitude sensitivity, which corresponds to eight orders of magnitude better in terms of energy flux.

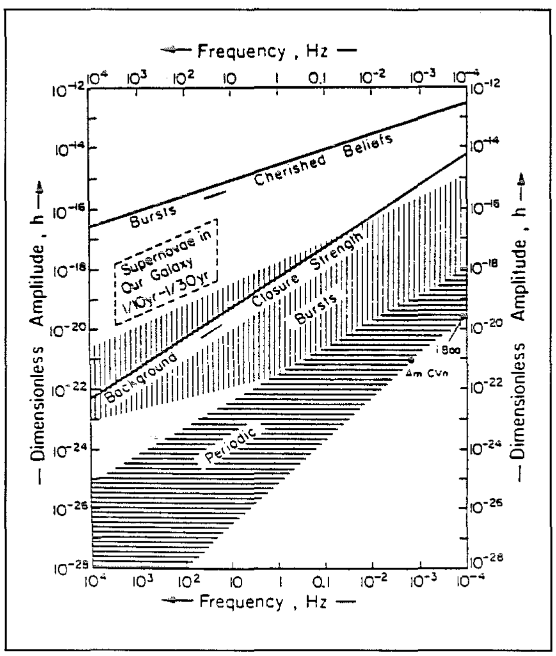

Estimates of the strengths of gravitational waves arriving at the earth over a range of frequencies. Regions where bursts or pulses at rates of order one per month (from black hole or neutron star formation) and periodic signals (from binaries and pulsars) might be expected are shown by shading. The lines indicate upper limits on bursts and stochastic gravitational wave background suggested by "cherished beliefs" such as mass-energy conservation, observed mass distribution in the galaxy and universe, and so on.

K.S. Thorne, Rev. Mod. Phys., 52, 285 (1980)

But, if we could achieve sensitivities of this order, quite a range of possibilities for seeing gravitational wave signals of different kinds might open up. In addition to supernova and stellar collapse signals at reasonable rates, we might see signals from collisions between black holes, from pairs of newly formed neutron stars rapidly rotating about one another, and from other possible sources. There are also continuous signals emitted by pulsars, known and unknown. A pulsar spins at a very steady rate — about 30 revolutions per second for the Crab Nebula pulsar, for example — and unless it is perfectly symmetrical about its rotation axis it will emit continuous gravitational radiation. In general, if the pulsar is precessing as well as spinning it will emit gravitational radiation at a frequency close to the repetition frequency of its radio pulses, and also close to twice this frequency. The strengths of the waves depend on the degree of asymmetry of the pulsar and so are hard to predict accurately, but the amplitudes are probably in the region of 10^{-26} for the Crab and the Vela pulsars. This is small but much easier to detect than a single pulse, since the signal is continuous and the frequencies known quite well. In fact I think that detection of pulsar and supernova signals presents problems of roughly equal magnitude, and when one is solved, the other will be near solution too. We would also expect a whole spectrum of gravitational wave signals at lower frequencies coming from more massive black holes — with masses up to a million times the sun's mass — in distant galaxies and quasars.

Can we build equipment sufficiently sensitive to pick up some of these signals? That's the task we set ourselves. The factor of improvement that we're expecting we may have to get over the previous experiments corresponds roughly to the difference in sensitivity between the human eye and the 200-inch telescope at Palomar Observatory, so the task may take some time.

The method that we've chosen to develop does not use bars at all. It goes back to the idea of watching free masses move relative to one another as the gravitational wave goes past. And if we use light to detect the motion, we can put the masses a long distance apart so that we get bigger effects to observe — an important advantage over the bar type of gravitational wave detector. Our experiments at Caltech are starting off with three test masses 40 meters apart — and maybe we'll make it several kilometers before we're finished. We hang each of the masses by wires, like a pendulum, so that it can make small movements in a horizontal direction fairly freely. We arrange the three masses in the form of an L, so that if a gravitational wave traveling in a vertical direction makes the two masses on one side of the L move apart it will at the same time make the pair forming the other arm of the L move together. Thus one distance gets longer and the other shorter. We need a way to measure that change, and the method we have chosen is based on an optical interferometer.

