Gravitational Wave Sources: An Overview

Bernard F. Schutz

Albert Einstein Institute, Golm, Germany
Department of Physics and Astronomy, Cardiff University, Wales

Abstract. With full-sensitivity operation of the first generation of gravitational wave detectors now just around the corner, and with the LISA space-based detector entering its final design stage, I review the wide variety of predicted sources from the perspective of what further theoretical work may be needed to assist in their detection. Some sources, such as binary black holes, require good theoretical models from which search templates for matched filtering of the data streams can be computed. Others, such as searches for un-modelled bursts, require clever robust search algorithms not tied to detailed waveform models. Still others, such as searches for continuous waves from pulsars, are compute-bound and need improved efficient computer algorithms. The sensitivity of initial ground-based detectors will depend in part on how good we are at searching the data. In the longer term, the amount of information we can extract from the LISA data stream will depend in part on how good we are at removing strong signals so that we can recover the weaker ones as well.

INTRODUCTION

The LIGO and GEO600 detectors are well advanced in commissioning, and have already gathered data in coincidence. The VIRGO detector will come on-line within a year or so, and bar detectors continue to monitor the sky for large gravitational wave bursts. Moreover, the low-frequency LISA detector is now very firmly under development as a joint NASA-ESA project, with a launch expected in 2011.

In parallel with the hardware development there has been a major effort to build software for data analysis. LIGO and GEO have collaborated on the public-domain LIGO Analysis Library (LAL), which is a public library designed to cover all aspects of data analysis, with strict protocols regarding compatibility and verification. Data analysis is not simply a matter of keeping up with the incoming data stream; gravitational waves are expected to be buried under the noise of the initial detectors and may only be recognized by clever extraction methods.

All this work is guided by our expectations about what sources we can expect to see. Indeed, detecting gravitational waves is likely to require a mix of five key ingredients:

1. good detector technology
2. good waveform predictions
3. good data analysis methodology and technology
4. coincident observations in several independent detectors
5. coincident observations in electromagnetic astronomy

Source studies aim at items 2. and 5., and at understanding what information is likely to come from observations. They underpin the work that goes into item 3. Source

3

studies require input from both general relativity and astrophysics. Unfortunately, the nonlinearity of general relativity makes detailed computation of expected waveforms from most sources difficult. This nonlinearity is unavoidable; it is an essential part of the problem, since almost all sources are driven to radiate by motions governed by their own self-gravitation. The astrophysics input to our source-prediction efforts is important because it helps focus effort on the most interesting and/or promising sources. But our astrophysical information is limited because we have not yet observed in gravitational waves. Electromagnetic astronomy provides constraints and estimates of what we can see, but there are large uncertainties both in the event rates and in the intrinsic amplitudes we can expect.

In the next section I will review the mathematical techniques that are used in general relativity to make estimates of source strengths and waveforms, and in subsequent sections I will review what we know and don't know about the main expected sources. Although I regard the fifth item above — coincident observations with electromagnetic astronomy — as extremely important, there is not enough room in this talk to cover it. Several other speakers at this meeting have discussed this issue in depth. I shall address this theme in another conference proceeding[1]. For a recent, more detailed review of source theory, see Cutler and Thorne[2].

THE TOOLS OF THE TRADE

Many gravitational wave sources involve the orbital motion of two compact bodies, which might be white dwarfs, neutron stars or black holes. Relativists address the *two-body problem* analytically with two methods:

- post-Newtonian methods (weak-field/low-velocity approximation);
- radiation-reaction methods (*restricted 2-body problem*, where one mass is very small).

The simpler *one-body problem* (isolated neutron star or black hole oscillations) can be studied with perturbation approximations. In fact, it seems that even the last stages of the merger of two black holes can be approximated as a single perturbed black hole, even before the horizons merge. One might call this "cosmic modesty", because what appears to be going on is that the effective potential for photon motion near black holes peaks outside the horizon, and so when two black holes are close enough that they share a common "barrier" to photons, then to an outside observer they look like a single distorted hole. The photon barrier seems to draw a veil across the messy interior. This is called the close-limit approximation[3].

Outside these approximation regimes, and for most other sources, numerical simulation studies are our only hope. Nonlinear simulations of black-hole merger are getting better but still have a long way to go before they begin to provide convincing waveform predictions. I will review this work below.

Downloaded 27 Feb 2008 to 194.94.224.254. Redistribution subject to AIP license or copyright; see http://proceedings.aip.org/proceedings/cpcr.jsp

Quasi-Newtonian estimates

Although detailed predictions require considerable analytic and numerical effort, it is helpful to have simpler tools that offer quick estimates for wave strengths. Newtonian gravity modified by the quadrupole formula for gravitational radiation reaction is simple to use and provides considerable insight into the characteristics of many sources. I give here a few key formulas and show how we can make some powerful inferences from them.

Every self-gravitating system has a natural frequency of oscillation, given roughly by

$$\omega_{\text{natural}}^2 \sim \pi G\rho \qquad \text{or} \qquad f_{\text{natural}}^2 \sim G\rho/4\pi, \tag{1}$$

where ω_{natural} is the natural angular frequency (radians per second) and f_{natural} the circular frequency (Hz) associated with the motion. In this formula ρ is the mean density of the system averaged over the region engaging in the motion.

This is a surprisingly robust formula, if one only wants estimates good to factors of two or so. For the Earth, for example, it gives the orbital frequency of a satellite skimming to the top of the atmosphere, but it also gives the frequency of the vibrations of the Earth when earthquakes excite it. Thus, it tells us that, to within a factor of about 2, if a satellite is directly over the site of an earthquake, it will arrive at the diametrically opposite side of the Earth at the same time as the seismic wave arrives! This might at first seem unlikely, since the speed of a seismic wave depends on the pressure inside the Earth, while a satellite is in free fall. But the pressure in the Earth is determined by the weight of the Earth, *i.e.* by the same mass and self-gravity that determines the orbital period of the satellite. This fundamental frequency (usually multiplied by two) will dominate the spectrum of gravitational wave emission whenever the restoring force in the emitting motions is the self-gravity of the source. This happens in all sources except gravitational-wave pulsars, where the frequency is set by the rotational speed of the neutron star; in this case twice the natural frequency gives an upper bound on the emitted frequency.

Gravitational waves are emitted in the quadrupole approximation with an amplitude given by the following formula:

$$h = \frac{2G}{c^4 r}\ddot{I}^{\text{TT}}, \tag{2}$$

where I^{TT} is the transverse-traceless part of the quadrupole tensor. I have left off indices on this equation, since we are working only approximately. An upper bound derivable from this formula is[4]

$$h \leq \left(\frac{GM}{rc^2}\right)\left(\frac{GM}{Rc^2}\right), \tag{3}$$

where M is the mass of the gravitational wave source, R its typical size, and r the distance from the source to the observer. This is simply the product of two dimensionless gravitational potentials: the typical internal potential governs the dynamics that produces the radiation, and the external potential (at the observer's location) provides the $1/r$ fall-off characteristic of all radiation fields. It is remarkable that these two potentials are all that is required for the upper limit. The expression clearly shows, by the way, that the
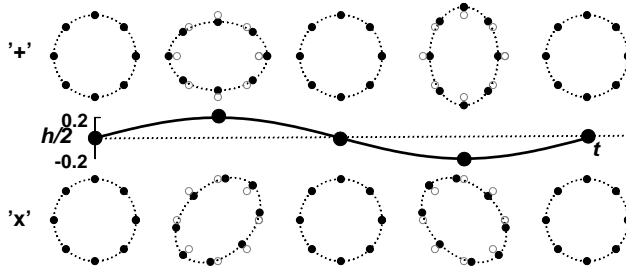
FIGURE 1. The two orthogonal transverse polarizations of a gravitational wave, illustrated by their action on a ring of free particles in empty space.

wave amplitude h is always a small fraction of the Newtonian gravitational potential of the source. Although this is an upper limit, it is a limit that is actually attained if the source is maximally non-symmetric, for example by an equal-mass binary system with circular orbits.

