Class. Quantum Grav. **13** (1996) A219–A238. Printed in the UK

Gravitational-wave sources

B F Schutz

Department of Physics and Astronomy, University of Wales College of Cardiff, Cardiff, UK
and
Max Planck Institute for Gravitational Physics, The Albert Einstein Institute, Potsdam, Germany

Abstract. The sources of gravitational radiation that seem to be most likely to be detected by planned ground- and space-based gravitational-wave detectors are reviewed, with particular attention to the likely information that observations of them will reveal. A particular emphasis is placed on observations by LISA, the fundamental-physics Cornerstone project of the European Space Agency's Horizon 2000+ programme.

PACS numbers: 0430D, 0480, 0480N, 9555Y, 9555

1. Introduction

In this paper I will survey sources of gravitational radiation that seem likely to be detected over the coming decades by both ground- and space-based detectors. I will particularly emphasize the sources that may be seen by LISA, the European Space Agency's project to place a large-scale laser interferometer in space as a Cornerstone mission in its Horizon 2000+ programme. This detector is described elsewhere in this volume.

I will begin with a brief introduction to gravitational waves, emphasizing the information that they carry about their sources. Then I look at the whole range of sources in a new way, showing how a rough estimate of the natural timescales associated with a source of a given mass and size leads to it being a candidate for either space-based (low-frequency) or ground-based (high-frequency) detection. I then discuss specific sources, and in particular the kind of information we can expect to learn about them from observations.

2. Gravitational waves

Gravitational waves that come to us from astronomical bodies are, by the time they reach us, weak and plane-fronted. General relativity tells us that the interaction of such a wave with a detector is described entirely by the strain it would induce in an array of freely-moving particles in empty space. The effect of a wave is to change distances between such particles. The strain associated with the change $\delta\ell$ of the distance ℓ between two particles is the relative distance change, which (to within a factor of 2) is equal to the amplitude h of the gravitational wave that induces the strain:

$$h = 2\frac{\delta\ell}{\ell}. \quad (1)$$

For the interaction to be as simple as this, the two free masses must be much closer to one another than the wavelength of the gravitational wave, so that the distance (which is a proper distance in the language of relativity) can be determined on a timescale which is short compared to that on which the wave amplitude changes. This condition is always satisfied for ground-based detectors, but not always in space. LISA, for example, is large enough that at its high-frequency limit of sensitivity its reaction to a gravitational wave is more complicated than this, and gravitational-wave detection schemes based on ranging to interplanetary spacecraft are usually at the opposite limit: the distances between the spacecraft and the Earth (the two free masses in such experiments) are much larger than the wavelengths they are sensitive to. Solid-mass detectors, such as the bar detectors that made the first serious attempts to detect gravitational waves, also have a more complex response, since they are not made of free particles but rather of atoms held tightly together by interatomic forces.

I will not consider these complications here, because they make signal extraction from such detectors more complicated but they do not change the physics of the sources. The effect of the wave on two particles depends on their orientation in space. The effect of a gravitational wave on an arbitrarily oriented array of particles is, in fact, characterized by only *two* time-dependent numbers, the amplitudes of its two *polarization* components, which are conventionally called h_+ and $h_\times$. The polarization diagram, figure 1, is similar to ones that can be found in most standard textbooks.

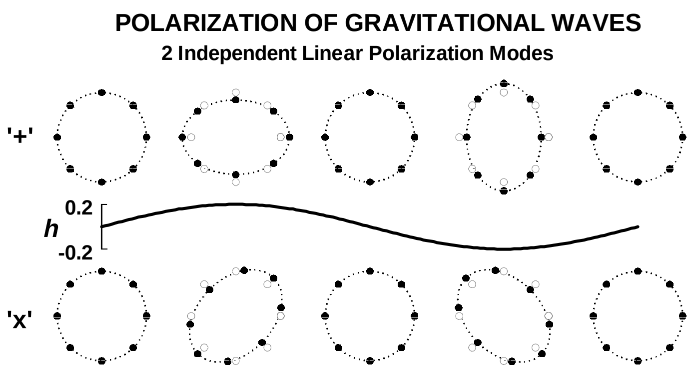

Figure 1. Illustration of the polarization of a gravitational wave. Two linearly independent polarizations of a gravitational wave are illustrated by displaying their effect on a ring of free particles arrayed in a plane perpendicular to the direction of the wave. The wave-form itself is drawn between the two sequences, for a wave with the (large) dimensionless amplitude $h = 0.2$. Shown to scale are the distortions in the original circle that the wave produces if it carries the $+$-polarization (above) and the $\times$-polarization (below). The motion of each particle can be discovered by comparing it to its original position, shown as the 'shadow' circles. In general relativity, these are the only two independent polarizations. They are transverse to the direction of the wave.

As we will see below, the polarization of the gravitational wave contains important information about the source, so we generally want to measure it. All practical detectors are linearly polarized, responding to a linear combination of h_+ and $h_\times$. If an incoming wave is a short burst of radiation, then a single detector cannot measure its polarization or determine its direction. For this reason, ground-based detectors looking for supernova bursts or coalescences of binary stars will work in networks. Four interferometers, for example, are sufficient to determine the polarization, direction and amplitude of any wave they detect.

If the incoming wave lasts long enough for the detector to accelerate appreciably during the observation, such as for a continuous signal from a pulsar or binary star system, then a single detector can do much more. Any Earth-bound detector or one in space will accelerate. If the detector changes its orientation, then the linear combination and projection of polarizations will change; this induces an *amplitude modulation* into the detector output. Similarly, the acceleration of the detector towards the source will produce a *phase modulation* (Doppler shift) of the incoming signal's phase, which will be detectable if the detector changes its position significantly (relative to the gravitational wavelength) as it accelerates. This complicates the extraction of the signal from the data (a simple Fourier transform will not do), but once the modulations are determined then the direction to the source and its polarization can be inferred from the output of a single detector. This will be important for LISA, since it will be operating on its own, and it will also probably be important for the first high-sensitivity ground-based interferometer, which may not have a network of detectors to operate with at first. This is dealt with at length in other references [1, 2].

The signal, once detected, contains a variety of information. Some of it will depend on the specific physics of the source, but some will be interesting in general. Here are some examples.

Frequency. The frequency of the dominant gravitational radiation from systems where the dynamics is driven by self-gravitation is, to within factors of the order of unity, the natural frequency

$$f_{dyn} = \frac{1}{2\pi}\left(\frac{GM}{R^3}\right)^{1/2} \sim (G\rho)^{1/2}/\pi, \tag{2}$$

where M is the system's mass, R its typical size and ρ its characteristic density. Solving for the density, we get

$$\rho \sim 10^{16}\left[\frac{f}{300\ \mathrm{Hz}}\right]^2\ \mathrm{kg\ m^{-3}}. \tag{3}$$

Notice that, for frequencies within the range of proposed detectors, 10^{-4}–10^{4} Hz, this ranges from nuclear-matter density at the high end down to the density of water at the low end. This illustrates the enormous range of kinds of potential gravitational-wave sources, and shows why there is interest in the low-frequency band that is only accessible from space as well as the high-frequency band that ground-based detectors aim for. We will return to this point later.

Systems whose dynamics are not governed directly by self-gravity need more careful treatment. For example, a rotating, lumpy neutron star will radiate, but its rotation rate may be much less than the dynamical frequency, and it radiates only because its lumps are held up by stresses in the crust.

Rate of change of frequency. If the gravitational-wave frequency of a fairly steady source changes in time, this could be due to the loss of energy and angular momentum to the gravitational waves. This can give further interesting information. In particular, if two stars of masses M_1 and M_2 in a circular binary orbit† emit radiation, then in the quadrupole approximation the amplitude of the waves and hence the energy carried away by them depend on the masses only through what is called the *chirp mass*

$$\mathcal{M} = \mu^{3/5} M^{2/5}, \tag{4}$$

† Circularity is a good assumption for systems emitting at frequencies observable from the ground, because they lose eccentricity as they spiral together. At lower frequencies, these formulae need to be generalized in a straightforward way to include eccentricity.

where μ is the reduced mass $M_1M_2/(M_1 + M_2)$ and M the total mass $M_1 + M_2$ of the binary. The formula is

$$h = 1.5 \times 10^{-19} \left[\frac{f}{1\ \mathrm{Hz}}\right]^{2/3} \left[\frac{\mathcal{M}}{1\ M_\odot}\right]^{5/3} \left[\frac{r}{1\ \mathrm{kpc}}\right]^{-1} \cos\left(\int_0^t f(t')\,\mathrm{d}t' + \phi_0\right), \qquad (5)$$

where ϕ_0 is an arbitrary initial phase, and where the frequency of the signal changes at a rate given by

$$\frac{\mathrm{d}\ln f}{\mathrm{d}t} = 0.126 \left[\frac{\mathcal{M}}{M_\odot}\right]^{5/3} \left[\frac{f}{100\ \mathrm{Hz}}\right]^{8/3}\ \mathrm{s}^{-1}. \qquad (6)$$

This effect has of course been observed in the famous Hulse–Taylor binary pulsar [6]. There are two remarkable deductions that one can make from the two equations above. The first is that, if one can measure the time dependence of the frequency, then one can directly measure the chirp mass $\mathcal{M}$. This contains information about the masses, but not enough to determine each of them uniquely unless it is combined with other information, as discussed in the following. The second deduction is that, if one can measure the time dependence of both the frequency and the amplitude, then one can eliminate $\mathcal{M}$ and deduce r as well [3]. We conclude: *binary systems that change their frequency during an observation because of gravitational radiation reaction are standard candles, systems whose distance can be deduced from their intensity.* We refer to the change in frequency as a *chirp*.

Polarization. It is not always appreciated that the polarization of the waves contains some of the most useful information. This is because it is so closely related to the sense of the mass motions in the source. This is described in figure 2, where we see that the ellipticity

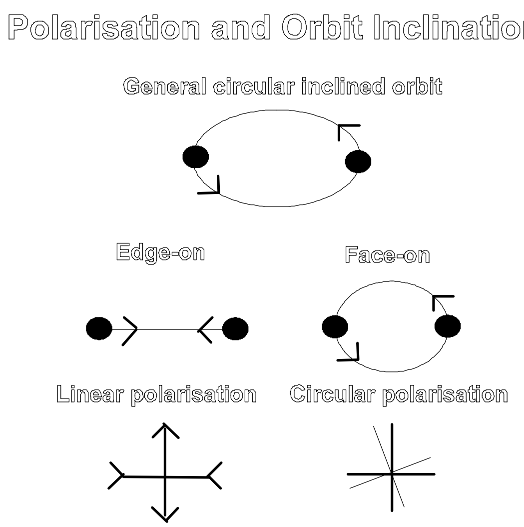

Figure 2. The polarization of a wave tells us directly the inclination of a binary orbit. Any binary system might be seen edge-on (left-hand side) or face-on (right-hand side). The polarization of the wave, or in other words the motions produced by the wave in the detector, follow the apparent source motions projected on the sky. For the system seen edge-on, the stars simply appear to move back and forth along a line, and the polarization is linear in the '+' orientation (bottom left). For the system seen face-on, the stars are seen to move on circular orbits and the polarization is circular too, with the same sense (bottom right).

of the polarization tells us the inclination of the orbit of a binary star source, or alternatively the inclination of the rotation axis of a rotating neutron star source. This is an important quantity that is usually difficult to deduce from optical observations of systems. An optical, x-ray or radio observation of a binary usually only gives us the radial velocity of one of the components, from which one can only deduce the *mass function*, which is a combination of the masses of the stars and of the inclination of the orbit. If the same binary is observed gravitationally, the polarization will allow us to eliminate one of these three unknowns. If, in addition, the binary chirps during the observation, then the chirp mass $\mathcal{M}$ gives a third observable that will allow one to determine the individual masses of the stars. This illustrates the complementarity between electromagnetic and gravitational observations, which is one of the strongest motivations for pursuing gravitational-wave detection.

Amplitude. The amplitude of a gravitational wave is a combination of its intrinsic strength and the distance to the source. If we know the intrinsic strength (as in measuring the chirp rate of a binary) then we can deduce the distance. If we know the distance (say from independent optical observations), then we can deduce the intrinsic strength and therefore something about how relativistic and asymmetric the system is.

Phase. The phase of the signal contains information about the positions of objects within the source. For example, if the source is a binary star, then the phase tells us where the stars are in their orbits and how the orbit changes with time. The frequency and chirp rate (above) are really part of the phase, being just the first and second time derivatives of the phase function. In highly relativistic binaries, the phase function will contain post-Newtonian orbital information that can lead to the determination of the individual stellar masses and possibly even of their spins [4, 5]. The details of the phase affect how well we can predict the signal and therefore how easy it is to detect.

There is not only astrophysical information in gravitational-wave observations, but also information about fundamental aspects of gravitation theory. For example:

Speed of propagation of gravitational waves. If a supernova explosion or other gravitational-wave generating event is detected by other means, such as the optical brightening of the star, then the arrival times of the electromagnetic waves and the gravitational waves can be compared to constrain their relative propagation speeds. If gravitational waves propagate slower than light, this would indicate a mass for the graviton. We already have upper limits on this mass from the fact that Newtonian gravity works well in the solar system and the fact that the Hulse–Taylor pulsar [6] is well described by Einstein's theory using massless gravitons. Both of these constrain the graviton mass to be less than 10^{-20} eV. At this mass, a graviton would arrive 0.1 s later than a photon emitted simultaneously from a distance of 1 kpc. At the distance of the Virgo Cluster, the time delay would be 30 min. At 1 Gpc, the time-lag would be 1 day. Whether observed delays can be interpreted to mean a mass depends on the source. For example, a supernova emits its gravitational waves several hours before it brightens up, so it would not give interesting information unless it were at gigaparsec distances. On the other hand, if a coalescing binary is associated with a gamma-ray burst, then they should be emitted with less than a second between them, and the constraint at 1 Gpc would be a very strong one.

Alternative theories of gravity. Most alternative theories that would fit with standard models of theoretical physics would differ from general relativity by the inclusion of extra fields, and these would be detectable by the fact that they would have different polarization states. Polarization observations by networks of detectors, or long-duration observations by single detectors, could provide evidence of this. The ratio of unexpected polarizations to expected ones would be small, so the detector would need to be sensitive. LISA probably has the best chance of seeing such effects.

Cosmological background. We will discuss the cosmological background below, but it is worth noting here how fundamental an observation it would be. It would have to come from an era well before the time of recombination, when the microwave background was formed, and it would come to us pristine: its statistics and spectrum would give direct information about the physics at the high energies at which it was formed.

3. Dynamics of gravitational-wave sources

The gravitational-wave sources we will consider span eight orders of magnitude in frequency, from 10^{-4} to 10^4 Hz. This is similar to the range from high-frequency radio waves (10 GHz) to x-rays (10^{18} Hz). Over this range, therefore, we should expect considerable variety. But there is also a lot that is systematic. The dynamics of most sources are dominated by their self-gravity, and this means that they will radiate at a frequency f_{gw} which is normally about twice their natural dynamical frequency f_{dyn}, given in equation (2).

This Newtonian formula will not be exact for highly relativistic systems, and there are some cases where gravitational waves come out at other multiples of f_{dyn}, but this formula is nevertheless a good guide to the relation between mass and size that we might expect for a system radiating gravitational waves at any given frequency. The exceptions are rotating neutron stars, which radiate at a multiple of their rotational frequency (which can be much less than the dynamical frequency) and a cosmological stochastic background, which was produced at a frequency that roughly obeyed equation (2), but which has been redshifted by the cosmological expansion to a much lower frequency.

The dynamics of gravitational-wave sources can usefully be described in figure 3, which shows lines of constant frequency f_{gw} in the mass–radius plane for three important frequencies: 10^{-4} Hz, the lowest frequency accessible to LISA; 1 Hz, roughly the boundary between what can be detected from the ground and from space and 10^4 Hz, the upper limit to what can in practice be observed from the ground. The upper part of the diagram is therefore the space-accessible region; the lower part, the domain of ground-based detectors.

In the diagram we place a number of interesting possible gravitational-wave sources. At the low-mass end, the natural vibrations of a typical neutron star and stellar-mass black hole radiate in the ground-based band; these should be excited when the objects are formed. The Sun lies in the space band, and indeed its natural vibrations could be detected by a space detector, through the near-zone Newtonian gravitational oscillations they produce rather than through their gravitational waves. Binaries in this mass range are discussed below. At the high-mass end, a 10^6 $M_\odot$ black hole would radiate in the space band. These vibrations could be excited by the formation of the hole or by a neutron star falling into such a hole.

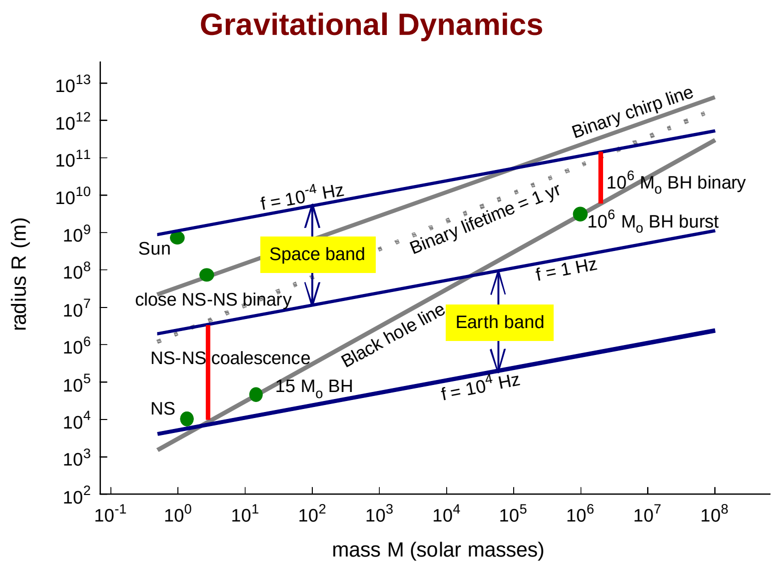

Figure 3. This diagram shows the wide range of masses and radii of sources whose natural dynamical frequency is in the band detectable from space or the ground. The three heavy lines delineate the outer limits of the space band at gravitational-wave frequencies of 0.1 mHz, 1 Hz and 10 kHz. The 'black-hole line' limits possible systems: there are none below it if general relativity is correct. The 'chirp line' shows the upper limit on binary systems whose orbital frequencies change (due to gravitational-wave energy emission) by a measurable amount (30 pHz) in one year: any circular binary of total mass M and orbital separation R that lies below this line will 'chirp' in a one-year observation, allowing its distance to be determined. The curve labelled 'binary lifetime = 1 yr' is the upper limit on binaries that chirp so strongly that they coalesce during a one-year observation. These lines and the indicated sources are discussed more fully in the text.

There are other useful lines in this diagram. The most important is the *black-hole line*, drawn for

$$R = \frac{2GM}{c^2}. \tag{7}$$

The region of the diagram below this line does not contain any physically realizable systems: a system forms a black hole when it reaches this line from above. The space-accessible frequency region contains black holes above about $10^4\ M_\odot$ up to $10^8\ M_\odot$, which means that space detectors can, in principle, confirm the present astrophysical consensus that most galaxies contain one or more giant black holes. The ground-based region is limited to smaller black holes, as we expect to form from normal massive stars.

Two other lines in the diagram refer to the chirping of a binary system, as discussed above. The line called 'binary lifetime $=$ 1 yr' is the line along which the characteristic timescale for the frequency to change, as inferred from equation (6), is one year. Binary systems below this line are systems which can be followed right to coalescence during a reasonable observation period. This is a line on which R^4/M^3 is constant. Notice that all solar-mass binary systems observable from the ground will coalesce within a year. A typical coalescing neutron-star binary is illustrated in the diagram. From space, we can expect only binaries of massive black holes, $M \geqslant 10^6\ M_\odot$, to coalesce during an observation, as shown.

Just as important, but less dramatic, is just seeing a binary system 'chirp', i.e. change its orbital frequency. Here the criterion is not that its coalescence timescale be the observation time, but rather that its frequency should change by an observable amount during the same one-year observation. This means that its frequency change need only be as large as the frequency resolution of a one-year observation, 3×10^{-8} Hz. The resulting ‘chirp line’ is a line of constant R^{11}/M^7. The diagram shows the chirp line appropriate to a one-year observation. It shows that chirping without coalescence is important for space-based detectors; ground-based detectors will be able to follow any chirping system right to coalescence. A space-based detector may well observe solar-mass binaries chirp, as indicated by the ‘close NS–NS binary’ point drawn in the diagram. (Here the word ‘close’ means an orbit that is as compact as one expects to find at any one time in the Galaxy; the stars are not as close to one another as in a coalescing binary, which can be expected to occur in one in every 10^5 galaxies in any year.) A space detector should also detect chirping in binaries consisting of massive black holes. The observations will be able to determine distances for these objects. This will be particularly interesting for LISA, as we describe below.

4. High-frequency sources detectable from the ground

I will discuss the sources detectable by ground- and space-based instruments separately, because in most cases the physics is rather different. But it may be of interest to see their relation to one another, so I have plotted a diagram to show them together, for burst sources of short duration (figure 4).

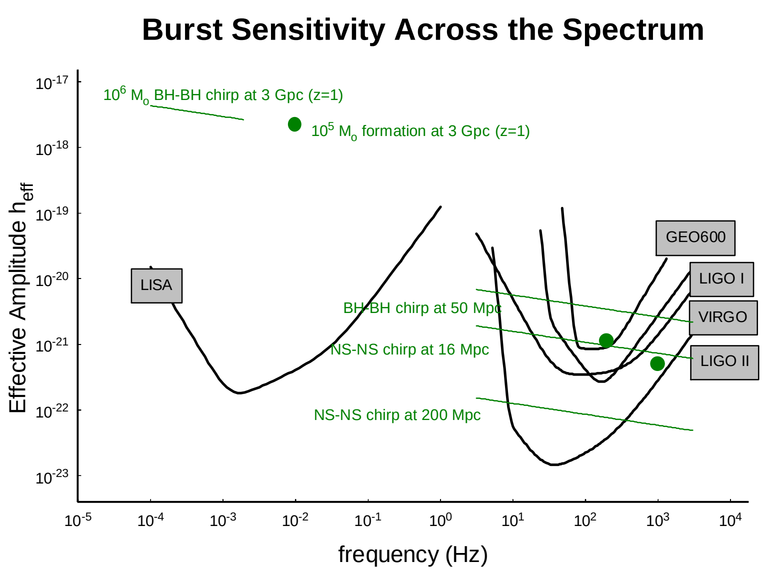

Figure 4. The sensitivity of detectors to bursts at high and low frequencies. Various possible sources are illustrated, along with expected sensitivities of planned detectors. The vertical axis is the effective amplitude of the waves in a broad bandwidth, and the noise limits of detectors are their noise in a bandwidth equal to the observing frequency.

4.1. Supernovae and gravitational collapse

Among the most violent events known to astronomy, supernovae were the driving motivation for the early development of bar detectors. They are still a major target for bar-detector development, and they are of course an important potential source for ground-based interferometers as well, which in their first stages should be able to see such a burst with an amplitude of 10^{-21}, and which can be improved to an amplitude of 10^{-22}, using technology already available today. Supernovae (of type II) are triggered by the gravitational collapse of the interior, degenerate core of an evolved star. The result is thought to be a neutron star or black hole. The collapse releases an enormous amount of energy, at least equal to the binding energy of a neutron star, about 0.15 $M_\odot c^2$. Most of this energy should be carried away by neutrinos, and this is supported by the neutrino observations at the time of supernova SN1987A. One way of calibrating the sensitivity of detectors is to calculate the amplitude of the gravitational wave that would be produced if a certain fraction of the released energy were converted into gravitational waves. For any weak gravitational wave of amplitude h, the energy flux carried by the wave is given by the simple expression

$$F = \frac{c^3}{16\pi G}|\dot{h}|^2. \tag{8}$$

This formula is exact, independent of approximations used in calculating how the waves are generated: it depends only on the fact that the waves are weak when they reach the Earth. If a burst of gravitational waves is emitted by a source that is a distance r from the detector, and if the waves carry a total energy E emitted predominately at a frequency f_{gw} and spread over a timescale τ, then it is easy to use this flux formula to show that the amplitude h will be

$$h = 5\times 10^{-22}\left[\frac{E}{10^{-3}M_\odot c^2}\right]^{1/2}\left[\frac{\tau}{1\ \mathrm{ms}}\right]^{-1/2}\left[\frac{f_{gw}}{1\ \mathrm{kHz}}\right]^{-1}\left[\frac{r}{15\ \mathrm{Mpc}}\right]^{-1}, \tag{9}$$

$$= 3\times 10^{-18}\left[\frac{E}{10^{3}M_\odot c^2}\right]^{1/2}\left[\frac{\tau}{1000\ \mathrm{s}}\right]^{-1/2}\left[\frac{f_{gw}}{1\ \mathrm{mHz}}\right]^{-1}\left[\frac{r}{3\ \mathrm{Gpc}}\right]^{-1}. \tag{10}$$

These equations show that a burst that emits $10^{-3}M_\odot c^2$ in gravitational-wave energy, or in other words less than 1% of the available energy, would be detectable by the first generation of ground-based detectors at the distance of the Virgo Cluster (18 Mpc) provided its energy comes out around 300 Hz. Bar detectors operating at 1 kHz with effective sensitivities of 10^{-21} could see these as well, provided the energy comes out at kHz frequencies. This sensitivity might be achieved by the newly proposed spherical-mass detectors.

If one knows what the structure of the waveform is before one searches for it, then the effective sensitivity one can achieve through pattern matching (matched filtering) is roughly $h\sqrt{n}$, where n is the number of cycles in the waveform. In the equations above, the duration τ is proportional to n: $\tau = n/f_{gw}$. If we substitute this into the equations and solve for the effective amplitude, we find

$$h_{eff} = h\sqrt{n} = 5\times 10^{-22}\left[\frac{E}{10^{-3}M_\odot c^2}\right]^{1/2}\left[\frac{f_{gw}}{1\ \mathrm{kHz}}\right]^{-1/2}\left[\frac{r}{15\ \mathrm{Mpc}}\right]^{-1}, \tag{11}$$

$$= 3\times 10^{-18}\left[\frac{E}{10^{3}\ M_\odot c^2}\right]^{1/2}\left[\frac{f_{gw}}{1\ \mathrm{mHz}}\right]^{-1/2}\left[\frac{r}{3\ \mathrm{Gpc}}\right]^{-1}. \tag{12}$$

These show that the effective sensitivity of a detector really depends more on the energy in the waves than on their duration. It also, of course, depends on the frequency, both because this equation depends on frequency and because the sensitivity of the detectors is a function of frequency.

The second stage of development of ground-based interferometers should make it possible to see such bursts at distances of many tens of Mpc, a volume that contains tens of thousands of galaxies. If only 1% of gravitational collapse events produce bursts of the strength we have assumed, then the first generation should see roughly one event per year. The second should see more than one per day. These sensitivity limits are indicated in figure 4.

Our detectability estimate above used the seemingly arbitrary figure of $10^{-3}M_{\odot}c^2$ for the energy emitted because it is difficult to make reliable predictions of this figure from present theory. Since spherical motions do not emit gravitational waves, the strongest sources are likely to be the most asymmetric. Unfortunately, modern computers are still not able to perform realistic simulations of gravitational collapse in 3D, including all the important nuclear reactions and neutrino and photon transport. It is not clear whether we will have good simulations before we have observations to settle this question.

4.2. Coalescing binaries

Coalescences of compact binaries consisting of either neutron stars or black holes are among the most important potential sources for ground-based interferometers. They may also be detectable in the future by the new spherical solid antennae. What makes them different from and more easily detectable than gravitational collapse events is that the radiated gravitational waveform during the inspiral phase can be modelled very accurately, and most of the energy is emitted at a much lower frequency. Significant gravitational-wave emission is expected from the ‘plunge’ phase after the stars reach the last stable orbit and fall rapidly towards one another, and from the merger event, but both of these phases are as yet poorly understood. Neither of these final phases will be likely to generate more signal than the inspiral phase, so the detectability of such systems rests on tracking their orbital emissions.

It helps to try to observe these systems at the lowest possible frequency. Even though the amplitude is less at low orbital frequencies (larger orbital diameters), the systems dwell at these frequencies very much longer, and this makes the signal easier to recognize.

Efforts to build detectors with good low-frequency sensitivity down to, perhaps, 10 Hz will be particularly helpful in detecting coalescing binaries, since they emit much more of their power at low frequencies. A typical neutron-star binary can be tracked for several minutes if the gravitational-wave signal is initially detected at 40 Hz. Constructing good theoretical waveform templates for such an event will not be easy, because post-Newtonian effects on the orbit will be important [4, 5].

The major uncertainty about coalescing binaries is their event rate: how many of these rare events are to be expected? Estimates of event rates are based on careful studies of pulsar detections in binary systems [7, 8] and on theoretical studies of binary evolution [9, 10]. The two methods give different answers, but they are not necessarily incompatible. Observations of pulsars tell us about the statistics only of systems that last long enough as pulsars to have a reasonable chance of our finding one near enough to us to observe. These systems apparently have a formation rate of about 10^{-6} yr^{-1} in the Galaxy. Studies of evolution, however, suggest that they are only the tip of the iceberg, and that a much larger population of systems with very short orbital periods exists, which live for such a short time that we do not have much probability to observe them. These systems are formed at a rate that may be as large as 10^{-4} yr^{-1} in the Galaxy.

The pulsar-based rate is a lower limit, and it converts into about one event per year out to a distance of 200 Mpc. If the larger rate stands up to examination, it will bring the nearest coalescence in one year in to perhaps 45 Mpc, or three times the distance to the Virgo Cluster.

A certain fraction of systems should contain black holes instead of neutron stars. In fact, since the formation of a black hole is less likely to disrupt a binary system (there is much less mass loss), the fraction of binaries containing a black hole is likely to be much larger than the fraction of gravitational collapse events that result in black-hole formation; the fraction containing a single black hole could be as large as 1 in 3 [8]. Since the radiation amplitude would be about 4 times higher, the event rate for such systems could actually be larger than for neutron-star systems.

First-stage interferometers that are sensitive at the 10^{-21} level down to 100 Hz are unlikely to see binary neutron-star systems even if the higher event rate holds: the effective amplitude is just about at their noise level for events at 50 Mpc. But black-hole–neutron-star binaries are a real possibility for these detectors. Second-stage interferometers sensitive down to, say, 40 Hz should see neutron-star events at 1 Gpc, which could result in hundreds or thousands of observations of such binaries per year.

Observations of coalescing binaries will give systematic information about the neutron-star population in distant galaxies (including their mass distribution) and will demonstrate conclusively whether gamma bursts are associated with coalescence events. If the coalescence event itself is observed, there will be abundant information about the neutron-star equation of state; in particular, the maximum orbital frequency dependence on the equation of state.

Most spectacularly, as we have pointed out earlier, observations will reveal the distance to the binary system. If stage-two detectors see a reasonable number of events within 100 or 200 Mpc, it should be possible to associate events at least statistically with clusters of galaxies. Optical redshift measurements of the clusters, combined with the distance measure, will give another method of determining the Hubble constant [3]. We will see below how similar methods used by a space-based detector could determine the deceleration parameter q_0 of the Universe.

4.3. Gravitational radiation from individual neutron stars

Binary neutron-star systems radiate by virtue of their orbital motion. Individual neutron stars, whether isolated or in binary systems, can also radiate gravitational waves if they spin (which almost all of them do) and if they are somehow significantly non-axisymmetric (which may or may not be common).

The non-axisymmetries may come from irregularities in the crust, perhaps from strains that have built up as the stars have spun down or perhaps from irregularities associated with their formation that became frozen in as the star cooled; or they can be dynamical, such as normal modes of pulsation that are excited in some way, or precession that is driven by the accretion of angular momentum. We will describe these possibilities below. One source of non-axisymmetry that is certain to exist, the off-axis magnetic field of an active pulsar, is unfortunately not large enough to produce significant radiation.

The observed spindown of radio pulsars is presumed to be driven primarily by the emission of a wind of energetic particles and of low-frequency electromagnetic waves from the spinning dipole magnetic moment. The radio waves that we observe carry very little energy, so we have no direct observations of the spindown mechanism. This leaves room for the possibility that gravitational radiation contributes a significant amount to the spindown as well. By using the spindown to place an upper limit on h (essentially from equation (9), where the energy is the kinetic energy of rotation of the pulsar and the timescale τ is the spindown timescale), one finds upper limits on the gravitational radiation from the Crab and Vela pulsars that is of the order of 10^{-24}. The upper limits on h for all the pulsars

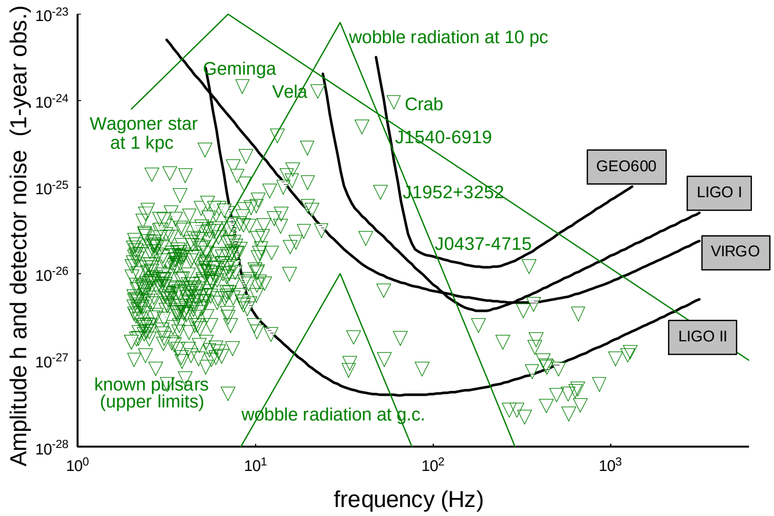

Figure 5. Pulsar spindown upper limits compared with the noise that is expected for the ground-based detectors in a year of observing. Pulsars would only be near these upper limits if a substantial fraction of their loss of rotational energy were accounted for by gravitational radiation. Also shown are upper limits on radiation from Wagoner stars and from accretion-driven precession inside, for example, Thorne–Zytkow objects.

whose gravitational-wave frequency exceeds 7 Hz and for which a spindown rate has been measured are plotted in figure 5.

The energy loss rate is also related to the non-axisymmetric ellipticity $\epsilon = 1 - a_2/a_1$, where a_1 is the semimajor axis of the equatorial section and a_2 the semiminor axis. The formula shows that the ellipticity is given as a function of the spin period P of the pulsar and its rate of change $\dot{P}$ by

$$\epsilon = 0.0057 \left[\frac{P}{1\ \mathrm{s}}\right]^{3/2} \left[\frac{\dot{P}}{10^{-15}}\right]^{1/2}. \qquad (13)$$

For the Crab this gives $\epsilon = 7 \times 10^{-4}$, again assuming that $\dot{P}$ is produced by gravitational waves. While this is probably sustainable by the crust of the star, the true ellipticity might be a factor of 100 or so below this. Pulsar J0437-418 is the nearest millisecond pulsar, and is a good candidate. Its ellipticity would have to be only about 10^{-8} to produce radiation at an amplitude that GEO600, for example, could detect if it implemented narrow-banding.

When a neutron star is formed in gravitational collapse, it will probably be rapidly rotating, say with a spin rate of 200 Hz or more, and it should be formed with random irregularities that will manifest themselves as excited normal modes of oscillation of the new star. The subsequent behaviour of these modes depends on whether they are stable or unstable.

A stable mode will damp out quickly, in a time of the order of 100 ms or so. The primary frequency will be determined by the mode frequency, not the spin of the star. If the energy in such a mode is a fraction of the binding energy released in the collapse, the effective amplitude (equation (11)) will be similar to that of a gravitational collapse that releases the same amount of energy. Even if the energy in such a mode were small compared to that of the original collapse burst, the radiation might nevertheless be detectable in relatively nearby events, say a supernova in the Virgo Cluster observed with a stage-two detector. Recall, as we saw in equation (11), that the effective amplitude does not depend on how slowly or quickly the mode damps away, just on the total energy it radiates over time.

Very rapidly rotating neutron stars can have unstable normal modes by virtue of the CFS instability [11]. It is not clear whether viscosity in the neutron star will suppress these instabilities, but newly formed neutron stars are likely to be hot enough that their viscosity is negligible. These modes will be present when rotation rates reach about 1 kHz (depending on the equation of state), but the mode frequencies themselves are likely to be low, perhaps a few hundred Hz. Being unstable, they grow until they radiate away enough angular momentum from the star to reduce its spin to below the instability point. Here, the energy carried away would be large, a good fraction of the binding energy, and the frequency would be low, so the effective amplitudes of these modes could be large enough to make them detectable from the Virgo Cluster even by first-stage interferometric detectors. The radiation might last only a few seconds. Certainly, detectors should search for such radiation associated with every optical supernova detection.

Accreting neutron stars can be driven to emit strong continuous gravitational-wave signals in at least two ways: by the excitation of the unstable normal modes of a sufficiently rapidly rotating neutron star, or by driving a ‘wobble’, or precession, through the accretion of angular momentum not aligned with the spin of the star. We consider each process in turn, after making some remarks about where accreting neutron stars can be found.

Accreting neutron stars form the central object of most binary x-ray sources in the Galaxy. These fall into two distinct groups, the low-mass and high-mass binaries, according to the mass of the companion star. Accretion rates of $10^{-10} M_\odot$ yr^{-1} are usually sufficient to produce the observed power. The x-ray phase of such binaries may last only 10^4 yr, so the change in the mass of the neutron star during this phase will be negligible. But the accreting material brings in spin angular momentum, and neutron stars are observed to spin up significantly during this period.

In high-mass systems, the transfer of mass to the neutron star also brings the two stars closer together. When the neutron star enters the atmosphere of its high-mass companion, the x-radiation is absorbed and we no longer see the object. The system is called a Thorne–Zytkow object. From the outside it looks like a normal giant star. Eventually the neutron star spirals close to the core of the giant, and when the giant becomes a supernova, a Hulse–Taylor binary could result. In the Thorne–Zytkow phase, the neutron star can accrete at much higher rates than it had been doing in the x-ray binary phase, up to the order of $10^{-8} M_\odot$ yr^{-1}. Its x-rays do not get out of the stellar envelope, which looks like a standard red giant from the outside. This phase can actually last longer than the visible x-ray phase.

If accretion has spun a star up to the CFS gravitational radiation instability point, Wagoner [12] pointed out that it will go non-axisymmetric and radiate any further accreted angular momentum away in gravitational waves. Because the angular-momentum accretion rate is proportional to the mass accretion rate, the gravitational-wave luminosity of such a star accreting from a thin disc will be proportional to its x-ray luminosity. This means that the gravitational-wave amplitude can be inferred from the x-ray flux, without needing to know the distance to the system. The frequency of the radiation, however, depends on the unstable mode, and is at present poorly known. The mechanism could presumably operate at other instability points.

A star that is not spinning fast enough or is too cool to be subject to the CFS instability can nevertheless radiate gravitational waves if it accretes angular momentum not aligned with its spin axis. In high-mass x-ray binaries, there seems to be no reason to expect that the spin of the neutron star, which is probably significantly affected by the random hydrodynamic processes in the collapsing core that give the star its kick, should be aligned with the orbital angular momentum. When it accretes material from its companion, this will come in with angular momentum in the direction of the orbital angular velocity. The result will be to induce a certain amount of precession, depending on how rigid the crust of the star is. To an outside observer, the precession frequency is nearly the frequency of the spin of the neutron star, and the gravitational radiation comes out mainly at the precession frequency (not twice this frequency), at least for small precession angles. A paper on this is in preparation [13].

4.4. Stochastic background: our earliest view of the Big Bang

Just as the Big Bang left us the cosmic microwave background radiation, it is also likely to have left a background of gravitational radiation. Because gravitational waves interact so weakly with other matter, this radiation is genuinely primordial: apart from a cosmological redshift, it is unchanged since it was produced. This is an important difference from the microwave background, which was thermalized and strongly coupled to matter until the epoch of recombination. While the microwave background comes to us from about 10^5 years after the Big Bang, any gravitational-wave background will come directly from a much earlier time, possibly only 10^{-35} s after the Big Bang.

The radiation is stochastic, consisting of many individual components superimposed on one another in a random way. It can be adequately characterized by its energy density as a function of frequency. The conventional way of doing this is first to define its spectral energy density $\rho_{gw}(f)$ to be the energy density per unit frequency, and second to normalize this to the energy density that is required to close the Universe, ρ_c, by the equation

$$\Omega_{gw} = \frac{f}{\rho_c}\frac{\mathrm{d}\rho_{gw}}{\mathrm{d}f}. \tag{14}$$

We note that $\rho_c = 2 \times 10^{-43}$ J m^{-3} for a Hubble constant of 100 km s^{-1} Mpc^{-1}.

This definition is a natural one for radiation that is produced by physical processes that have no natural length scale. For such scale-free radiation, the energy density Ω_{gw} will be independent of frequency. Some models of the stochastic background, however, suggest a spectrum whose amplitude is independent of frequency [14], in which case $\Omega_{gw}(f) \propto f^3$. This favours ground-based detectors.

A background of gravitational waves appears in a single detector as simply another source of noise. There are two ways to identify it:

- In a single detector, if one has confidence in the characterization of the instrumental and environmental noise, and if the observed noise is larger, then one could attribute the excess noise to a background of gravitational waves. For ground-based detectors, the expected gravitational-wave noise level is so low that it will not be seen against expected instrumental noise sources. But for space-based detectors this may not be the case, and in the section on LISA below we discuss this in more detail.
- Using two detectors, one can cross-correlate their output data streams, essentially by multiplying them together and integrating over an observation time T. The instrumental noise, assumed independent, largely cancels out, while the gravitational-wave noise, being the same in both detectors, sums systematically. This technique works, provided the two detectors are close enough together to respond in the same way to any given component of the stochastic gravitational-wave field. In practice, detectors will be separated by significant distances, and this causes their mutual gravitational-wave response to decorrelate somewhat. Nevertheless, as we describe below, the technique can still dig deeply into instrumental noise and will set stringent limits, or make sensitive detections.

A two-detector correlation experiment relies on the independence of the instrumental noise, so in particular experimenters must ensure that the detectors are spaced far enough apart that any ground-vibration noise does not correlate between them. We noted above that the sensitivity of a pulsar search increases as the square root of the observation time. In that case, one is searching for a coherent signal against noise. But for a stochastic background, one is looking for a random, noisy background against instrumental noise. For this reason, the sensitivity to h grows more slowly, only as $T^{1/4}$. The sensitivity to the energy density is the square of this, so it grows as the square root of T. The signal should also grow with the bandwidth B of the correlation experiment, and in fact in the same way as with time, since both variables correspond to just adding in more independent data. The result is that the minimum detectable stochastic h in a correlation experiment between two identical and independent detectors is

$$h_{stoch} = \frac{\sigma}{(BT)^{1/4}}, \tag{15}$$

where σ is the RMS noise in a single detector over the same bandwidth. If the two detectors are not identical, then σ is replaced by the geometric mean of the two different values of σ. For bandwidths of 1 kHz and observation times of 10^7 s, this means an improvement of a factor of 10^5 in energy density sensitivity, or of 300 in amplitude sensitivity, over the broadband sensitivity of an interferometric detector.

Our equation for the flux of gravitational waves, equation (8), converts to an order-of-magnitude estimate of the energy density in the stochastic case simply by dividing by c. For an amplitude of 10^{-22} (the broadband sensitivity of an advanced detector), the energy density divided by the closure density is 10^{-4}. A cross-correlation experiment will improve this limit by 10^5, so it should be able to detect a stochastic background with an energy density as small as 10^{-9} of the closure density. In fact, more accurate calculations [15] show that with the sensitivity and optimum bandwidth of the advanced LIGO interferometers, we should be able to go to about 10^{-10} of the closure density.

Bar detectors can also look for a stochastic background. Recently Astone *et al* [16] have looked at the different observing possibilities for bars and interferometers working together, and we conclude that it could well be worthwhile doing a joint cross-correlation experiment to search for stochastic gravitational waves, at least before the full-scale advanced interferometers become operational. If in our above discussion the two detectors are an interferometer and a bar, the correlation can only be done over the bar's bandwidth, so the data from the interferometer will have to be filtered to this bandwidth before the correlation is performed. What we have shown, however, is that the resulting sensitivity depends only on the broadband burst sensitivities of the two detectors, and not otherwise on the bar's bandwidth. In other words, a 10^{-21} bar and a 10^{-21} interferometer can do just as good a search as two 10^{-21} interferometers, provided they are looking at the same central frequency (since the gravitational-wave amplitude may depend on frequency). Several bar–interferometer combinations have been studied in the PhD thesis of Compton at Cardiff [17]. This result may lead to bar–interferometer collaborations in the near future.

5. Low-frequency sources detectable from space: LISA

Sources of gravitational waves that emit in the low-frequency regime accessible from space are mainly either stellar-mass binary systems with relatively large separations and therefore weak gravitational fields (far from the black-hole line in figure 3), or massive systems that are highly relativistic and therefore almost inevitably contain black holes. The exceptions are the Sun (which influences a detector through its time-dependent near-zone Newtonian field) and the stochastic background.

All of these sources tend to be long-lived. Even a black-hole coalescence has a natural timescale of several months, and non-relativistic binaries in this window live for thousands or millions of years. Therefore the appropriate way to display the strength of a source compared to the sensitivity of a detector is to assume, say, a 1 yr observation time, and show the intrinsic amplitude of the source against the noise of the detector after integrating for a year (see the data-analysis section below). We draw such a diagram for the LISA detector in figure 6. Shorter-duration events are shown at an amplitude that correctly represents their signal-to-noise ratio if they are extracted by optimal pattern matching. Most long-lived sources have a fixed frequency; those that do not are shown at the highest frequency they reach.

In this diagram we show the LISA noise curve in two ways: the lower, light curve is the actual noise in a 1 yr integration. The upper, bolder curve is a fairer representation of LISA's sensitivity. It is set at a threshold of a S/N ratio of 5, taking into account antenna-pattern effects as LISA rotates during its orbit of the Sun: any source above this threshold will be detectable with confidence over almost the entire sky. Sources below this line will be detectable if there is independent information about them (such as their exact frequency) that will allow one to lower the confidence threshold. And when we discuss the cosmological stochastic background, the light 1σ noise curve is the appropriate comparator.

LISA makes use of the phase and amplitude modulation of signals to return directional and polarization information. For weak sources, phase modulation provides no resolution below about 1 mHz, where the gravitational wavelength becomes comparable to the orbital radius of 1 AU. Accuracy should be at the level of a few tens of degrees above this. But for strong sources, such as the massive black-hole coalescences described below, angular accuracy can be better than 1 arcminute. Amplitude modulation can be used to determine not only the polarization of the signal but also its direction. Below 1 mHz, this still provides directional accuracy at the level of a few tens of degrees, depending on the signal-to-noise ratio.

The second LISA interferometer provides independent information about the signal. If the signal is strong, then even if there is common noise between the two interferometers, the information will not be degraded, and we can treat the two interferometers as genuinely independent. In this case, there is direct polarization information in the two different signals, and this can be combined with the amplitude modulation to improve the polarization sensitivity and the directional information. Studies are underway at this time in order to quantify this.

5.1. Binary systems

One of LISA's main targets will be galactic binary systems, particularly those containing neutron stars and/or white dwarfs. Although all three known Hulse–Taylor pulsar systems emit orbital gravitational radiation at frequencies somewhat too low for LISA, LISA has a much greater range than radio pulsar surveys. The statistical analyses of pulsar binary observations mentioned above [8, 7] suggest that there should be of the order of 100 neutron-star binaries in the Galaxy within the LISA frequency range, and LISA would be able to see them all. If the larger estimates based on evolution calculations [9] are right, then there may be thousands, and one or two may even lie below the 'chirp line' in figure 3. LISA will be able to determine their distance directly from the chirp information.

With so many neutron-star binaries visible, there should accordingly also be tens or hundreds of neutron-star–black-hole binaries and even a few black-hole–black-hole systems, which will be stronger and therefore even more easily detectable. In fact, it is likely that there will be of the order of one black-hole–black-hole binary at a high-enough frequency to be visible from the Virgo cluster, which is shown in the diagram.

There are other binaries that ought to be even more plentiful. Some known x-ray binaries and cataclysmic variables are in the range of LISA; in fact, if they were not detected by LISA, it would be disastrous for general relativity. Many of the Thorne–Zytkow stars which might provide high-frequency sources from accretion onto the embedded neutron star could also radiate in the LISA band from their orbital motion.

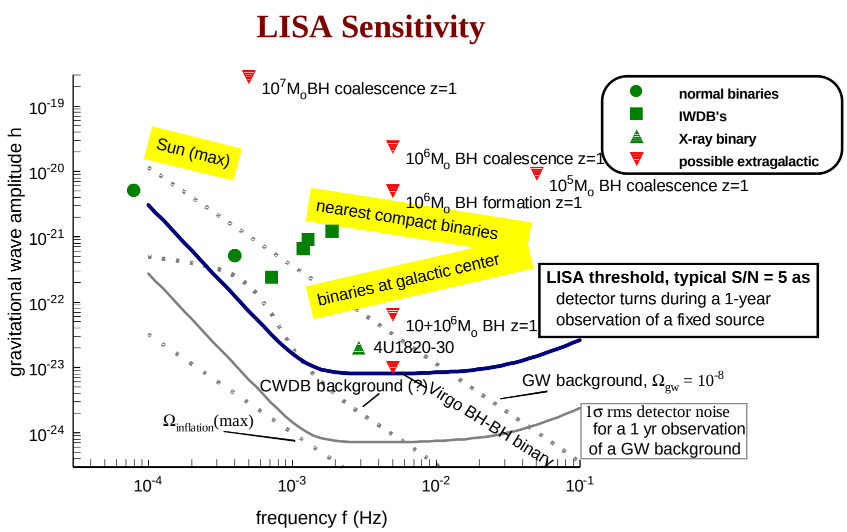

Figure 6. Strength of various sources and the sensitivity curve of LISA. Plot of the intrinsic amplitude of likely gravitational waves against their frequency. Most LISA sources will be approximately monochromatic. The full curve is the one-year sensitivity curve, the amplitude that could be detected by a single (two-arm) LISA interferometer in a one-year observation with confidence, i.e. with a signal-to-noise ratio of 5, allowing for the rotation of LISA during an observation. Below it, drawn faintly, is the actual RMS noise level for a one-year integration. The gravitational-wave amplitude h is shown for different types of periodic and quasiperiodic sources. The strongest sources in the diagram are binaries of massive black holes at cosmological distances, observed as they coalesce due to the orbital emission of gravitational waves. They have been placed in the diagram rather arbitrarily at their coalescence frequency and at an amplitude that correctly shows their SNR in relation to the LISA sensitivity curve, for a distance $z = 1$ with $H_0 = 75$ km s^{-1} Mpc^{-1}. The expected signals from some known binaries are indicated by circles and boxes, and are identified in the text. The nearest neutron-star and white-dwarf binaries at any frequency should lie in the band labelled 'nearest compact binaries'; the band below that shows the amplitudes expected from 'typical' white-dwarf binaries near the galactic centre. It is possible that the shortest period stellar-mass black-hole binary in the Virgo Cluster is in the position shown by the point lying on the sensitivity curve. A possible cosmological background left from the Big Bang is shown here at an energy density per decade of frequency today that is 10^{-8} of the total needed to close the Universe, again for $H_0 = 75$ km s^{-1} Mpc^{-1}. An upper limit to that generated by inflation is also shown. If there is a confusion limit due to galactic binaries, as discussed in the text, then it might appear as shown. For these backgrounds, the LISA RMS noise level (lower curve) is the appropriate sensitivity guide. The band labelled 'Sun (max)' is where solar g-modes might produce strong near-zone (Newtonian) gravitational perturbations observable by LISA. (Figure reproduced from [18].)

Importantly, there should be a large number of white-dwarf binaries, which are very difficult to detect by optical observations. Our present observational limits on their population are weak, and it is possible (or even likely) that they will be so plentiful that they will provide a confusion-limited background at low frequencies [19]. LISA observations of binaries would provide a rich harvest of astrophysical returns. One of the most interesting pieces of information would be the polarization of the signal. This will tell us the three-dimensional inclination of the orbital plane. For a known binary, whose mass function is known from spectroscopic observations, and whose primary mass is estimated from models, knowing the inclination determines the mass of the secondary. Then the intrinsic amplitude of the gravitational waves from the system will determine the distance to the binary. This extra information will be crucial for modelling such systems. Notice that this can be done even if LISA does not observe the 'chirping' of the system. Instead, it makes use of independent electromagnetic observations of the system.

5.2. Stochastic background

This radiation would appear in LISA observations simply as a noise in an individual detector. Provided it is above other noise sources, and provided we have confidence in identifying or limiting other noise sources, then LISA could make deep searches, particularly in the frequency range above 1 mHz. COBE observations of the microwave background fluctuations can be adapted to provide upper limits on the amount of radiation that could have been produced by parametric amplification of quantum fluctuations by inflation. This is shown in the figure, and it cuts close to but below the LISA noise curve. However, there is very likely to be a background due to white-dwarf binaries at this frequency that will overwhelm this cosmological one, in which case no detector would be able to see this cosmological background.

5.3. Massive black holes

The model that active galactic nuclei contain supermassive black holes has gained wide acceptance among astronomers in the last decade. These holes may have masses up to $10^9\ M_\odot$ or more. But active nuclei are rare, and most galaxies may have seen only modest amounts of activity in their past. However, there is growing evidence that ordinary galaxies and perhaps even dwarf galaxies commonly contain more modest black holes in the mass range 10^5–$10^7\ M_\odot$ (which we shall refer to as simply massive black holes, in contrast with supermassive black holes). As figure 3 shows, this is the mass range that a space-based detector would be sensitive to. The supermassive holes radiate at too low a frequency, but the massive black-hole range can radiate at LISA frequencies in at least three ways: as binary systems that coalesce, as large black holes that accrete smaller-mass compact stars and holes, or in the process of formation of the holes themselves.

A binary merger of two such massive black holes is the strongest source we anticipate for LISA, and one of the most interesting. A merger at a redshift of one of holes with a mass $10^6\ M_\odot$ would have an amplitude signal-to-noise ratio of more than 10^4. This is so strong that LISA will be able to give a position to less than 1 minute of arc, which should allow the galaxy cluster containing the merger to be identified optically, even at $z = 1$. The combination of the redshift obtained from this optical identification and the distance to the source provided by LISA itself from the chirp rate of this signal would allow a determination not only of the Hubble constant—which ought by then to be known by other means—but even more importantly of the deceleration parameter q_0 of the Universe, and it should find both of them to the unprecedented accuracy of better than 1%. From this one can infer the mass density of the Universe or decide whether there must be a cosmological constant. This is certainly one of the most exciting prizes for the LISA project.

Massive black holes in galactic centres should also occasionally swallow up stars. While main-sequence stars and giants are so large that they will be torn apart by tidal forces before they reach the horizon, neutron stars and stellar-mass black holes will remain essentially point particles that follow very complex orbits until they finally fall into the hole. These are not easy to model exactly, but with approximate matched filters that follow portions of the orbit, it should be possible to see these event at redshifts of 1 or so. They should be much more plentiful than black-hole mergers. These are particularly interesting because they can also test strong-field gravity near black holes. They can determine whether the massive hole is well described by the Kerr metric, which in general relativity is the unique solution for a vacuum black hole.

6. Conclusions

Gravitational-wave detection is one of the last challenges of general relativity: studying the waves goes to the heart of the dynamics of gravity. Nothing could be more fundamental in the study of classical gravitation theory. In addition, the interpretation of the source information leads to further fundamental physics information: tests of the speed of gravitational waves, searches for further gravitational fields, tests of the uniqueness of the Kerr solution for rotating black holes and searches for a cosmological background of gravitational waves. Placing a gravitational-wave detector in space, where it can see the low-frequency waves emitted by binaries and distant black holes, is the goal of ESA's Horizon 2000+ project LISA, in which ESA has firmly focused its attention on doing fundamental physics in space. Although efforts will be made to bring the rather distant launch date of LISA forward, even a launch in 2015 will be worth waiting for, because LISA will test gravity and glean astronomical information that can be obtained in no other way. Ground-based detectors should, by that time, have detected gravitational radiation directly and examined at some level its consistency with general relativity. LISA will then open up the low-frequency window, with a sensitivity far surpassing that of the ground-based detectors, and for me the most interesting part of these observations is likely to be to see what LISA discovers that we did not anticipate.

References

[1] Schutz B F Data processing analysis and storage for interferometric antennas 1991 *The Detection of Gravitational Waves* ed D G Blair (Cambridge: Cambridge University Press) pp 406–52

[2] Schutz B F The detection of gravitational waves 1996 *Astrophysical Sources of Gravitational Radiation* ed J-P Marck and J-A Lasota (Paris: Springer)

[3] Schutz B F 1986 Determining the Hubble constant from gravitational wave observations *Nature* **323** 310–1

[4] Cutler C, Apostolatos T A, Bildsten L, Finn L S, Flannagan E E, Kennefick Dl, Markovic D M, Ori A, Poisson E, Sussman G J and Thorne K S 1993 The last three minutes: issues in gravitational wave measurements of coalescing binaries *Phys. Rev. Lett.* **70** 1984

[5] Blanchet L, Damour T, Iyer B R, Will C M and Wiseman A G 1995 *Phys. Rev. Lett.* **74** 3515

[6] Taylor J H and Weisberg J M 1989 Further experimental tests of relativistic gravity using the binary pulsar PSR 1913+16 *Astrophys. J.* **345** 434–50

[7] Narayan R, Piran T and Shemi A 1991 Neutron star and black hole binaries in the Galaxy *Astrophys. J.* **379** L17

[8] Phinney E S 1991 The rate of neutron star binary mergers in the Universe: minimal predictions for gravity wave detectors *Astrophys. J.* **380** l17

[9] Tutukov A V and Yungelson L R 1993 *Mon. Not. R. Astron. Soc.* **260** 675
[10] Yamaoka H, Shigeyama T and Nomoto K 1993 *Astron. Astrophys.* **267** 433
[11] Friedman J L and Schutz B F 1978 Secular instability of rotating Newtonian stars *Astrophys. J.* **222** 281–96
[12] Wagoner R V 1984 *Astrophys. J.* **278** 345
[13] Schutz B F in preparation
[14] Brustein R, Gasperini M, Giovannini M and Veneziano G 1995 Relic gravitational waves *Phys. Lett.* **361B** 45
[15] Flanagan E E 1993 *Phys. Rev.* D **48** 2389
[16] Astone P, Lobo J A and Schutz B F 1994 *Class. Quantum Grav.* **11** 2093–112
[17] Compton K 1996 *PhD Thesis* University of Wales Cardiff, Cardiff, Wales
[18] Bender P *et al* 1996 *LISA: Pre-Phase A Report (MPQ 208)* Max-Planck-Institut für Quantenoptik, Garching, Germany
[19] Hils D, Bender P L and Webbink R F 1990 *Astrophys. J.* **360** 75