In principle we might use a Michelson interferometer. We could attach a partially reflecting beam splitter to the mass at the corner of the L, and illuminate it by a beam of light from a source such as a laser, so that half of the light passes through the beam splitter toward the mass at the end of one arm of the L, and the other half is reflected toward the mass at the end of the other arm. Mirrors on the two end masses reflect the beams back to the beam splitter, where they recombine and pass to a photodiode. The fringe pattern we get there depends on the relative phase of the light in the beams. If in phase, we get high intensity, if out of phase we get darkness. So if the masses move and make the difference in the path lengths change slightly, we may see a change in output from the photodiode. In principle this is fine, but once again the smallness of the effects causes problems. The wavelength of light is about 5×10^{-7} meters, and the gravitational wave motion is probably going to be smaller than that by a factor of 10^{10} or more, so we're not going to see whole fringes moving past. We can't even expect to see a motion of the fringe pattern by a reasonable fraction of the distance between one fringe and the next. We're going to have to look for changes corresponding to 10^{-10} of a fringe or less. And that's a bit difficult.

There are two fairly fundamental limiting factors for this kind of gravitational wave detector. One is the quantum limit to the sensitivity of the apparatus set by the uncertainty principle. If we try to measure the position of one of the test masses with very high accuracy, we make its momentum uncertain — giving it a little bit of a kick. If we look at the mass again a little later,

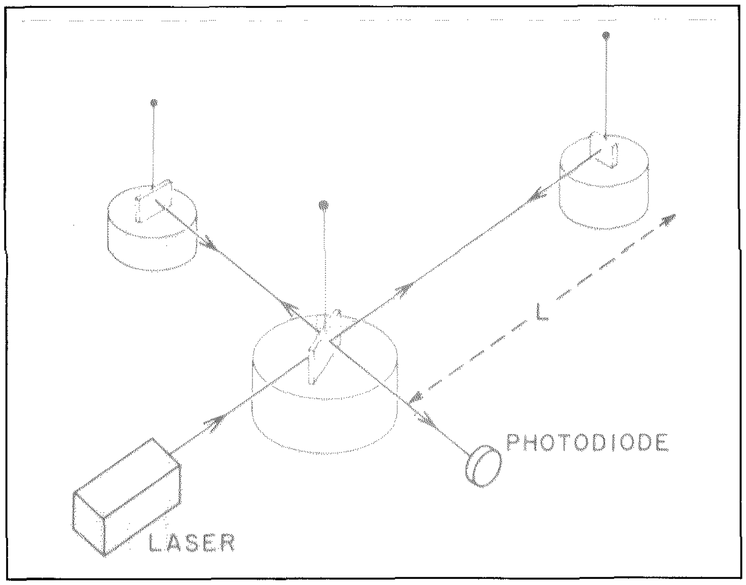

This drawing shows a possible arrangement of a simple, but insensitive, laser interferometer gravitational wave detector using three suspended test masses, with baseline L.

and we find it has moved, we don't know whether that's because we kicked it or because a gravitational wave went past. For test masses of reasonable size, and a gravitational wave pulse lasting a millisecond, this quantum limit corresponds to a wave amplitude of the order of 10^{-21} when the masses are 40 meters apart. So this is not a very serious limit in this case, and it becomes less important still if the distance between the masses is increased.

A much more serious problem comes from fluctuations in the output from the photodiode, arising from what might be called "photon counting error." The light detected by the photodiode can be regarded as made up of photons, and as they arrive essentially randomly at the diode, there will be fluctuations in the intensity of light recorded. These set a limit to the smallest change in relative phase of the interfering light beams which can be detected, and thus to overall sensitivity. This limit depends on the intensity of the light, and on the time available for the measurement — which for a gravitational wave pulse corresponds roughly to the pulse duration. For 1 millisecond pulses and a laser giving one watt of light, the sensitivity limit for detectable motions is about 10^{-15} meters. This is very good by ordinary standards, and indeed the first pioneering experiments on gravitational wave detection by laser interferometry were done using methods similar to this by Robert Forward of Hughes Laboratories some ten years ago. However we're now aiming at much higher sensitivity, and we have to reduce this photon noise by a large factor. If we tried to do it just by stepping up the laser power, we would need several megawatts of light from our laser continuously. This is hardly practicable, and

continued on page 24

The Search for Gravitational Waves

. . . continued from page 9

we would have difficulty paying the power bill.

But there are several other tricks which could make the system more sensitive. One is to make the light travel back and forth many times in each arm of the interferometer, making many separate reflections at curved mirrors on the test masses. If the masses move, there will then be a much bigger change in the total light path, giving a larger fringe shift to be observed. This idea was suggested by Ray Weiss of MIT, who is developing gravitational wave detectors of this type.

The number of times it is useful to make the light travel back and forth may depend on the losses in the mirrors, or, if losses are very small, on the time the light spends in the interferometer. If the light is kept in for too long, and the movement of the mirrors due to the gravity wave reverses its sign, there's going to be some cancellation of signal. In an instrument with 40-meter arms and mirrors allowing about 600 reflections of the light, the photon noise limit to sensitivity corresponds to a gravity wave amplitude of about 10^{-19} for 1 millisecond pulses and a laser power of 1 watt. With longer arms and lower-loss mirrors, so that light storage time is the limiting factor, the photon counting limit to sensitivity becomes of order 2×10^{-21} with either 4 km arms and 80 bounces of the light or 400 meter arms and 800 bounces. This would be getting somewhere near our target sensitivity — if there were no other problems.

Several technical difficulties with multiple reflection systems became apparent in early work at Munich and Glasgow, among them the finding that scattering at the mirrors can allow small but troublesome amounts of light to take unplanned paths through the system. There is also the simple fact that in a system with long arms the mirrors and the vacuum system containing them and the beams may have to be quite large — and correspondingly expensive. This led us to devise another approach. We use fairly small, highly reflecting mirrors on the test masses, arranged to form optical cavities between one another, in which the light bounces back and forth from a single spot on each mirror. If we arrange that the length of each cavity is precisely adjusted so that the light from the laser resonates within it, then there is a considerable buildup of light intensity. A slight change in the length of one of the cavities will change the phase of the light inside it, and we can observe that change of phase. In principle the system is somewhat similar to the previous one, except that here all the reflecting beams are superimposed on top of one another. With mirrors of equal loss this arrangement should give slightly better sensitivity than the multiple-reflection Michelson system, and it's potentially less expensive and more practicable on a large scale. The mirrors and beam pipes can be quite small. Even with 10-kilometer arms the mirrors would be only a few inches in diameter. The system does require, however, that the wavelength of the laser light be very well controlled to match the extremely narrow resonances of the long optical cavities.

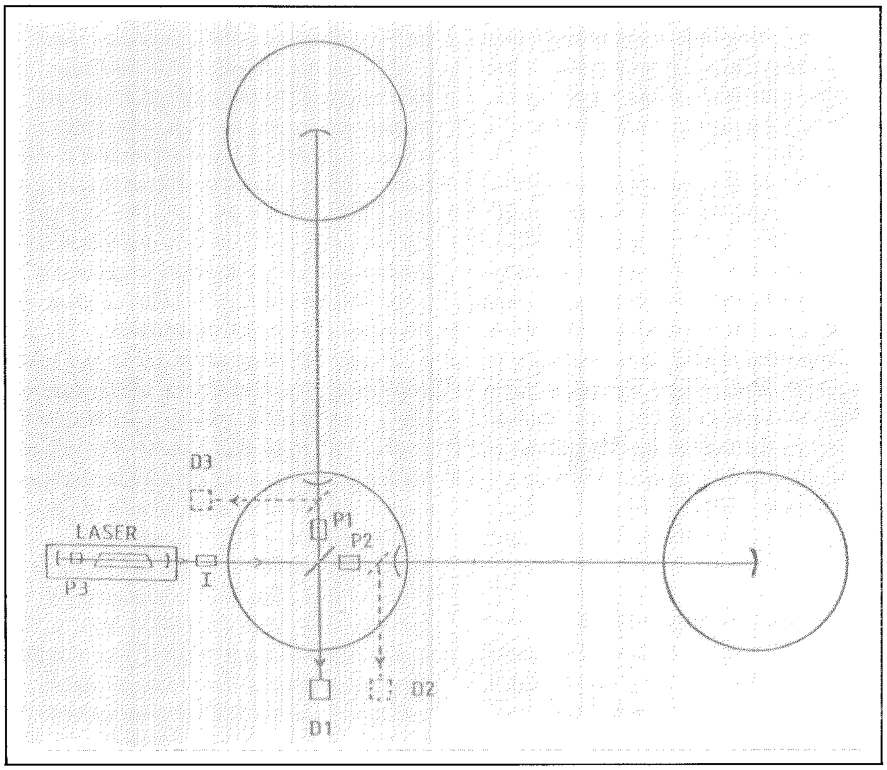

The simplified diagram above illustrates an optical cavity gravitational wave interferometer. Reflection of light many times between curved mirrors mounted on the test masses (indicated by circles) increases phase difference in the light coming back from each arm caused by motions of the masses. The phase difference is measured by photodiode D1, aided by electro-optical phase modulators P1 and P2 driven at a high radio frequency. Signals from auxiliary photodiodes D2 and D3 measure small deviations of laser wavelength from resonance with the cavities and may be used both to control average length of each cavity and reduce fast fluctuations in laser wavelength.

How do we do all this in practice? The first problem was to stabilize the frequency of a fairly powerful argon ion laser sufficiently well. This led us to come up with a new way of stabilizing a laser — or at least one we had not heard of before. We phase modulate the light from the laser at a high radio frequency, and direct the light coming back from the first mirror of one of the optical cavities onto a photodiode. There are two components in this light: one of them light from inside the cavity, and the other light which has just come from the laser and has bounced off the first cavity mirror. If the laser is exactly in tune with a cavity resonance, these components will be out of phase with one another and there will be no signal from the photodiode at the modulation frequency; but if the laser wavelength lies slightly to one side of the resonance, the phase difference will change, and a signal will appear at this frequency. The amplitude of the signal is a measure of the error in the wavelength of the laser, and we can feed this back to an electro-optical crystal inside the laser itself to rapidly adjust the wavelength of the laser closer to the cavity resonance. This technique was first tested at JILA (the Joint Institute for Laboratory Astrophysics) in Boulder, Colorado, in collaboration with John Hall there, and very shortly afterwards, we had it operating at the University of Glasgow, where we had begun experiments on laser

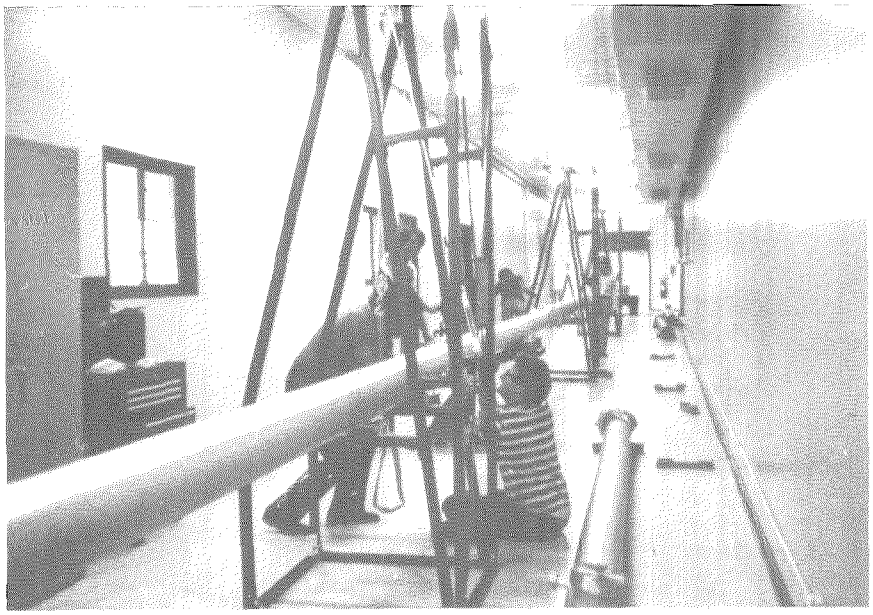

Graduate and undergraduate students assemble one of the 40-meter lengths of vacuum pipe.

gravity wave detectors some years earlier. It worked extremely well. And in the complete gravitational wave detector we could use the same basic method to keep the length of the cavity between the second pair of masses adjusted to the wavelength from the laser — and use the error signals to sense the effects of the gravitational wave. With these encouraging results we set to work at Caltech to start up the new project here from scratch.

We wanted to make the detector a fair size, so the first thing to do was to persuade Caltech to make us a building to put it in. And they made us a very elaborate lean-to shed along two sides of an existing building, giving us two paths 40 meters long in nice, quiet, stable, temperature-controlled conditions. And inside that building we started to make up the system. It was put together mostly by students — undergraduates and graduate students — largely under the direction of Stan Whitcomb, assistant professor of physics.

One of the practical problems was how to suspend the test masses inside the vacuum tanks. Each mass has to be free to move in response to the gravity wave, and yet at the same time it must be controlled in its orientation so the attached mirrors remain accurately lined up to reflect the light beam. That's rather awkward to do without introducing additional disturbances, but we did it by detecting the orientation with small auxiliary lasers and moving the wires from which the mass is suspended. It is hung by three thin wires from a little metal block connected to the moving coil drives of small loudspeakers which can tilt or rotate the block in any direction. The mass follows the block like a puppet on its string. This enables us to line it up as accurately as if it were quite rigid, whereas actually it's almost totally free to move in the field of the gravity wave. This tricky job was worked out in detail mostly by graduate student Mark Hereld.

To cut down vibrations from the ground, the vacuum tanks are mounted on special isolation pads with foundations separate from the building. The suspension inside each tank is supported on a stack of alternate layers of lead bars and pieces of soft rubber. Little toy rubber cars, which were very inexpensive and of just the right resiliency, separate the lead bars. It would have been nice to use models of expensive cars like Mercedes Benzes everywhere, but we couldn't find enough of these, and toy dune buggies worked just as well.

We are fairly confident that this kind of apparatus will eventually achieve the kind of sensitivity mentioned earlier, though it will probably have to be made larger still to really see the gravitational waves from frequent events outside our galaxy. We regard this first Caltech detector largely as a prototype for future instruments with longer baselines between the masses. Even for larger detectors, the sensitivity estimates mentioned earlier for millisecond pulses might seem a little marginal. It is important to note that new ideas and new technical tricks that can improve performance significantly are still turning up. For example, just a short time ago an idea came up that may give us the last little twist that we need to do the job for gravitational wave pulses. If we take the case of the simple Michelson interferometer as an example, we would usually employ a phase modulation technique which operates best when the system is locked to a dark fringe, with minimum light coming out to the photodiode. In this situation most of the light from the laser has to be going somewhere else, and in fact it passes out from the other side of the beam splitter and is wasted. We now propose to use additional mirrors to bring this light back into the system again in just such a way that it reinforces the light coming from the laser. If we do this correctly, the light will build up in intensity inside the system to an extent limited only by the small losses in the mirrors and other components. Only a tiny amount has to come out to the photodiode — until a gravitational wave comes along, changes the phase condition, and allows the extra light to come out at just the right time to give us a signal. It sounds very nice!

Of course, it probably won't all work out quite as well as that, but when you estimate the possible sensitivity that might be achievable with very low loss mirrors and a large system, it looks pretty promising. With a detector with 10-kilometer arms we might be able to get a sensitivity of order 10^{-22} for millisecond pulses, even with currently available laser powers. For longer pulses sensitivity should improve. So we may have something in reserve. And we have also very recently come up with further possibilities for improving sensitivity for pulsar signals in the same basic kind of apparatus. There are still very many practical problems but I think we can overcome them one by one.

There are also other ways of looking for gravitational waves. At low frequencies, waves with periods in the range of seconds to hours or larger, laboratory experiments face a background of fluctuating gravitational gradients from nearby moving objects — or even from changes in distribution of air density in the atmosphere — which would limit sensitivity; and there are obvious advantages in doing the experiments in space. Indeed a lot of effort, at JPL and elsewhere, is going into the most direct kind of experiment, in which the radio signals used to track the motions of distant spacecraft relative to the earth are examined for effects of gravitational radiation bursts. Fluctuations in radio propagation velocity set limits to sensitivity, but there could be relatively strong gravitational wave bursts present, and these experiments are very interesting. Laser experiments done totally in space could be more sensitive, but are more complicated and expensive, and therefore further off. It has been suggested that one might put an interferometer somewhat like our ground-based ones into space. Very long baselines and high sensitivity could be achieved by using laser interferometric Doppler techniques to measure changes in distance between shielded test masses in widely separate spacecraft, and the possibilities for observing the definitely predicted gravitational waves from known binary stars this way are being seriously considered in several laboratories, including Caltech, JPL, and JILA. These are very interesting possibilities, but they will take time to materialize, and I think it is likely that we may find gravitational waves first on earth perhaps with the kind of apparatus that we are working on now.

Is all this effort worthwhile? I think so, and not just for the practical applications of some of the spin-offs, such as low-loss optical systems and new ways for stabilizing lasers. And not just to check the correctness of relativity theory in another interesting type of prediction — although that's certainly well worth doing in itself. But gravitational waves almost certainly exist, and if we can observe and analyze them they will give us a new window in astronomy. If we see gravitational waves from a supernova we will be seeing into the heart of stellar collapse, into a region hidden to radio and light waves by the surrounding material. We might be able to see what happens in the formation of a neutron star or a black hole, in a way that is impossible by other means. If, by slightly different techniques, we could see the gravitational wave analog of the big bang radiation, we should learn about very early stages in the big bang. And above all we may find things that we have never thought of yet. The future looks promising — but the experiments are certainly challenging ones. □

Acknowledgments

This kind of research essentially involves group efforts, and I would like to acknowledge the contributions to the experimental work described of my colleagues at Caltech, particularly Stan Whitcomb, Robert Spero, Dana Anderson, Mark Hereld and Yekta Gursel, and also colleagues at Glasgow including James Hough, Henry Ward, and Andrew Munley. The Caltech project is supported by the National Science Foundation.