The quadrupole approximation allows us to estimate the effect of the source on our detectors, but it does not let us estimate the effect on the source itself of radiating the often huge amounts of energy that our sources can emit. The computation of energy and the associated radiation reaction was at one time a controversial subject in general relativity, but today it is well-understood. The simplest formula to use is the formula for the energy flux (per polarization, see below), for example from the Landau-Lifshitz pseudo-tensor or more rigorously as justified by Isaacson (see [5]):

$$F = \frac{1}{32\pi} \frac{c^3}{G} |\dot{h}|^2, \tag{4}$$

where the over-dot denotes a time-derivative. This formula has a form that is quite standard among classical field theories. One always expects that the flux will be proportional to the square of the time-derivative of the field; this is true for a scalar field and for the Poynting flux of electromagnetism[1], and it is clearly true also for general relativity. The factor c^3/G is simply a dimensional factor, and is the only combination of the two fundamental constants available for this problem that can produce a quantity with the dimensions of flux. The factor of $1/32\pi$, however, is harder to justify. This is the only part of the formula where one has to do some genuine tensor analysis in general relativity in order to get it right!

---

[1] The Poynting vector is often expressed as a product of the time-derivative of the vector potential and its space derivative, but space derivatives are proportional to time derivatives for wave motion.
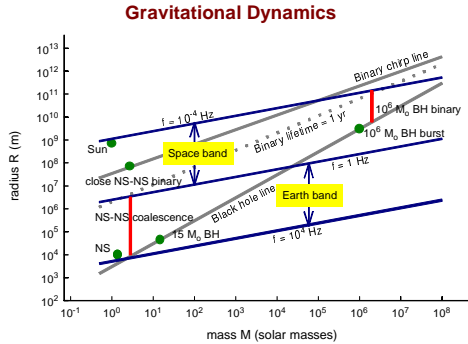
6

FIGURE 2. The mass-radius plane of sources divided into groups that radiate in the frequency ranges accessible from the ground or from space. The regions that chirp or coalesce within one year are also shown. In all cases the assumption is that the system radiates at its natural dynamical frequency, Equation 1, using the mean density of the system of the given mass and size.

This formula is reliable for a wave of small amplitude travelling in flat space, which means far from its source. But the total energy being lost by the source, obtained by integrating this flux over a sphere far from the source, represents the decrease in Newtonian energy of the source. We will come back to this.

Wave polarization

Gravitational waves in general relativity, like electromagnetic waves, have two transverse independent polarizations. These are usually illustrated as in Figure 1, by showing the action of the wave on a ring of free particles oriented transversely to the wave propagation direction. The effect of a wave is tidal: it distorts a circle into an ellipse, without changing its area to first order. The fractional distortion is half of the wave amplitude h. The two polarizations create ellipses rotated by 45° with respect to one another. These are orthogonal, and any other polarization is a linear combination of these two.

It is important to realize that the motions of test particles in response to the wave mimic the motions that took place in the source that generated the wave. If we look from the detector back in the direction of the source, then if we project the motions in the source onto the plane of the sky (again, transverse to the wave direction) and remove overall expansions (the distortions preserve areas), then what we are left with will be faithfully reproduced by the motions of test particles in our detector.

To be concrete, imagine that the source is a binary system oriented so that we see it in the orbital plane. Then the motion of the stars, projected onto the plane of the sky, is simply back-and-forth along a line. If the line is oriented along our x-axis, then the wave will be polarized in the sense of the upper diagram in the picture. Now, suppose we are looking at the same binary system from a vantage point above its orbital plane,

7

along its axis of rotation. Then we have both polarizations, separated by a phase: this is a circularly polarized wave. Clearly, by measuring the degree of circular polarization in a gravitational wave, we can infer the orientation of the plane of the orbit. This is information that is usually hard to obtain from optical observations of binary systems.

Luminosity and radiation reaction

If we replace the time-derivative in Equation 4 with $2\pi f$, where f comes from Equation 1, and we then integrate the flux over the whole sphere surrounding the source (by just multiplying by $4\pi r^2$), and finally we replace ρ by $M/\frac{4}{3}\pi R^3$ and h by its upper limit, then we get an upper limit on the luminosity of the source:

$$L \leq \frac{3}{32}\frac{c^5}{G}\left(\frac{GM}{Rc^2}\right)^5.$$
(5)

Notice two things about this equation. The first is the very strong dependence on the compactness of the source: the fifth power of the internal Newtonian potential. Second, the factor c^5/G provides the only dimensional part of the expression, so it has dimensions of luminosity. It is a natural luminosity, and is in fact the upper bound on all luminosities, gravitational or electromagnetic. The factor of $3/32$ is not exact, but will suffice for our work here.

From the luminosity we can easily estimate a characteristic timescale for changes to take place in the source as a result of its energy loss. This is just the time it takes to radiate a substantial fraction of the gravitational potential energy of the system:

$$\tau_{GW} = \frac{1}{L}\frac{GM^2}{2R} \geq \frac{16}{3}\frac{R}{c}\left(\frac{GM}{Rc^2}\right)^{-3},$$
(6)

where the lower bound comes from the earlier bound on L. This measures the lifetime of a gravitational radiator. It is often called the *chirp time*, since the radiation from a binary system will rise in frequency as the stars spiral together, and τ_{GW} gives the time-scale on which this chirping signal changes its frequency.

Measuring distances with chirping binaries

A remarkable fact is that a binary system that chirps is a standard candle: one can use its radiation to determine the distance to the binary[6]. One can see how this works from our formulas. The three measurable quantities h, f, and τ_{GW} depend on just three variables: M, R, and r. Measuring all three then allows on to solve for the variables and in particular to determine the distance r. This fact remains true if we do the calculation better, not just using bounds, and even when the masses of the binary stars are not equal. If the source is at a cosmological distance, then the variable r becomes the *luminosity distance* to the object.

This opens up important possibilities for future observations, particularly with LISA. If a distant binary of supermassive black holes is observed to coalesce, then we know its distance. If we can, by other means, obtain a redshift then we are able to do cosmology, investigating the evolution of the acceleration of the Universe[7, 8].

Summary: gravitational dynamics across the spectrum

In Figure 2 I summarize the results of the previous sections' estimates in a chart whose axes are the mass M of a system and its size R. In this log-log plot, lines of constant dynamical frequency are straight lines of slope 1/3. Three such lines are shown, for $f = 1$ kHz, $f = 1$ Hz, and $f = 10^{-4}$ Hz. These roughly divide the diagram into the frequency band of Earth-based detectors and that for space-based detectors.

Running through this diagram is the "black-hole line", a line with slope 1, below which there are no objects. It slices through the Earth-based observing band at relatively low masses: there are simply no sources visible from Earth with masses larger than about $1000 M_{\odot}$. Ground-based observing is a branch of stellar astronomy. The black-hole line does not close off the space band until around $10^8 M_{\odot}$, which allows LISA to see binary black holes in the interesting range of $10^5 - 10^7 M_{\odot}$.

Another line through the diagram is the "binary chirp line". Any system below this line will chirp during a 1-year observation: it will change its dynamical frequency by a measurable amount. All binary systems observable from the ground will therefore be standard candles, as will all the massive systems observable from space. Still more compact are systems that change catastrophically in one year; they are delineated by the "Binary lifetime = 1 yr" line. Interestingly this enters the LISA band at $10^6 M_{\odot}$: binaries more massive than this will coalesce within a LISA observation time.

Notice the great range of frequency spanned by planned detectors: 8 decades. This would take us from X-rays down to radio waves in the electromagnetic spectrum. It is not surprising that there is a great variety of sources that have been discussed.

Kicks and binary lifetimes

To illustrate the utility of our simple approximation formulas, let us ask two further questions. The first regards the possibility of "kicks" or recoils when two black holes merge; the second asks how close two black hole must be to merge within a Hubble time. Both are interesting questions for the theory of the evolution of black holes in galactic centers.

When two black holes merge, their product will acquire a linear velocity in their center-of-mass frame if the gravitational waves they emitted while merging carried off linear momentum. If the holes have equal mass and zero spin, then by symmetry no special direction is picked out and there will be no kick. If the holes have very different mass then the radiated momentum will be proportional to the smaller holes' mass, so the subsequent velocity will be scale with the ratio of the two masses. The ideal case will be when the holes differ in mass by a factor of a two or so. This might happen in the early

9

stages of formation of central black holes, if they are formed by hierarchical merging of smaller black holes. A strong kick could expel the merged hole from the center of the galaxy, shutting down the process.

Now, quadrupole radiation itself carries no net linear momentum, because it picks out no special direction. Instead, one needs quadrupole-octupole coupling, which means that the energy flux that is asymmetric enough to lead to a kick is down from the quadrupole flux by a factor of order v/c. The momentum flux is obtained from this by dividing by a further factor of c. Allowing for angular factors, it seems unlikely that the net momentum luminosity could therefore be larger than about $0.2Lv/c^2$, where L is given by Equation 5. The momentum has to be radiated within half of the last orbital period P_{orb}; otherwise cancellations in directions will occur. We thus get, as a limit on the recoil momentum p_{recoil},

$$p_{\mathrm{recoil}} < 0.1 \frac{LvP_{\mathrm{orb}}}{c^2}.$$

Now, we have estimates for all these quantities, depending on the compactness of the system, which is to say the separation between the holes when they begin to plunge together. If we divide by M to get the recoil velocity, then the bound becomes

$$\frac{v_{\mathrm{kick}}}{c} < 0.1 \left(\frac{GM}{Rc^2}\right)^4. \tag{7}$$

For non-spinning black holes, the factor GM/Rc^2 probably does not exceed 1/3, which leads to a bound of $v < 10^{-4}c$, or $300\,\mathrm{km\,s^{-1}}$. This could propel the final hole entirely out of a galaxy. Spinning black holes, which can approach one another more closely, seem to offer even better chances for strong recoil. Better estimates of recoil will soon come from numerical simulations of black hole mergers.

The second question is, to what separation do stellar interactions need to bring black holes in a binary system to ensure that gravitational radiation reaction will bring them to merger within a Hubble time? The answer is to set τ equal to the Hubble time and solve for R. The result is

$$R_{\mathrm{max}} = 5 \times 10^{-4} \left(\frac{M}{10^6 M_\odot}\right)^{3/4} \mathrm{pc}. \tag{8}$$

SOURCES

In this brief review we can only catalog the variety of sources that have been discussed and make a few remarks about what we need to understand about them to improve our chances of digging their signals out of the noise. I shall group them into three customary classes: burst sources, whose duration is short enough that one does not need to compensate for any phase modulation caused by the motion of the detector; continuous sources, where phase modulation during an observation is important; and stochastic sources, which produce random radiation over a broad spectrum. Sources of all classes are expected in both low and high frequency bands.

Bursts

Coalescing binary systems of neutron stars.

The first interferometers (LIGO, GEO600, VIRGO) should be able to see black-hole binary systems in coalescence out to 100 Mpc, and the weaker double neutron-star systems to about 20 Mpc[2].

While estimated event rates make it seem unlikely that there will be a neutron-star coalescence within this volume of space in, say, five years, it is not impossible. (For a list of references, see [2].) Rates are estimated on the basis of the statistics of observed binary pulsar systems, so they are lower limits on what might be there. Real rates could be higher for a number of reasons.

- Pulsar surveys are known to be incomplete, sparsely covering high galactic latitudes, and particularly limited in searches for binaries by the large amounts of computer time required for such searches. Quoted rates use estimates of their incompleteness, but these may be too conservative.
- Our Galaxy may be untypical of our neighborhood out to the Virgo Cluster. In particular, nearby starburst galaxies may have significantly higher rates.
- The first generation of star formation (Population III) may have formed binaries but left them scattered over the galactic halo, too far away to show up in pulsar surveys.
- Not all neutron-star binaries may evolve into binary pulsars. For example, there may be a population of systems where the oldest neutron star has accreted enough during its companion's giant phase to become a small black hole rather than staying as a detectable millisecond pulsar. Such systems would not be seen in pulsar surveys.
- Neutron-star binaries may have other ways of forming than as binary pulsars. It is possible that other stable states of matter for compact stars exist, such as strange-quark stars, and if these do not form pulsars then they will not turn up in surveys.
- If the MACHO objects that are seen to cause microlensing in the Galactic halo are compact enough, they could contribute a population of binaries at low masses, possibly down to $0.1M_\odot$.

These considerations suggest that searches for such systems ought to span a much bigger mass range than the ones associated with known pulsar binaries. To catch MACHO systems one should go down to $0.5M_\odot$ or lower if possible; to catch small-black-hole systems one should go up to $3M_\odot$.

Coalescing binary systems of stellar-mass black holes.

Black-hole binaries with component masses of order $10M_\odot$ are probably more attractive as early sources than neutron-star binaries. Although stellar evolution probably forms fewer such systems in any given volume of space, they can be seen 5 to 10 times further away, so the volume event rate is amplified by 100 to 1000. The reason for opti-

mism is that it has recently become clear that globular clusters manufacture black-hole binaries efficiently[9]. The black holes sink to the centers of the clusters due to relaxation, and their density there can be high enough to form binaries and to harden them up by three-body encounters. Eventually an encounter releases enough gravitational energy from the hardening binary to give the system the escape speed, and then they leave the cluster. Many of these systems will be compact enough to coalesce in a Hubble time. The first detectors could see these out to about 100 Mpc, and an event rate of one every two years is not impossible[2].

As with neutron-star binaries, there are ways in which the event rate could be higher than expected, for example if the galactic dark matter contains a big population of $100M_\odot$ black holes that form binaries (there are few observational constraints on this mass range) and/or if Population III star formation led to massive black holes (thought to be likely). Moreover, the globular cluster history of galaxies is not understood; if the population of clusters was much higher in the past, then so was the black-hole binary manufacturing rate. Again, searches should be pushed up to the highest possible masses.

Unlike neutron-star binaries, the detection of black-hole binaries will be seriously affected by our ignorance of the waveforms, particularly for holes with masses above a few tens of solar masses. In this upper mass range, the plunge and coalescence phase radiates in the most sensitive frequency band of the interferometers; without good matched filtering we may easily lose a factor of two or more in sensitivity and therefore in range, or a factor of ten in event rate.

Advanced LIGO should see an abundance of these sources, both neutron stars and black holes. Its range for $10M_\odot$ black-hole binaries may reach to redshift $z = 0.4$, and even neutron-star systems will be visible to 0.3 Gpc ($z \sim 0.1$).

Coalescing binary systems of supermassive black holes.

Coalescing black-hole binaries of much larger mass are the source that has created the most interest in the LISA project. Observing at frequencies 10^6 times lower than LIGO, LISA will be sensitive to the dynamics of black holes with masses 10^6 times larger than those LIGO can study. These events are so much stronger that LISA will see coalescences of black holes in the mass range 10^4–$10^7 M_\odot$ essentially anywhere they occur in the Universe. Although not strictly "burst" sources, since they typically last long enough for the motion of LISA to induce phase modulation, I will treat them here because physically they are so closely related to the stellar-black-hole coalescences.

The rate of coalescence is, however, just as uncertain as for the smaller systems. We know that most galaxies possess black holes in this mass range, and that galaxies merge very often. These mergers bring the central black holes of these galaxies together in the merged core, but it is not known if stellar dynamical friction is normally strong enough to bring them close enough to merge from gravitational radiation reaction within a Hubble time. (See Equation 8.)

If each galaxy has experienced one such merger in its lifetime, then LISA should see about one such event per year. The rate could of course be much higher, if mergers of galaxies are more frequent or if the central black holes are formed by the merger

12

of smaller black holes. There is reason to believe that Population III star formation led naturally to black holes of $1000M_\odot$[10]; to form holes of $10^6 M_\odot$ directly seems to require the help of dark matter. It will be interesting, therefore, to see whether LISA detects mergers of, say, $10^4 M_\odot$ black holes with $100M_\odot$ holes. These will be important clues to the early history of supermassive black holes.

Simulations of black-hole coalescence.

As mentioned above, and explained in more detail below, the detection of black hole coalescences will be made easier, the more information we have about what kind of waveform to expect. Groups around the world have been studying this problem for decades, beginning with the pioneering work by Eppley and Smarr[11] on the head-on collision of Schwarzschild black holes. But head-on collisions are not realistic, and fully three-dimensional simulations are needed. These stretch both the available computer resources and the ingenuity of scientists.

Here is a personal list of the most important challenges facing those who do simulations, and where we stand in terms of meeting them.

- Numerical stability of the field equations. The Einstein system contains 10 equations, not all independent. There is an infinite variety of ways of choosing a minimal combination to integrate on the computer. The adoption of the BSSN form of the equations has doubled or tripled the run-time for the longest simulations involving orbiting black holes, but these do not last forever. All three-dimensional simulations eventually show instabilities, after something like $50M$, where M is the light-crossing time across the black hole. There is a debate whether these come from bad boundary conditions, but the weight of evidence is that they are "constraint-violating modes", which are modes of the equations that grow when the constraint equations are not exactly satisfied[12]. Although unphysical, they will always appear in any finite-precision solution. Ways of controlling such modes are urgently being sought. Some relief from these problems may be provided by the Lazarus approximation[13], in which cosmic modesty (see above) is used to replace merging black holes with a single perturbed black hole. This prolongs the life of a simulation (hence the name Lazarus) but it does not remove the instability.

- Gauge conditions. The freedom to change coordinates has not yet been robustly explored in this field, partly because of concern about introducing instabilities, and partly because it is necessary to give a physical interpretation to the coordinate-dependent quantities that are calculated by the computer. Many groups now impose radial shift conditions that are designed to stop coordinates falling into the black holes. In addition, the group at the Albert Einstein Institute (AEI) has recently introduced a rotating coordinate system that follows orbiting black holes in such a way that they do not need to move across cells of their finite-difference grid[14]. This seems to have significantly helped to prolong the run-time.

- Outer boundary condition. With a finite computational domain, it is important to have a realistic outer boundary condition. With simple wave equations this is not difficult, but so far there is no satisfactory boundary condition in general relativity

13

that will, among other things, remove gravitational waves from the interior without significant reflection. The solution so far is to place the boundary far enough away that any reflections arrive only after the simulation fails for other reasons, for example instabilities. But one cannot guarantee that the simple conditions used at the moment do not introduce their own instabilities.

- Inner boundary conditions: excision of the singularity. No numerical simulation can cope with the singularity inside the black hole, so a region around it must be cut out. If this region is entirely inside the horizon, then in principle this has no effect on the evolution of the simulation outside, because of causality. But in a code with only finite accuracy, there is the possibility that something will "leak" from inside the horizon to outside, and this is particularly worrying if the inner boundary conditions — necessarily rather arbitrary — leads to instabilities in the solution. Stable three-dimensional excision has been demonstrated for moderate run-times for holes that do not move across numerical meshes[15], but the goal of moving holes with excised interiors is proving to be more difficult to achieve[16].

- Computer memory. The most efficient way to use scarce computer memory in finite-difference codes is to introduce adaptive mesh refinement, where grid points are added or removed according to the demands of numerical accuracy. This could particularly help to keep the outer boundary far enough away that reflections do not disturb an integration. Several groups are working on varieties of this technique. Unfortunately, they are not all working together, so progress on this purely computational issue is slower than it might be.

- Initial data. This is the most subtle and most physical of the problems facing simulations, and it is probably the key one which, if solved, would allow simulations to provide important physical information even with their numerical limitations. The problem is simply that we do not know what the right initial conditions are for two black holes in a quasi-circular orbit, just before they begin their plunge to coalescence. The AEI group has been using a popular method of determining the starting conditions (separation and velocity of the holes) that goes back to Bowen and York[17]. But the simulations do not behave as expected: the holes plunge together very quickly, and even when started further apart still plunge together almost head-on. It seems that this method does not give the holes enough circular velocity, or enough separation, or both. A promising alternative has been developed by the Meudon group, within the EU Network on Gravitational Waves[18]. It matches better to the post-Newtonian predictions of Buonanno and Damour[19]. The AEI group, which coordinates this network, is currently integrating this new initial data into its code, and simulations should be available later this year.

From the perspective of the AEI, which has the largest numerical relativity group at the present time, the field of numerical relativity worldwide looks under-funded and too dispersed; this is especially true of the US. It is important not only to provide more funding and job security for postdocs and students; it is also important that the groups receiving the funding coordinate their efforts. For example, the very successful EU network referred to above has regular meetings at which they review and renew their extensive organization charts of who works with whom, on projects that all agree are important. The numerical simulations are performed within a single computational

environment, the Cactus Computational Toolkit. I believe that this is a field that needs groups or collaborations of a critical size, and this is currently lacking in the US. This must be remedied urgently if numerical relativity is to provide the kind of information that we need for detecting gravitational waves from black hole binaries in the next few years.

Other burst sources

A number of other sources for ground-based detectors fall into the category of bursts.

- Gravitational collapse. This has been the motivation for building gravitational wave detectors from the days of the first bar detectors, but it is still not a source that we can make definite predictions about. Numerical simulations of gravitational collapse suggest that amplitudes from events in the Virgo Cluster will be too small even for Advanced LIGO[20], but these simulations do not yet include all the physics that is needed; in particular they do not treat neutrino transport, convection, and asymmetries in the proto-neutron-star in a three-dimensional manner. Simulations of gravitational collapse to supermassive black holes, on the other hand, suggest that these events might well be visible to LISA if that is the way that such objects form[21].

- Neutron star vibrations. The normal mode frequencies of neutron stars are mostly above 1 kHz, with the exception of r-modes (which I treat under continuous waves, below) and g-modes (which have not received very much study). It would be extremely interesting to be able to detect radiation from neutron-star modes, because their frequencies would inform us about their interiors. Just as the Sun has revealed its interior secrets through helioseismology, so too could neutron stars reveal at least the broad picture of their interior physics if we could detect two or three mode frequencies. It is possible that these modes could be excited by glitches in nearby neutron stars, but detecting them will require a third-generation special-purpose broadband high-frequency detector.

- Unexpected sources. Whenever a new electromagnetic window has been opened in astronomy, among the first observations have been systems that were completely unexpected. The first X-ray rockets saw not only the Sun but also Sco X-1. The first gamma-ray detectors in space detected gamma-ray bursts. The earliest radio observations detected what later turned out to be giant radio galaxies, powered by black holes. Our hope is, of course, that history will repeat itself with gravitational waves, especially considering the fact that the leap to gravitational radiation from electromagnetic observations is considerably greater than the leap from optical observations to X-rays. There are many hidden corners of the Universe that may contain compact systems of which we have had no warning; in particular, we do not understand the dark matter in galaxies, and there is no reason to think that it consists of just one homogeneous component, spread relatively smoothly through space. However, to detect such sources may require sophisticated signal analysis, and certainly requires careful combination of the observations of several detectors

working simultaneously. This is a subject that is receiving considerable attention within the LIGO LSC. It could be just as important for LISA as for LIGO.

Data analysis issues.

Coalescing neutron-star binaries are probably the best-understood of all the likely sources of gravitational waves. Although the hydrodynamics of their coalescence is not yet satisfactorily modelled, this occurs at such a high frequency that their detectability is not affected by this ignorance. Work over the past decades on the post-Newtonian orbital problem has progressed to the point that we have excellent templates with which to do matched filtering[22].

One remaining issue for neutron-star searches is that the parameter space for realistic systems is very large, especially if we go down to masses below $0.5M_\odot$ (see above) or introduce neutron-star spins. Searches will have to be done in a hierarchical manner, but it seems that this will not present a big problem.

Searching for black-hole binaries is a different matter, because their larger mass shifts their coalescence radiation into a more important part of the LIGO observing band. For holes of several tens of solar masses, the plunge and coalescence could account for half of the received signal power, while for larger masses it will completely dominate. In the absence of good waveform predictions, scientists are considering ways to extract at least some of this signal from the detector noise[2].

LISA's supermassive binaries will not need sophisticated data analysis to detect: they will stand out above its broad-band noise without any filtering. The problem for LISA will be to remove this huge signal accurately enough to permit good extraction of weaker signals from other expected LISA sources, which we discuss below. This will require the accurate determination of orbital parameters, masses, and spins. It may not be possible to remove the coalescence events themselves very accurately, but they only last a few minutes.

Data analysis for gravitational collapse is similar to that for unexpected sources: in both cases very little is known about what waveform to expect. Data analysis teams are exploring time-frequency methods, wavelets, filtering for simple short model waveforms, and other approaches. For neutron-star normal modes the data analysis will not be difficult. One must search the two-parameter space of frequencies and damping times, but this is much smaller than other searches demand.

Continuous sources

Continuous sources add two (related) complications to the search problem. First, a source that radiates continuously must inevitably have a weaker amplitude than a strong burst source, which carries a catastrophically large amount of energy away from its source. This means that we will typically need to observe such sources continuously for long periods of time to accumulate enough signal to make them detectable. This leads to the second difficulty. The long observation time leads to modulation of the signal by

16

the motion of the detector, and this modulation pattern depends on where the source lies on the sky. Therefore, the number of filters gets multiplied by the area resolution factor that this modulation gives, which is a function of the length of the observation and the frequency of the waves. The weakness of the signal demands accurate and long search templates, while the modulation multiplies by a large factor the number of templates that we need to use.

Gravitational wave pulsars.

All pulsars must be sources of gravitational radiation, because their tilted magnetic fields create a time-dependent quadrupole moment as they spin. But the amplitude of this radiation, for the inferred dipole field strengths of known pulsars, is so weak that it will be beyond the reach of even Advanced LIGO. Our upper limit in Equation 3 is not a good approximation to the radiation from a pulsar, because only a small fraction of the mass participates in emitting the radiation, and because the frequency is determined by the spin frequency rather than the natural dynamical frequency. We can therefore obtain a suitable formula by multiplying the limit by a factor ε, which represents the fractional deviation from axisymmetry of the shape of the star (called its ellipticity), and by multiplying by $(f_{\mathrm{psr}}/f_{\mathrm{natural}})^2$, the square of the ratio of the spin frequency of the star to its natural dynamical frequency given in Equation 1. (We use the square because there are two time-derivatives in Equation 2.) This gives the estimate

$$h_{\mathrm{psr}} \sim \varepsilon \left(\frac{v_{\mathrm{rot}}}{c}\right)^2 \left(\frac{GM}{rc^2}\right), \tag{9}$$

where v_{rot} is the rotational speed of the pulsar at its equator, or in other words $f_{\mathrm{psr}}R$.

A pulsar will be detectable if one can build up enough signal-to-noise in, say, a year to see it. This signal builds up as the square-root of the number of cycles of radiation received from the pulsar in the observation period, which will be twice the number of times the pulsar spins in a year. For initial interferometers, whose broadband sensitivity is roughly $h_{\mathrm{bb}} \sim 10^{-21}$, the sensitivity limit for a one-year observation will be of order

$$
\begin{aligned}
h_{\mathrm{yr}} &\sim h_{\mathrm{bb}} \left(2 f_{\mathrm{psr}} \times 1 \ \mathrm{yr}\right)^{-1/2} \\
&\sim 1.3 \times 10^{-26} \left(\frac{f}{100 \ \mathrm{Hz}}\right)^{-1/2}.
\end{aligned}
\tag{10}
$$

If we equate this to h_{psr} from Equation 9 — and take for the pulsar a standard radius of 10 km, a standard mass of $1.4 M_{\odot}$, and a distance of 1 kpc — then we can solve for the minimum detectable ellipticity:

$$\varepsilon_{\min} \sim 6 \times 10^{-7} \left(\frac{f_{\mathrm{psr}}}{100 \ \mathrm{Hz}}\right)^{-3/2}. \tag{11}$$

This formula is only a very rough guide. The noise in interferometers is not flat, so that at lower frequencies (say, at the Crab's frequency of 60 Hz) or at higher frequencies

17

(millisecond pulsars), the limits on ε get worse. Moreover, this limit is the value of ε that will give a signal equal to the noise level; for the signal to be detectable with confidence it must be three to ten times stronger, depending on the kind of search being performed. Importantly, the scaling shows that millisecond pulsars are detectable with smaller values of the strain than young pulsars.

Physical estimates of how much ellipticity a neutron star can sustain depend very much on the mechanism. The neutron-star crust is not very strong, and the mass in the crust is small, so that a crustal irregularity probably could not produce an ellipticity exceeding about 10^{-7}[23]. Stars with solid cores could, in principle, sustain ellipticities up to around 10^{-3}, limited by the square of the ratio of the radius of the core to the radius of the star. But such stars exist only with certain equations of state. Strong buried magnetic fields could produce observable ellipticities. Strange-quark stars are smaller (hence more relativistic) for a given mass, and could in principle radiate much more strongly. Most promising is a mechanism suggested recently by Cutler[24], in which a buried toroidal magnetic field forces the star into effectively a prolate shape, which is unstable and tips over so that the long axis is in the equator of rotation. Ellipticities up to 10^{-7} are possible for plausible values of the strength of the buried field. This could lead to detectable radiation, especially from millisecond pulsars.

Pulsars known from radio observations are all spinning down, and this energy loss sets a limit to their ellipticities, which comes from assuming that all the lost rotational energy is going into gravitational waves. This is almost certainly not the case for young pulsars, as the few measured values of the braking index (second time-derivative) confirm, but at least this gives upper limits on ellipticity. For the Crab the spindown limit is around 10^{-3}, but for millisecond pulsars, which radiate more effectively and have smaller energy losses, the spindown limits are near ellipticities of 10^{-6} or smaller. This means that it is possible that gravitational radiation contributes a significant part of the spindown of such stars. In Cutler's model, referred to above, gravitational waves could account for the biggest part of the spindown of millisecond pulsars.

Searches for gravitational wave pulsars are relatively easy to perform for known radio pulsars, because the data reduction for a known position and frequency is very fast. So, too, are searches for radiation from known point-sources, such as X-ray sources that are probably young neutron stars. Far more difficult, however, is a wide-area search for radiating neutron stars not associated with known radio pulsars. During an observation time of one year, a millisecond pulsar's signal will be modulated so much by the motion of the Earth that it can be located to within one arcsecond on the sky. This means that every one of the 10^{13} square arcsecond patches on the sky must be searched separately. And for young pulsars, one has to search over a wide range of spindowns[25]. The GEO gravitational wave group at the AEI is developing software designed to perform such searches efficiently, but even with teraflop computers it is unlikely that we will be able to survey the whole sky[26]. Searches will have to be confined to selected directions, such as globular clusters, the spiral arms of the Galaxy, or the galactic center. There is considerable potential here for surprises. We do not understand the physics of the neutron-star interior well enough to be confident that neutron stars are very symmetrical.

Low mass X-ray binaries.

There has been considerable study of neutron star r-modes, which are velocity-dominated modes that are destabilized by the emission of gravitational radiation (the CFS mechanism) but which are damped out if viscosity is sufficiently large. The current thinking is that young neutron stars are probably immune to this instability, but it may well operate in low-mass X-ray binaries (LMXBs)[27, 28]. Indeed, the evidence is that LMXBs do not spin up through accretion as much as they might, and it is possible that they are in a steady state of radiating gravitational waves that carry away the angular momentum that they accrete.

Although these sources are similar to gravitational-wave pulsars, they may be more difficult to detect. We know their positions, of course, but they are in binaries whose orbital parameters are not usually well enough known to remove the phase modulation exactly, so some searching will have to be done. The most problematic issue, however, is that the variable accretion rate in LMXBs probably feeds through into a slowly varying spin rate, and this is not easy to model.

If we succeed in modelling these systems well enough to track the phase of the signal for a time of order 4 months, then Advanced LIGO should be able to detect the radiation from Sco X-1, provided of course that the system's spin is limited by gravitational waves[2]. This is an exciting target to shoot for!

Compact white-dwarf binaries in LISA observations.

In the LISA observing band, most sources are long-lived, so removing orbital phase modulation will be a necessity for almost all observations. Fortunately for the data analysis, although not for the scientific return, the amount of modulation is small. A wave with a frequency of 1 mHz has a wavelength of almost exactly 2 AU, equal to the diameter of the effective telescope aperture LISA (or any other detector) synthesizes as it orbits the Sun while receiving waves from a continuous source. So below this frequency, phase modulation is negligible, and even at 10 mHz there are only some hundreds of resolution patches on the sky. So searching over these patches and removing modulation is not expected to present any problems.

The negative aspect of this is that LISA does not get good directional sensitivity. Contrast this with the case of a gravitational wave millisecond pulsar, where the directional sensitivity is at the arcsecond level but the problem of removing modulation makes wide-areas searches very difficult.

LISA will be surrounded by tens of thousands of Galactic white-dwarf binary systems compact enough to radiate in its observing band. Some of these are systems that are already known, and are therefore guaranteed sources: if LISA does not see radiation from them, then something is seriously wrong with general relativity. Studies of the likely statistics of the whole population suggest that below 2 mHz, the abundance of systems coupled with the failing angular resolution prevent LISA from resolving the binaries in frequency: each frequency resolution bin will contain more than one system radiating strongly enough to be registered. LISA will therefore be confusion-limited in

this range, and the signals from the galactic population create a noise level several times higher than LISA's instrumental noise, against which LISA has to look for supermassive binaries and other sources.

This noise drops below LISA's instrumental noise above about 2 mHz. All systems in the Galaxy radiating in this range will be resolved by LISA. Moreover, since they will all chirp (see Figure 2), LISA will measure distances to them all. Below 2 mHz there will still be many systems near enough to stand up over the confusion noise, so that altogether several thousand identifiable white dwarf and neutron-star binaries are expected. This survey of the galactic population will provide rich insight into galactic structure and into the evolution binary systems.

As argued above, finding the resolvable binaries will pose no problems for data processing. But as for the supermassive black hole systems, a bigger challenge will be to remove them from the data stream without contaminating it, so that unbiased searches for other sources can reach their optimal sensitivity. This will not be trivial: signals in LISA's comoving detector frame are spread over several frequency resolution bins, and many of them will overlap. An even more challenging problem will be do resolve and remove the signals that are close to the confusion limit. Indeed, with careful processing it may be possible to reduce the confusion noise by factors of 1.5 or 2 from current estimates.

Gravitational captures.

LISA will see another very important black-hole source: the captures of stellar-mass black holes by the supermassive black holes in the centers of galaxies. The big black holes sit in the centers of clusters of stars that contain a normal proportion of smaller black holes from normal stellar evolution, and occasionally one of these will encounter another star that puts it on a collision course with the central black hole. Some of these nearly-radial encounters will take the infalling body so close to the center that the loss of energy to the gravitational waves that it emits as it loops around the center is enough to trap it, so that it returns repeatedly, each time becoming more and more bound. In the last year of the evolution of this highly eccentric orbit, the body may execute thousands of orbits, all of them within three or four gravitational radii, before ultimately falling into the central hole[29].

The interest in this radiation is that the orbit is very sensitive to the mass and spin of the central hole, and it therefore constitutes an excellent probe of the geometry near the hole. It is nearly a test particle, so its orbit to lowest order (in the mass ratio) is a geodesic of the Kerr geometry of the central hole. By watching this evolve over tens of thousands of orbits, we will be able to map out the geometry precisely, and test one of the most fundamental theorems of general relativity: black hole uniqueness, otherwise known as the no-hair theorem.

Unfortunately, this will be a difficult source to detect. The radiation amplitude is proportional to the mass of the small black hole, which may be 10^{-5} of that of the central black hole. The event rate in any volume of space is, of course, much higher, since these encounters occur regularly in the history of the central hole. The nearest

events in one year may be no further away than redshift $z = 0.1$. The radiation will also have a higher mean frequency than that of the coalescence of supermassive black holes, staying much closer to the final coalescence frequency; this helps in two ways, first in placing the signal in the most sensitive region of LISA's spectrum, and second in improving the gain of matched filtering because there are more cycles. In a 1-year observation, then, such a system at $z = 0.25$ might have a signal-to-noise ratio of several tens, but probably not much higher than 100.

To achieve this level of detection, we will have to match the signal cycle-for-cycle over the whole complicated evolution of the orbit. Here is where the real challenge is. At present we do not know even how to calculate the orbit. The calculation requires us to find the equation of motion of the captured hole correct at least to first order in its mass; this is where we will see not only radiation-reaction effects but any conservative departures from geodesic motion which could affect the received signal phase. There has been important progress recently in defining an algorithmic way to do this; the next step is to implement this in an efficient computer program[30].

Even after it becomes possible to compute the orbits, we will have to face the huge size of the parameter space governing these orbits. Tiny changes in the parameters of the initial orbit will lead to major differences in the waveform months later. Even geodesic orbits around Kerr are not generally closed or even planar, and when the perturbations are taken into account they will be even more complex. Spin-orbit and spin-spin effects can be significant. Estimates of the number of distinguishable orbits outrun an conceivable computing power, so hierarchical searches of the kind that are being devised for the gravitational-wave pulsar problems will be needed. We have only a preliminary idea of even the degree of complexity of this problem.

If these barriers can be overcome, we shall be able to detect the nearest of these events. But more distant events will create a kind of confusion noise, similar to but more complex than that of the compact binary systems. The signals from capture orbits will chirp over a broad spectrum, so that each source is completely intertwined with every other one. The detection and removal of even the strongest sources is likely to require a simultaneous fit to all their orbits, and it is not at all clear how deeply this can be done.

We have here a gravitational wave background reminiscent of the cosmological X-ray background, which was unresolved in early X-ray missions, and which has only recently been shown to be composed of point sources at great distances. This situation is fundamentally related to Olbers' Paradox, the observation that in an infinite flat homogeneous universe the night sky would be infinitely bright; this would be true in gravitational waves as in light. The only reason that the sky is dim is that the universe is finite, in the sense that sources do not extend beyond a finite cosmological redshift. We always observe with an instrument that has limited resolving power, whether in angular measures or in frequency or in both, and it follows that the instrument will fight against a bright background sky if the universe contains many sources per resolution element.

In ground-based gravitational wave astronomy we are accustomed to sparse sources, but when LISA observes gravitational capture events there will probably be many hundreds or even thousands happening in the universe at any time; and the more distant ones provide an "Olbers limit" on our sensitivity to the nearer ones. This sensitivity limit will be at a higher level than instrumental noise, but it depends on that noise, because if the instrument were more sensitive then we could resolve more of them. Indeed, the

LISA design team has been discussing whether to improve the design sensitivity of LISA in order to guarantee that at least a few events are detectable. Since the sensitivity depends in part on knowing the complexity of the filtering problem, we are going to have to make an educated guess depending on the (incomplete) understanding of this problem at the time when the design has to be fixed.

Stochastic Gravitational Waves

The universe is filled with random gravitational radiation, the superposition of all the waves that have been emitted by every dynamical process, from supermassive black hole coalescences to feathers waving in the wind. We have discussed in the previous section how waves emitted by discrete sources blend together into a confused sea. LISA's observations of compact white dwarfs and of gravitational captures dabble at the shores of this sea. In this section we will discuss backgrounds that are composed of so many confused and overlapping sources that they can only be described statistically. These sources can be systems that we are familiar with from astronomy, or they can go all the way back to the Big Bang. We therefore distinguish between astrophysical background and a genuine primordial cosmological background, although of course astrophysical backgrounds fill all of cosmology as well. They just arose later than the primordial background.

The strength of a background is usually quoted not in terms of h but in terms of the energy density of the random waves per unit frequency, as a fraction of the closure density of the Universe. This is called Ω_{gw}.

From an operational point of view, the method of detecting a background is independent of its source. Ground-based detectors will work by cross-correlating the output from two detectors that are far enough apart that there is little common instrumental or environmental noise, but close enough that they respond to the gravitational-wave background in the same way. Ideally, "close enough" means nearer than a fraction of a gravitational wavelength. But if this is not practical, then they will still respond in a common way to gravitational waves that arrive from direction transverse to the line joining the detectors; the larger the separation of the detectors, the smaller is the solid angle from which correlated radiation can come, and the smaller is the sensitivity of the pair.

The best-placed pair of detectors for detecting a background are the two LIGO detectors at Washington and Hanford. For *detecting* a background, I discount the two co-located detectors at Hanford, because it would be impossible to verify that an observed correlation was not due to a common disturbance on the site. These detectors can, of course, set a stronger *upper limit* on a background than the separated ones, provided the amount of correlated noise is sufficiently small. Advanced LIGO could detect a background at a level of a few times 10^{-9} of closure. This would represent an improvement on our current best limit of 10^{-8} from pulsar timing, a limit at a very much lower frequency[2]. If the Hanford detectors have little common noise and there is no detectable background at this level, then LIGO may be able to set upper limits as deep as 10^{-10}.

For LISA the detection problem is simpler. LISA will see a background if it is stronger

22

than instrumental noise. (Although there is a possibility to correlate two "independent" detectors, each consisting of a different pair of arms, it is possible to show that the noise from their common arm will give the correlation exactly the same sensitivity as that of a single detector.) It is important in this regard that LISA has an independent way of assessing its instrumental noise: of the four LISA gravitational wave signal streams, one — called the Sagnac signal — is independent of incident gravitational waves in the low-frequency limit[31]. If the Sagnac noise has the expected amplitude and slope at low frequencies, then any excess noise in the other data streams that passes consistency tests can fairly confidently be taken to be external, although it will certainly require considerable effort to achieve this fair degree of confidence. LISA should be able to see a background at a level of $\Omega_{\mathrm{gw}} = 10^{-10}$.

Astrophysical backgrounds.

The compact white-dwarf background is an astrophysical background, although it is not a cosmological background because it is dominated by sources in the Galaxy, whose local density is very much higher than the mean cosmological density of white dwarf binaries. If LISA is to set limits on a primordial background, it will first have to distinguish the local background from the cosmological one. This will be done in two ways: first, the local background is composed of systems belonging to the same population that LISA will resolve at higher frequencies and in its local neighborhood. The statistics of the thousand or so resolved systems will give a good prediction of what to expect from the confusion-limited background. Second, the local background is anisotropic, and LISA's rotating antenna pattern should cause characteristic intensity changes at frequencies where the local background dominates.

At higher frequencies the background from the cosmological distribution of ultra-compact white-dwarf binaries and from neutron-star binaries is expected to lie somewhat below the LISA noise curve, and is expected to be fairly flat in frequency[32]. If these estimates are wrong and this noise is strong enough for LISA to detect, it will be difficult to distinguish between such a background and a primordial one with a different spectrum than expected (see below). The arguments will rely a great deal on modelling the background. Unfortunately, the only neutron-star or ultra-compact white-dwarf systems that LISA will resolve will be in our Galaxy, so that modelling will have to assume that our population is typical.

Astrophysical backgrounds begin to fall off in strength above about 0.1 Hz, both because there are fewer systems that reach this high frequency and because the ones that do pass through the higher band relatively quickly: they are mainly early-stage coalescing neutron-star binaries. Above about 1 Hz these events are probably resolvable in a one-year observation[33].

Primordial background.

Inflation predicts a primordial gravitational wave background, created by the parametric amplification of pre-inflation quantum fluctuations in the gravitational field. These are companions to the fluctuations in the inflaton field that gave rise to the initial density fluctuations that led to galaxy formation. Observations of the cosmic microwave background show us that, at the epoch of decoupling, the scalar fluctuations with wavelength of the order of the cosmological horizon size were of order 10^{-6}. One might expect the tensor fluctuations to have produced gravitational waves of a similar amplitude at that time, but not larger. Their energy density (given by the energy flux in Equation 4, divided by c) would then have been of order 10^{-12} of the closure density, since the energy density depends on the square of the amplitude and their frequency was the Hubble "constant" at that epoch. As the universe expanded after that, the energy density of this background redshifted away as $(1+z)^{-4}$, while the closure density has redshifted approximately as $(1+z)^{-3}$, so the fraction Ω_{gw} has gone down by a further factor of 1000 to 10^{-15}. Since there is no intrinsic length scale in inflation or in the primordial quantum fluctuations, the spectrum (energy density per unit frequency) is essentially flat, independent of frequency. From our discussion above it is clear that this is a difficult target to shoot for.

However, inflation is only the simplest scenario for the early universe that is consistent with what we observe today. There are many alternatives. Superstring-inspired cosmologies, which modify the pre-inflationary tensor fluctuation spectrum, can produce very strong stochastic waves with spectra shaped (by some arbitrary parameters) to peak at any frequency one likes. Brane cosmologies can also produce strong waves, while ekpyrotic models produce no primordial gravitational radiation at all[2].

It is clear that detecting primordial gravitational waves would be one of the most fundamental outcomes of the development of gravitational wave detectors. The strength and spectrum would discriminate among all these models and provide some of the most important hints to guide the development of grand unified theories. From our discussion of astrophysical backgrounds, the primordial background may turn out to be much weaker than the astrophysical backgrounds below 1 Hz. For this reason, NASA is encouraging LISA scientists to look ahead already to try to sketch the design of a LISA follow-on that would observe very deeply in the gap region between LIGO and LISA, the range from 0.5–5 Hz.

CONCLUSION

We have surveyed a broad spectrum of sources of gravitational waves and seen how to make rough order-of-magnitude estimates of their strengths and characteristics. I don't want to give the impression that these rough calculations are all we know; in almost every case there is a great deal of detailed modelling, without which we might have less confidence in the rough-and-ready calculations employed here. Many of the subjects that I have treated superficially are discussed in much greater detail by other speakers at this meeting.

But at the same time, it is important to have simple, physically accessible ways of understanding potential sources, especially when it comes to fitting models to the observations. It may of course happen that our detectors will show us only what we expect to see and nothing more. But it seems to me to be more likely that with increasing sensitivity we will soon come face-to-face with waveforms that we did not expect and whose properties we can only incompletely measure. At that point we will need all the physical intuition we have available in order to understand what new and unexpected things the Universe is doing.

ACKNOWLEDGMENTS

REFERENCES

1. Schutz, B.F., "Confronting Theory with Gravitational Wave Observations", in *Gravitation: a Decennial Perspective*, ed. Ashtekar, A., to be published.
2. Cutler, C., and Thorne, K.S., in *General Relativity and Gravitation* eds. Bishop, N.T., and Maharaj, S.D. (World Scientific, Singapore, 2002). Also gr-qc/0204090.
3. Baker, J., Abrahams, A., Anninos, P., Brandt, S., Price, R., Pullin, J., Seidel, E., *Phys. Rev.* D **55**, 829 (1997).
4. Schutz, B.F., *Am. J. Phys.*, **52**, 412–419 (1984).
5. Misner, C.W., Thorne, K.S., and Wheeler, J.L., *Gravitation* (Freeman & Co., San Francisco, 1973).
6. Schutz, B.F., *Nature*, **323** 310 (1986).
7. Schutz, B.F., in *Lighthouses of the Universe: The Most Luminous Celestial Objects and their Use for Cosmology*, M. Gilfanov, R. Sunyeav, E. Churazov, B. Leibundgut (ed s)(Springer, Berlin, 2001), 207.
8. Holz, D.E., and Hughes, S.A., *Phys. Rev.* D, to be published. Also astro-ph/0212218.
9. McMillan, S., and Portegies Zwart, S., *Astrophys. J. Lett.* **528** L17 (2000).
10. Madau, P., and Kuhlen, M., in *Proceedings of the XXI Texas Symposium on Relativistic Astrophysics*, ed. Pacini, F., to be published.
11. Smarr, L., v Cadev z, A., DeWitt, B., and Eppley, K., *Phys. Rev.* D, **14**, 10, 2443 (1976).
12. Alcubierre, M., Brügmann, B., Dramlitsch, T., Font, J.A., Papadopoulos, P., Seidel, E., Stergioulas, N., and Takahashi, R., *Phys. Rev.* D, **62**, 044034 (2000).
13. Baker, J., Campanelli, M., and Lousto, C.O., *Phys. Rev.* D **65**, 044001 (1991).
14. Seidel, E., *et al*, to be published.
15. Alcubierre, M., and Brügmann, B., *Phys. Rev.* D **63**, 104006 (2001). Also gr-qc/0008067.
16. Gomez, R., *et al*, *Phys. Rev. Lett.* **80**, 3915 (1998).
17. Bowen, J., and York, J.W., *Phys. Rev.* D **21**, 2047 (1980).
18. http://www.eunetwork.org
19. