

Gravitational Waves: Astronomy's Newest Frontier

Bernard F. Schutz

*Department of Physics
University of Wales College of Cardiff*

Gravitational waves are Nature's weakest form of radiation, but instruments are now being developed that will have the sensitivity to detect waves from supernovae and collisions of neutron stars that take place in distant galaxies. Based on 3-km-long laser interferometers, and capable of sensing changes in length smaller than 10^{-18} cm, these detectors incorporate the latest optical technology and should reveal astronomical information that can be gathered in no other way.

Introduction

During the next decade, scientists are planning to build at several locations around the world some of the most unusual scientific instruments ever constructed. They will be enormous: two pipes of stainless steel, each 3-km long and 1 m in diameter, will be joined at right angles to make a giant "L". (see figure 1). They will be sophisticated: the pipes will enclose a very high vacuum, inside of which will run megawatts of single-frequency laser light, produced by highly advanced solid-state lasers. They will be the most accurate length-measuring instruments ever built: their sensitivity to variations of only 10^{-17} cm in the length of a 3-km arm will be equivalent to measuring 1-cm changes in the distance between the Earth and the centre of our Galaxy. They will produce one of the largest volumes of data ever generated by a scientific experiment, and yet the useful signals in the data will be extremely rare. And

finally, they will be expensive: £30 million or more for each instrument.

The object of this undertaking is to search for nature's most elusive form of radiation: gravitational waves. One of the last predictions of Einstein's theory of general relativity that has yet to be confirmed by direct observation, gravitational radiation, occupies an absolutely central place in theoretical physics. If gravitational waves do not exist, physics is in for a wholesale revision. But such is our confidence that they do exist (confidence based on the success of Einstein's theory to date and on a considerable body of indirect observational evidence for them) that our chief motivation in building these instruments is to use them as astronomical observatories. This is because gravitational waves carry unique information about their sources, information that we cannot get from observing any other form of radiation.

Attempts to detect gravitational waves are not new: detectors based upon metre-long cylinders of aluminium

were invented by Joseph Weber at the University of Maryland in the early 1960s, and they are still under active development. Indeed, if a supernova explosion occurs in our Galaxy within the next five years, these "bar" detectors may well make the first ever direct detection of gravitational radiation. But the sensitivity of such bars is limited by current technology, and the new generation of instruments based upon lasers offers much higher sensitivity over a far greater range of frequencies.

The year 1990 is a critical year in the development of these laser interferometric gravitational wave detectors. The scientific funding authorities of several countries will decide the fate of proposals to build five such instruments. Proposals have been submitted by a British-German collaboration an Italian-French collaboration, an American collaboration (to build two), and an Australian collaboration. At Cardiff my research group is involved both as co-proposers of the British-German project and also as coordinators of the development of the computer systems and software that will be used to analyse the data from all the instruments. The various detectors must operate as a single worldwide network; at least two are required to verify that a gravitational wave has been detected, and at least three to determine what direction it came from. Optimism is high that most of the proposals will be funded in 1990.

In this article, I will try to convey some of the reasons that these projects are being given a high priority by the scientific communities of many countries. I will begin with a short introduction to gravitational waves, followed by a description of how detectors work, with special emphasis on their technological challenges. Then I will review briefly the astronomical sources of the gravitational waves that we expect to find, and the sort of information that the waves will convey. For more detailed discussions, see Thorne (1987) or Schutz (1989).

Gravitational Waves

The simplest description of gravitational waves is that they are small ripples in the gravitational field that move through space with the speed of light. Waves are not present in Newton's theory of gravity, in which changes in the gravitational field are felt instantaneously everywhere in space. But Einstein's relativity requires that no

Fig. 1. The L-shaped form of a large laser-interferometric gravitational wave detector.

Fig. 2. The distortion of the Earth by the tidal force of the Moon.

influences travel faster than light, and this has the consequence that in any relativistic theory of gravity, changes in the field will move outwards at a finite speed. In general relativity, which is Einstein's replacement for Newton's theory of gravity (see Schutz 1985), this speed is the speed of light itself. General relativity is now well tested experimentally (Will 1986) and forms the basis of our predictions about gravitational waves.

The gravitational forces carried by the waves are called *tidal forces*. What this means is that only *differences* in the gravitational force of the wave across one's detector are measurable. The Earth falls freely in the overall gravitational field of the wave, just as it does in the gravitational field of the Moon, and just as a "weightless" astronaut does in the field of the Earth. The overall gravitational force therefore has no effect on the structure of the Earth (astronaut, detector), but any *differences* in the gravitational force across the Earth (astronaut, detector) can in principle be measured. On the Earth, the differences in the Moon's field give rise to the tides; in our detectors, the "tidal forces" exerted by the waves give rise to tiny distortions in the lengths of the arms.

Figure 2 shows how tides on the Earth work. It is not surprising that the oceans nearest the Moon should bulge away from the Earth, since they are attracted more strongly to the Moon than the Earth as a whole is. What is more surprising at first sight is that the oceans on the side of the Earth that faces away from the Moon should also bulge out away from the Earth. This is because these oceans are attracted more weakly toward the Moon than is the Earth as a whole, so the Earth is falling away from them, leaving them trailing out behind. The characteristic pattern of a tidal force is therefore an ellipse: a circular Earth is distorted into an elliptical shape.

Because gravitational waves exert tidal forces, their action on matter bears a close resemblance to the effects shown in Figure 2. In Figure 3 we see the basic picture of how a gravitational wave affects matter. Consider only the top row of diagrams at first. Imagine a ring of particles arrayed around a circle, with one in the centre, sitting out in empty space far from anything else. There are no physical connections among the particles. Now suppose that a gravitational wave arrives from above, travelling towards the paper. The effect of the wave is to distort the circle into an

Fig. 1

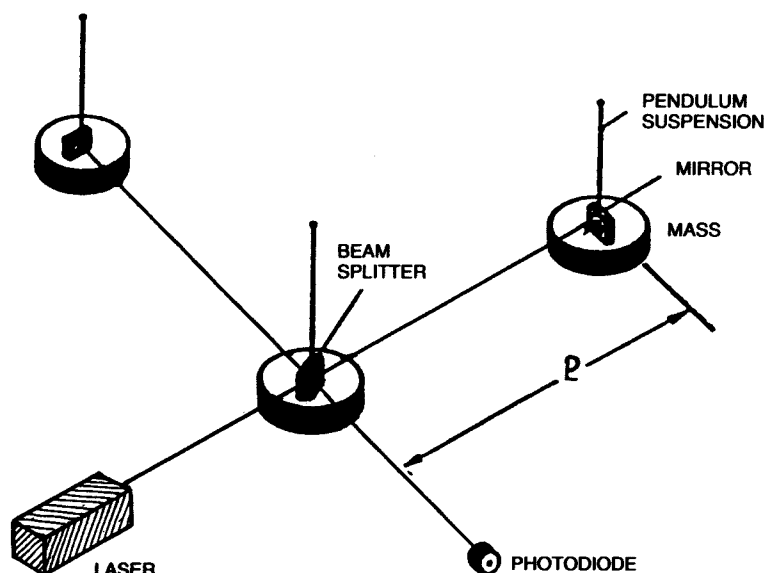



Fig. 2



ellipse. The axes of the ellipse stay fixed in direction, but oscillate in length.¹ Note the resemblance to Figure 2. The chief difference is that gravitational waves in general relativity are *transverse*, so that the distortions occur only in a plane perpendicular to the direction of travel of the wave, while the tidal forces on the Earth are forces in a plane containing the direction to the Moon.

There is nothing unique about the orientation of the ellipses in Figure 3, and in the second row of Figure 3 we show the action of a wave oriented at 45° to the first. These orientations are the *polarisations* of the gravitational wave, and, (just as in electromagnetism), any wave can be described as a combination of the two polarisations shown. We will see that the detectors themselves will respond only to one polarisation.

A remarkable and important property of gravitational waves follows from the fact that the force is a tidal force, i.e. that it is the difference in the gravitational force across the circle in Figure 3. If we make the circle twice as large, then the difference of the force across it will be twice as big, and the amount by which it will enlarge or contract will be doubled as well. The result is that the shape of the ellipse in the figure is independent of its size: it is a characteristic of the wave itself. If the radius of the original circle is called l and the maximum displacement of a particle is called δl , then the ratio $\delta l/l$ depends on the wave but not on l . We define the *amplitude* h of the wave to be twice this ratio:

$$h \approx 2 \frac{\delta l}{l}$$

This amplitude falls off as the wave expands away from its source. The amplitude at a distance r from the source is proportional to $1/r$.


From our study of possible sources of gravitational waves below, we will see that a typical gravitational wave amplitude at the Earth for a wave produced in an astronomical object is 10^{-22} . The smallness of this number is an illustration of the weakness of gravitational waves. When such a wave passes through a detector, it produces a distortion only some 10^{-22} of its length. If the detector is an aluminium bar 1m long, this corresponds to a change in its length of only 10^{-20} cm. The laser

¹For readers who are familiar enough with general relativity to worry about coordinate conditions and so on, the distances shown are proper distances as measured with conventional rulers by a local inertial observer sitting on the central particle. In this local inertial frame, the light that we will later imagine passing up and down between particles suffers no gravitational redshift.

Fig. 3. The tidal effect of a gravitational wave on a circular ring of free particles in empty space. The top row shows one polarisation, the bottom the orthogonal one. The polarisations are called "+" and "x", respectively, after the orientation of the axes of their ellipses.

Fig. 4. The action of a gravitational wave with the "+" polarisation hitting a cylindrical bar broadside. Note that a wave with the "x" polarisation will not stretch the bar at all, nor will a wave approaching along the axis of the cylinder. The bar is a polarised detector with maximum sensitivity to waves from the side.

Fig. 3



detectors now proposed are to be 3 km long, in order to take advantage of the tidal effect: gravitational waves produce distortions in these instruments that are 3000 times larger than in a bar.

Despite the weak effects that gravitational waves have on detectors, they carry huge amounts of energy. Even a wave with an amplitude of 10^{-22} at a frequency of 1 kHz carries an energy flux that is some 10^6 times brighter than the brightest star in the night sky. The trouble is that the energy goes right through us: the weakness of the gravitational interaction ensures that almost none of this energy is left in our apparatus. This gives gravitational waves great penetrating power. Gravity itself penetrates everywhere: one's weight does not depend on whether one stands on the ground or inside a lead-lined strongbox. Gravitational waves have a similar penetrating power. A detector will respond the same to waves regardless of whether we place it in the open air or deep underground. This also means that gravitational waves are not attenuated much by passing through matter, and this fact is both a boon and a curse for gravitational wave detection. It is a curse because, as we have just seen, gravitational waves passing through our detectors disturb them only very little. This makes the waves extremely hard to detect. It is a boon because the gravitational waves have been disturbed only a small amount by any other matter they might have encountered on their way to us. This makes them ideal carriers of astronomical information: we see them in exactly the form in which they were emitted.

How Detectors Work

The two basic forms of gravitational wave detector, bars and lasers, are both designed to sense the stretching produced by the tidal force of the wave, as in Figure 3. But they do so in very different ways.

Bar detectors

Since their invention by Weber during the 1960's, bar detectors have undergone a long and steady improvement in many laboratories (Weber 1960, 1967; Blair *et al.* 1989). Present operating bars have a sensitivity of about $h \approx 10^{-16}$. While this is much less sensitive than we would like, it may be sufficient to see rare, relatively nearby events, so that it may

well be that bar detectors will make the first observations of gravitational waves sometime during the next five years or so.

The basic design of a bar

Imagine a long cylinder of aluminium floating in empty space in place of our ring of free particles, and place the centre of the cylinder where the central particle was. Now let a gravitational wave with the top polarisation of Figure 3 arrive at the cylinder from the side, as in Figure 4. Then the stretching and compressing action of the tidal forces along the axis of the cylinder will tend to stretch and compress the bar, while the tidal forces along the perpendicular direction will have a small effect on account of the smaller size of the cylinder in that direction. The stretching caused by the wave will be resisted by the elasticity of the cylinder, and the result will be in general a fairly complicated oscillation of the cylindrical bar along its axis. One can measure the amplitude of the gravitational wave if one can measure the motion it induces in the bar.

Fig. 4



In practice, one cannot place the bar in empty space; it has to be in the laboratory. So, typically, the bar is suspended from a framework by a single wire wrapped around its middle. The framework is isolated as much as possible from ground vibrations. The bar is cooled to liquid helium temperatures (4.2K) to reduce thermal vibrations. Such cryogenic bars are typically 2m long and a tonne or so in mass. They are instrumented with a variety of methods of reading out their vibrations, different groups developing different techniques.

Limitations of bar detectors

Although bars are still under active development at a number of institutions, there are fundamental limitations on their performance that have influenced several groups to move to laser interferometers. One limitation is bandwidth: at present, bars operate as resonant detectors, and are sensitive to gravitational waves only within a few Hertz of the frequency of the fundamental mode of oscillation of the cylinder. They therefore lose much information that would be useful to astronomy.


A second, and more worrying, limitation is called the "quantum limit". Basically, for a typical bar, the energy deposited in it by a gravitational wave with an amplitude of about 10^{-20} is equal to the energy of one phonon (one quantum) of oscillation of the bar. This means that to detect gravitational waves at a level of 10^{-22} , one needs to measure delicate changes in the quantum state of the bar. While this is possible in principle,² it is difficult to implement for massive bars, and so far there is no clear path to breaking the 10^{-20} barrier for

bars. Nevertheless, a bar with a sensitivity of $h \approx 10^{-20}$ would be a very interesting detector to have; if a supernova exploded within our Galaxy, such a bar would have an excellent chance of registering it. The next generation of

²The theory of how to measure below the quantum limit is in fact one of the principal contributions that gravitational wave research has so far made to the rest of physics (Caves *et al.* 1980). Going by the generic name of "squeezing", it is of particular interest these days in quantum optics and telecommunications research (Leuchs 1988).

Fig. 5. The action of a gravitational wave with the "+" polarisation hitting an interferometer from above. A wave with the "x" polarisation will not change the relative lengths of the two arms. Waves approaching in the plane of the detector have varying effects, depending on their direction and polarisation. The antenna pattern of an interferometer is quite complex.

Fig. 5



bars, now under development, will be cooled below 0.1K, and could in principle reach this goal. At least until interferometers reach somewhere near 10^{-21} , bars will continue to be the subject of active development.

Interferometric detectors

An interferometer is closer to the idealised free particles of Figure 3 than is a bar detector. As Figure 5 shows, its two arms are marked out by the central particle of the polarisation ring and two particles on the ring 90° apart. As the wave passes, its effect is to shorten one arm and lengthen the other, and then to reverse this. An interferometer uses coherent light from a laser, divided between the two arms, reflected off mirrors attached to the end particles, and then recombined into an interference pattern to sense changes in the relative lengths of the two arms. In practice, one cannot have perfectly free particles on the Earth; but one can arrange that they are free to move in the horizontal direction by suspending them by wires from supports well isolated from ground vibrations. As long as the gravitational wave has a frequency much higher than the pendulum frequency of the suspension (about 1Hz), the mirrors will move as free particles. No resonance condition operates here, so these detectors are intrinsically broadband.

The technology of interferometry

Only in the last few years has it become possible to design an interferometer that could reach $h \approx 10^{-22}$. The success of these instruments is crucially dependent on optical technology. Of

paramount importance is getting high levels of light power in the arms of the interferometer. This is because of another kind of quantum limit: light comes quantised as photons, and any distance measurement using light is subject to an uncertainty in the position of the photon.

Only by using large numbers of photons (large light power) can this uncertainty be reduced. The distance changes we wish to measure are as small as 10^{-12} of the wavelength of the photons we will use, and this forces us to use an extremely large number of photons per measurement. With 1000 or more measurements per second, this quantum "shot noise" limit translates into a requirement for several megawatts of light power stored in the arms of the interferometer.

No lasers exist that can supply megawatts of power continuously, but by recycling all the light back into the system after it has been used for a measurement, power in the arms can be built up until the laser is simply supplying the inevitable losses due to mirror imperfections and so on (Drever 1983; Meers 1988). If the mirrors are good, and we have tested extraordinary mirrors made by British Aerospace and other firms that have losses of only about 2 parts in 10^5 per reflection, then the losses are small and lasers with more modest power outputs of 50-100W suffice. Solid state Nd/YAG lasers can supply such power, and will probably be used in the proposed detectors.

Other technologies are also required. The light path along the arms must be in a high-quality vacuum, in order to avoid changes of refractive index in air caused by localised fluctuations in air pressure.

The suspension of the mirrors must isolate them well from ground vibrations. And squeezing, this time of the quantum state of the photons in the detector, may allow us to break through the photon quantisation limit mentioned above. Ironically, although squeezing was first studied for bar detectors, it is only for interferometers that it has so far proved a practical proposition (Leuchs 1988). For further discussion of interferometers, see Jeffries *et al.* (1987).

Present prototype detectors

Research groups at Glasgow University, at the Max Planck Institute for Quantum Optics (in Garching near Munich), at the California Institute of Technology, and at Tokyo University all operate working prototype interferometers, with arm lengths that range from 10m to 40m. These have achieved sensitivities of around $h \approx 2 \times 10^{-18}$, comparable with bar detectors. A number of other groups operate special-purpose interferometers that are designed for the development of various techniques for full-scale interferometers; these include groups at CNRS in Orsay (near Paris), at INFN in Pisa, and at the Massachusetts Institute of Technology. The prototypes have been used to take data from time to time, but they are primarily engineering testbeds, and the groups involved do not anticipate making long observing runs until they have built the full scale detectors.

Proposed detectors

A collaboration between British and German scientists has led to recent submission of a joint proposal to their respective funding agencies, SERC and

Fig. 6. The prototype 10m detector at the University of Glasgow: the central mass is in the vacuum tank at rear right.

BMFT, to build a single large-scale interferometer. This proposal is typical of others that have been submitted in the USA, Australia, France and Italy, so I will concentrate most of the description on it. In the UK, the experimental scientists are located at Glasgow University. They have design support from the Rutherford-Appleton Laboratory, near Oxford. The analysis of the data will be done by my group at Cardiff. In Germany, the principal group is an experimental one at the Max Planck Institute for Quantum Optics.

In order to take advantage of the tidal nature of the gravitational wave force, the proposal is for an interferometer with arms 3km long. This is about as long as one can reasonably expect to find a site for in Europe. The vacuum pipes carrying the beams would be about 1m in diameter; this would allow at least two independent interferometers in the same detector. This is important because different interferometers must be used to achieve optimum sensitivity in the two gravitational wave frequency ranges 100-1000 Hz and 1-10 kHz. These frequencies will assume more significance

when we discuss gravitational wave sources, below. Seismic isolation should permit observing down to at least 100 Hz. In order to achieve sensitivity at frequencies as high as 10 kHz, we shall have to record the output at least 20,000 times per second. Over one year the detector will produce more than 1 Tbyte (10^{12} bytes) of data. This will have to be stored and analysed, a subject which I will discuss in the next section.

The cost of such a detector is determined largely by the cost of the vacuum system. The optics, lasers, and computers are a small fraction of the overall cost. For the 3-km detector the cost will be somewhere between £25 and £30M, depending on the site (which has not yet been decided). In the UK, planning permission has been obtained for building a detector in the Tentsmuir Forest in Fife. Sites in Germany are under investigation.

Handling the large volume of data

As we will see below, it is difficult now to predict just how often a detector of a given sensitivity would see gravitational

waves, but it is certain that most of the data that we take will be pure noise. Occasionally, once an hour if we are lucky, once a month if we are not, an identifiable gravitational wave will come in. Some of the most interesting gravitational waves will be below the usual noise level and will only be extractable by using computers to do sophisticated pattern matching (matched filtering). Given the rate at which data are taken, the number of patterns that have to be looked for (about 1000), and the length of each pattern (up to 2 seconds), it appears that a very fast minicomputer will be sufficient to perform the analysis in real time, so that gravitational wave events can be recognised immediately. A more difficult problem is that of searching for long wave-trains of gravitational waves, such as might originate from pulsars (see below). The sensitivity of such a search will be limited by computing power: even the fastest supercomputer could not fully analyse the data that a detector would take during a 2-week search.

Fig. 6




Fig. 7a




Fig. 7. Artists' impressions of some of the possible sources of gravitational radiation. (a) Pulsar. (b) Supernova. (c) X-ray binary. (d) Coalescing binary.

light will arrive at different detectors at different times, and the time-delays among the various detectors that observe a given event can be used to triangulate its position. The time delays between proposed detectors are typically as long as 20-30 ms, and can easily be resolved by detectors taking data at 20 kHz. Three interferometric detectors provide enough information to reconstruct the wave completely: its amplitude, polarisation, and direction of travel. Doing gravitational wave astronomy therefore demands a network of at least three and preferably more detectors around the world.

In Cardiff our responsibility is to coordinate the development of data exchange and analysis by this planned network. Many problems of signal analysis in networks are poorly understood at present, and the practical and political problems involved in data sharing have to be solved before the detectors go on-line.

Likely Sources of Gravitational Waves

We turn now to the sources of these waves. Without some confidence that there are waves of sufficient strength to be detected, no one would spend large sums of money on interferometers.

The worldwide network of detectors

The proposed detectors are not designed to operate alone: they would coordinate their observing as a single network. There are two important reasons for this. The first is that gravitational wave events will be so rare that one would not believe a detection if it occurred in only one detector: one can never exclude the possibility that some poorly understood local noise source has masqueraded as a gravitational wave. But if two detectors in different parts of the Earth see the same wave, then one can be confident that it is real. Even so, it is possible that random noise will occasionally look like a real event. The only way to avoid this is to set a threshold: to accept events as real only if they are so large that they would be expected to occur by chance less than, say, once per year. If further detectors are added to a network, so that events are accepted only if they register in three or four detectors, then the threshold can be lowered, since chance coincidences among several detectors will be much rarer. This allows waves to be detected from further away in the Universe, and so will lead to many more events being detected. Every new detector that is added to the network enhances the performance of all the existing ones.

The second reason for operating as a

network is that detectors are essentially omnidirectional, and individual detections cannot provide any information about the direction of travel of the wave or its intrinsic amplitude. But a gravitational wave travelling at the speed of

Fig. 7b





Fig. 7c



While there are many uncertainties, it is fair to say that at $h \approx 10^{-22}$ it is very likely that there will be hundreds of detectable gravitational wave events in a global network per year. What is more important than the sheer number of events will be the sort of information they will give us about some of the most puzzling and exotic objects in astronomy.


Pulsars, Supernovae, and Black Holes

It will help to begin by describing the relativistic objects that we expect to be involved in the generation of strong gravitational waves.

Pulsars and supernovae are stars that emit most of their radiation in one or two beams that sweep across the sky as their star spins. If the Earth is in one of the beams, we will see the star apparently turn itself on and off once per rotation, just like a lighthouse. (see figure 8). What makes pulsars extraordinary is their rate of rotation. Typical pulsars rotate up to 10 times per second, and the fastest reported one is spinning at nearly 2000 times per second! Only extremely compact stars can hold themselves together at such rotation rates, and the only known type of star that is compact enough is the neutron star.

A *neutron star* contains the mass of the Sun compressed into a ball that is about 10km in radius, roughly the size of a small city. Its small size gives it an intense gravitational field that is capable of holding it together against the centrifugal effects of such rapid rotation. Neutron stars are formed by the collapse of the core of a massive star that has burnt up all its available nuclear fuel and can no longer resist the inward pull of gravity. Such a collapse usually gives rise to a *supernova explosion*, in which some of the collapsing matter bounces off the newly-formed neutron star and

Fig. 7d

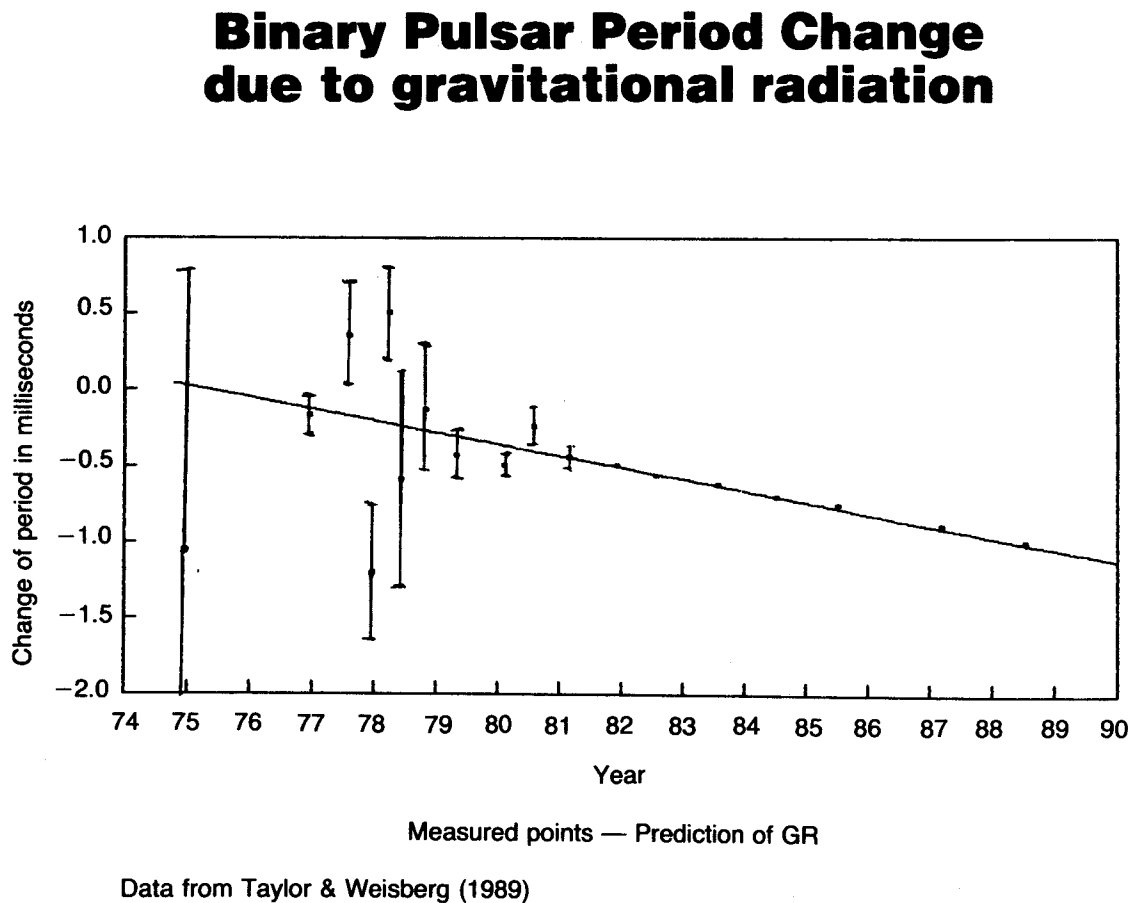


drives the outer layer of the star away in a spectacular explosion. (see figure 8). The rapid rotation of neutron stars is not a big surprise: most stars rotate slowly, but just as a skater spins faster by pulling her arms in, so too will a neutron star speed up its rotation dramatically during the collapse.

Supernovae are observed often in distant galaxies, but the supernova that was observed in February 1987 in the Large Magellanic Cloud, a small satellite galaxy of our Galaxy, the Milky Way, was the first one visible to the naked eye in hundreds of years. Observations of other galaxies suggest that there is

Fig. 8. The orbital period of the Binary Pulsar PSR1913+16. Its decrease is consistent with the predictions of general relativity (solid line).

Fig. 8



probably one supernova every 30 years or so in our Galaxy. Most of these are presumably hidden from view, perhaps shrouded in clouds of dust.

Binary pulsars are for our purposes, the most interesting since they are those that members of binary star systems. Most stars live in binary systems where two stars orbit each other under their mutual gravitational attraction. So it is not surprising to find pulsars in binaries. Among known pulsars, however, those in binaries are relatively rare: a dozen or so out of 450. There are two factors responsible for this: first, the supernova that forms the pulsar tends to disrupt the binary, sending the pulsar moving off through space at a good speed; and second, it is much harder for astronomers to find pulsars in binary systems, since the changing velocity in the orbit produces a varying Doppler shift of the pulse period that tends to wash it out in pulsar searches, where one looks for regularly repeating pulses over a long period of time.

In all known binary pulsar systems it seems that the companion star is compact: either another neutron star or

a *white dwarf*. (A white dwarf is a star with the mass of the Sun but the radius of the Earth: not as compact as a neutron star but still very small. Stars of small mass, like our Sun, tend to evolve into white dwarfs, without spectacular explosions). The first pulsar discovered in a binary system was³ PSR1913+16, which is still known as *the* Binary Pulsar. In many respects this is still the most singular and important system yet discovered (Hulse and Taylor 1975; Taylor and Weisburg 1989).

The companion of the Binary Pulsar is another neutron star, which is not a pulsar (at least, not beaming in our direction), and the stars are so close that the corrections that general relativity makes to Newtonian orbits are appreciable. One of these so-called post Newtonian features of orbits in general relativity is the *orbital precession*: an orbit that would be an ellipse in Newtonian theory remains basically elliptical, but the orientation of the ellipse rotates with time. This precession has been measured for the planet Mercury, where it is a tiny effect: some 32 seconds of arc rotation of the ellipse per century.

But in the Binary Pulsar, where the stars are so close together that they orbit once every 8 hours, the rotation is 4.5 degrees per year, so it can be measured accurately and easily.

By observing the orbital precession and other post-Newtonian effects, one can infer the individual masses of the stars, something that could not be done from observations of a strictly Newtonian system. For the Binary Pulsar, both stars turn out to have mass $1.4 M_{\odot}$. This is a significant number, since the Nobel-Prize-winning astrophysicist S. Chandrasekhar showed long ago that it is an upper bound on the mass of a white dwarf. This measurement is consistent with the theoretical prediction that the burnt-out cores of the massive stars that form neutron stars are incipient white dwarfs, and they only collapse to neutron stars once their mass has grown so large that they can no longer remain as white dwarfs.

The Binary Pulsar is important to us because its orbital motion gives off a considerable amount of gravitational radiation. The radiation has a frequency that is twice the orbital frequency: one

cycle per four hours. This is too low a frequency for the radiation to be observable from any ground-based instruments: seismic and other vibrational noise would mask it⁵. But the energy carried away by the waves has a significant effect on the orbit. It acts as a kind of friction, drawing the stars closer and closer together. As they spiral together, they move faster and faster, and the orbital period decreases. This change of period is measurable in this system, and it can be compared with the prediction that general relativity would make on the basis of the masses of the stars and the properties of the orbit deduced from the post-Newtonian effects. It agrees with the measurement of the gravitational radiation effects to within the observational uncertainty of one percent. (Taylor and Weisberg 1989). This is a strong indirect indication that general relativity correctly predicts the properties of gravitational radiation. This is naturally a great boost to our confidence in building these detectors.

X-ray binaries. Neutron stars in binaries do not always reveal themselves as pulsars. Sometimes they are intense sources of X-rays. If they orbit an ordinary star like the Sun, they may pull matter off the star and onto their own surface. As the matter falls toward the star, it tends to spin into a disc in orbit about the neutron star, (see figure 8). This disc heats up through friction at the very rapid orbital velocities it must reach near the star's surface, and this heating leads to the emission of X-rays. Many, if not all, of today's binary pulsars were once X-ray binaries.

Observations of X-ray binaries usually suggest that the compact object onto which the matter is falling has a mass near $1 M_{\odot}$, which is consistent with the object being a neutron star. But in a few systems, the mass comes out to be nearer $10 M_{\odot}$, and we know from theoretical studies that neutron stars cannot have a mass larger than about $3 M_{\odot}$. These massive objects at the centres of X-ray discs must be black holes.

Black holes are the final resting place of massive cores that cannot halt their collapse at the neutron star stage. The core contracts into such a small volume that the gravitational field is strong enough to trap even light: since nothing can travel faster than light, nothing can get out once it falls in, and the hole is truly black. The massive X-ray binary systems suggest that in at least some supernovae, either the core that collapses is larger than the maximum

neutron star mass, or, more likely, the rebound from the initial collapse is not strong enough to drive off the outer envelope of the star, which then collapses onto the neutron star core, pushing it into a black hole. The radius of a black hole is proportional to its total mass; for a $1 M_{\odot}$ black hole it is only 3km.

Any black holes formed by supernovae in isolated stars are not likely to be easy for us to find. They are just too small to have any observable effect on anything else. But if a black hole is formed in a binary system, then it may evolve into an X-ray source, just as for a neutron star. Although the gravitational field of a black hole is extraordinarily strong inside the hole, gravity well outside it is the same as gravity outside a neutron star. So the matter falling onto the hole will form a hot X-ray disc, and we will see the system as an X-ray source. The


Fig. 9. The form of the expected signal from the two seconds immediately before two neutron stars coalesce. For clarity, the frequency of the waves has been artificially reduced by a factor of 20.

X-rays come from matter that has not yet crossed into the hole; after it enters the hole, no radiation it emits will get out. We don't have much idea about how often supernovae lead to black holes, but about 1% of binary X-ray sources contain black holes rather than neutron stars; it is not unreasonable to guess that this would be true for isolated stars as well.

Black holes probably form in other circumstances as well. There is now a consensus among astronomers that the very bright, distant objects called *quasars* are powered by matter falling toward and into giant black holes in the centre of galaxies; these may range in mass from 10^4 – $10^9 M_{\odot}$. There is strong evidence that even our own Galaxy has a black hole in its centre.

This is where we shall leave our survey of exotic objects; now we turn to the variety of roles that such objects play in sources of gravitational waves.

Fig. 9



A Catalogue of Sources

Supernovae. Detecting Supernovae has historically been the primary goal of gravitational wave detector development. From what little we know about what goes on inside a supernova, we believe it will produce a broadband burst of radiation lasting only about a millisecond, with a typical frequency of 1 kHz. It could contain a sizeable fraction of the total mass-energy of the star, perhaps as much as 1% of Mc^2 when a neutron star of mass M is formed. Such a burst could be detected by the planned interferometers as far away as 40 Mpc.⁶

Coalescing binaries. The "Binary Pulsar" system described earlier is gradually shrinking. The two neutron stars will collide and coalesce in about 10^8 years. Just before they coalesce, they will be orbiting very close to one another (separated by 50-150km) and moving very fast, 50 or more orbits per second. (see figure 8). At this stage, the gravitational radiation emitted by the orbital motion carries away huge amounts of energy $5 \times 10^{-3} Mc^2$, comparable to a decent supernova burst. Because this energy is spread out over many cycles, the waves' amplitude is smaller than one would expect from a supernova and may not be visible above the detector noise. But it can be dug out of the noise by pattern matching, because we have very reliable calculations of what the radiation should look like.

The loss of so much energy to gravitational radiation causes the orbital shrinking to accelerate, so that the stars spend only about 2 seconds between the time their radiation reaches 100 Hz (the likely lower frequency limit of our detectors) and the time they coalesce. During this time, however, they may orbit one another hundreds of times. The fact that we expect wave trains with so many cycles allows pattern-matching to be particularly effective, improving the detectability of such signals by factors of 20-30.

When we do detailed calculations of the expected amplitude of the waves and the rate at which their frequency

changes due to the shrinking of the binary system, we come upon a major surprise: by observing waves from a coalescing binary, one can directly measure how far away it is, without knowing anything about the exact masses of the stars in the system. This is a rare circumstance in astronomy, where it is usually easy to measure the apparent brightness of an object but hard to know what its intrinsic brightness is, from which its distance could be deduced. In fact, at the distances that we expect to see coalescing binaries (out to 650 Mpc, ten times further than we expect to see supernovae) our conventional distance estimates are uncertain by a factor of two. Gravitational wave observations of coalescing binaries could pin down the cosmic distance scale to uncertainties of much less than 10%. We will return to this below.

Given the importance of such observations, what do we know about how often we might observe such events? After all, we will have to wait some 10^8 years to see the Binary Pulsar coalesce! Fortunately, the volume of space out to 650 Mpc contains tens of millions of galaxies, and we believe that our Galaxy contains hundreds of systems like the Binary Pulsar. Most should be invisible to us, either because they are too far away within the Galaxy to be seen (the Binary Pulsar is relatively near to us, and would not be detectable were it much further away), or because they contain pulsars whose beams do not sweep across the Earth. We do see one other system very similar to the Binary Pulsar, called PSR 2127+11C, and searches now underway may reveal more. Based on these considerations, we expect to detect 50-100 events per year in a global network of detectors. But this prediction is very uncertain, and we could see many fewer, or many more.

Not all binaries that coalesce will consist of two neutron stars. As we saw above, perhaps one percent of coalescing binaries will contain a black hole of $10M_{\odot}$ or more; it is possible that a further one percent of these might consist of two such black holes. The signal from the two-black-hole system would be some 6 times larger than that of two neutron stars at the same distance, so the range of the interferometer network for binary black holes would be more than 4Gpc (4×10^9 pc), approaching a cosmological redshift of 0.5. Signals from such systems will have been travelling for half the age of the Universe, and they will consequently tell us something about the Universe when it was much younger.

The formation of the giant black holes in galactic centres probably generates large amounts of gravitational radiation, but owing to the larger size of the black holes the frequency of the radiation is lower, below 1 Hz. Only gravitational wave detectors based in space could detect these events.

Pulsars and other continuous-wave sources. There are many possible long-lived or continuous sources of gravitational radiation in the frequency range accessible to the proposed detectors. These include pulsars with "lumps" in their crust; unstable pulsars spinning down after having been formed with too large an angular velocity; and neutron stars in X-ray binaries where an inherent instability is driven to radiate gravitational waves by the matter that is falling onto the star.

Continuous-wave sources are among the most interesting but least well-understood potential sources of gravitational waves. Although interferometers are broadband detectors, it is possible to configure their optics to tune them to a narrow range of frequencies, with a correspondingly higher sensitivity inside the narrow bandwidth. We expect to use these techniques to search for specific sources, such as pulsars of known frequency, but it appears that the sensitivity of searches aimed at discovering previously *unknown* pulsars (which might be too old to pulse any more, or which might be beamed in a different direction) will in the end be limited by the *computing power* we can bring to bear on the analysis of the data.

Because the sensitivity of an observation is limited by random detector noise, the longer we observe a continuous source, the better. Our sensitivity improves as the square root of the observing time. In principle, we could imagine running an experiment lasting up to 3 or 4 months of continuous observing, and such runs are planned to look for radiation from nearby known pulsars. But to do an all-sky search for unknown pulsars even the fastest supercomputer working today could analyze no more than two weeks' worth of data. Moreover, the computing power required increases as the *eighth* power of the sensitivity of a search, so that computing

⁶Pulsars are named by astronomers according to their coordinates on the sky. Thus, PSR1913+16 has right ascension 19 hours 13 minutes and declination +16 degrees.

⁷The symbol M_{\odot} stands for the mass of the Sun, 2×10^{30} kg.

⁸NASA has funded design studies for interferometric gravitational wave detectors in space that would be able to detect this radiation. Such detectors might well be launched early in the next century.

⁹The astronomer's usual distance unit is the parsec, abbreviated pc, which is about 3 light years or 10^{16} m. Galaxies have a typical size of ten or so kiloparsecs (10 kpc), and the typical distance between galaxies is about a megaparsec (Mpc). Within a distance of 40 Mpc there are thousands of galaxies and hundreds of supernovae per year.

power has to increase by 10,000 in order to improve our sensitivity by a factor of three!

Random background radiation. In addition to predictable, discrete sources of gravitational waves, it is possible that there is a "noisy" background of radiation at some level. This might come from thousands of weak sources whose radiation arrives at the detector at random, or it could come from epochs in the early life of the Universe that leave behind chaotic, random radiation. Some of the predictions concerning such radiation involve processes that lead to the formation of galaxies, so detecting the random radiation would give us some information about how galaxies formed, something we know very little about at present.

Searches for this radiation require two detectors. If there is a strong correlation between the noisy outputs of the two detectors, it is likely to be due to random radiation exciting both detectors in the same way, rather than to the (independent) sources of detector noise. The detectors must not be too far apart, since they must respond to the same random waves at the same time. A baseline between two detectors within Europe is much better than a trans-Atlantic baseline, and this is one of the reasons why we would like to see two detectors within Europe.

Unpredicted sources. As with the opening of any other window in astronomy one can be confident that there will be unexpected sources of gravitational waves at some level. If they are strong enough to stand out above the broadband noise, then they will be readily detected and studied. If they are weaker but have some structure, such as the coalescing binary signal, then they may still be found by looking for correlations between the outputs of two or more detectors. This is certainly one of the most exciting possibilities of gravitational wave detection: learning things that we could not even have imagined beforehand.

Information from Gravitational Waves

We saw earlier that the Binary Pulsar system provides strong evidence that gravitational waves exist and have the strength predicted by Einstein. Why, then, should we build expensive detectors to find them directly? The reason, of course, is the information that the waves

will carry to us about astronomical events.

Tests of general relativity

It is never good enough in physics to rely on indirect observations where direct tests are possible. Something might be wrong with our interpretation of the Binary Pulsar system, so that it does not really test general relativity at all. Seeing radiation effects in the Binary Pulsar only whets our appetite for the real thing.

Apart from verifying that gravitational waves really do exist, direct detections can perform three new tests of relativistic gravity:

Test of gravitational wave polarisation. Detectors are, as we have seen above, linearly polarised. If four detectors see the same event, there is redundant information among them that can be used to test whether the wave's intrinsic polarisation is consistent with the predictions of Einstein. This is a sensitive discriminator among different relativistic theories of gravity: they will all predict gravitational radiation, but with different polarisation properties.

Speed of propagation of gravitational waves. When a supernova occurs, the gravitational waves are emitted within a fraction of a second of the collapse of the core of the star. The optical brightening of the star occurs up to a day later, because the explosion has to travel relatively slowly through the envelope of the star to its surface before we can see it. If gravitational waves travel at the same speed as the light after they leave the supernova, then they will arrive within one day of each other. Over a distance of 15 Mpc (the distance to the nearest large concentration of galaxies, the Virgo Cluster), a difference in speed of only one part in 10^{10} would change the relative arrival times by more than a day. Optical and gravitational wave observations of the same supernova therefore provide a sensitive measurement of the speed of gravitational waves.

Test of strong-field gravity. A further test can be made if black hole coalescing binaries are detected. Computer simulation should soon be accurate enough to make detailed predictions of the dynamics of the merger of the holes, and of the radiation they emit, with only a few free parameters (such as the masses, spins, total angular momentum, and impact parameter of the collision). Given a reasonable signal-to-noise ratio, matching the observations to the predictions could provide a stringent test of strong-field gravity.

Gravitational wave sources

Besides testing general relativity, observations of gravitational waves will provide crucial information about their sources that could be obtained in no other way. Some of the things we expect to learn about are:

Hubble's constant: the rate of expansion of the Universe. The Universe is expanding. The expansion is remarkably uniform, so that galaxies recede from us on average at a speed that is proportional to their distance from us. This is exactly what is required if all matter was at one time concentrated at a single point, from which it exploded: the Big Bang. The constant of proportionality between speed and distance is called Hubble's constant, after Erwin Hubble, the great astronomer who discovered the expansion. Unfortunately, its value is still highly uncertain today: different astronomers claim measurements that differ by a factor of 2. Measuring Hubble's constant is equivalent to measuring the age of the Universe, since we cannot know how long it has taken a galaxy to reach a certain distance away from us unless we know the speed it has at that distance. Measuring Hubble's constant is therefore regarded as one of the most important outstanding problems in astronomy.

The reason Hubble's constant is uncertain is that it is hard to measure distances to astronomical objects. We can measure velocities of expansion by the redshift of spectral lines, and we can even measure *relative* distances fairly reliably by comparing the apparent brightness of similar objects (stars, clusters, or galaxies) in different places, but to measure an *absolute* distance reliably has proved very difficult, because it is usually difficult to know the intrinsic brightness of any object. This is where coalescing binaries become important: as we noted above, observations of gravitational waves from coalescing binaries will give us their distance directly. If the coalescing binary is not too far away, say, 100 Mpc, then it should prove possible either to (i) measure the recessional velocity of the galaxy that contained it, or (ii) to use a statistical method based on the fact that most stars will be in a relatively few clusters near the measured position of the source; either method will determine Hubble's constant and the age of the Universe quite accurately with only a handful of observations (Schutz 1986). A number of alternative ways of measuring Hubble's constant have been

proposed, and the launch of the Hubble Space Telescope in 1990 may allow some of them to be used, but none offers the simplicity and accuracy of the coalescing binary method.

Neutron star nuclear physics. Neutron stars are huge laboratories in which nuclear physics plays a crucial role in determining the structure of the stars. The interactions of neutrons in neutron stars are poorly understood and inaccessible to laboratory experiments, yet they are important to understanding ordinary nuclei. Gravitational wave observations of supernovae, coalescing binaries, pulsars, and X-ray binary gravitational wave beacons can provide information that will improve our understanding of the nuclear physics of everyday matter.

Inner secrets of the supernova. Observations of bursts of gravitational waves from gravitational collapses tell us a number of things about supernovae themselves. We could learn how many collapses do not produce visible supernovae; how often rotation plays an important role in the collapse; whether the collapse has formed a neutron star or a black hole; and what the mass and angular momentum of the compact object are. What is more, by alerting optical astronomers to the occurrence of a supernova within an hour of its detection, we will obtain for the first time observations of a supernova as it first brightens.

Cosmological mass distribution. Given a reasonable event rate, coalescing binaries are good tracers of the distribution of stars out to 500 Mpc or (for black holes) a few Gpc. Their distribution would indicate what inhomogeneities the Universe might have on distance scales of 100 Mpc or so, a scale on which we have no information at present. Discovery of a non-uniform distribution of galaxies would challenge present theories of how galaxies formed, and it would give us insight into a very early phase of the evolution of the Universe.

The early Universe. By confirming or ruling out a background of random gravitational radiation at some intensity level, gravitational wave observations will further probe the early Universe and the nuclear physics theories that predict such backgrounds.

These are a few, fairly predictable, payoffs should our planned detectors successfully observe gravitational waves. They could be supplemented by more speculative things, such as possible gravitational lensing of black

hole coalescing binary events or the measurement of the rate of deceleration of the expansion of the Universe, which is equivalent to measuring how much mass is contained in any volume of the Universe. They give a flavour of reasons why scientists in many countries are persuading their funding agencies to support these projects.

The Prospects

Work on gravitational wave detection presses on in many countries. Prototype interferometers continue to be improved in several laboratories; in many others, research into the optical techniques and the laser technology that will be needed for large-scale detectors is making encouraging progress. Proposals for full-scale interferometers have been submitted in six countries, and decisions are awaited in 1990. While it is impossible to predict the outcome of any of these decisions, especially in view of the uncertainty of science funding in some of them (notably Britain), it is significant that in several countries, including Britain and the USA, the scientific funding agencies have indicated that gravitational wave interferometers are at the top of their list of priorities for major new pure-science experiments.

Once a project is approved, construction of the kilometre-scale vacuum system could begin within a year, and could take three years. There could follow another two years or so of work on the optics before the instruments could begin to function as gravitational wave observatories, with sensitivity somewhere near $h \approx 10^{-21}$. To reach 10^{-22} , at which point coalescing binaries become detectable, requires an unpredictable amount of development, and could take several more years. No one has built instruments capable of measuring such small displacements before, and so no one can predict what obstacles may lie in the way of attaining a given sensitivity. Nevertheless, progress on the prototypes, which are themselves the most sensitive instruments for measuring distance ever built, has been remarkably steady, with an improvement in sensitivity by a factor of 10 every two years. The leap from 30m prototype to 3km observatory automatically brings an improvement of a factor of 100; the further factor of 100 required to reach 10^{-22} from today's sensitivity may therefore not take quite as long as we may fear. (Bear in mind that these are the words of a theorist: my experimental

friends, who will actually have to do the work, don't agree!)

Meanwhile, bar detectors are quite capable of observing supernovae in our own Galaxy, and at least three of them, at Stanford University, at Louisiana State University and at CERN (operated by Rome University) are, or soon will be, running full-time, ready to register the next nearby gravitational collapse. It would be justice of a sort if the discovery of gravitational waves were to be made by the venerable bars, rather than by the new generation of interferometers. But whether the Galaxy obliges this sense of justice will depend on how generous it is with supernovae. One thing seems worth betting on: we should see the birth of gravitational wave astronomy before the close of this century.

References

- Blair, D., Frasca, S., and Pizzella, G. 1989 *Gravitational Wave Data Analysis* (ed. Schutz, B.F.). Kluwer, Dordrecht, p.285.
Caves, C.M., Thorne, K.S., Drever, R.W.P., Sandberg, V.D., and Zimmerman, M. 1980, *Rev. Mod. Phys.* 52, 341.
Drever, R.W.P. 1983 *Gravitational Radiation* (ed. Demille, N. and Piran, T.) North-Holland, Amsterdam, p. 321.
Hulse, R.A., and Taylor, J.H. 1975 *Astrophys. J.* 195, L51.
Jeffries, A.D., Saulson, P.R., Spero, R.E. and Zucker, M.E. 1987 *Scientific American*, June, p.50.
Leuchs, G., 1988 *Contemp. Phys.* 29, 299.
Meers, B.J., 1988 *Phys. Rev. D*, 38, 2317.
Schutz, B.F., 1985 *A First Course in General Relativity*, Cambridge University Press, Cambridge.
Schutz, B.F. 1986 *Nature*, 323, 310.
Schutz, B.F. 1989 *Class. Quant. Grav.* (in press).
Taylor, J.H., and Weisberg, J.M. 1989 *Astrophys. J.* 345, 434.
Thorne, K.S. (eds. Hawking, S.W., and Israel, W. 300 Years of Gravitation, Cambridge University Press, Cambridge, p.330.
Weber, J., 1967 *Phys. Rev. Lett.* 18, 498.
Weber J. 1960 *Phys. Rev.* 117, 306.
Will, C.M. 1986 *Was Einstein Right?*, Basic Books, New York, N.Y.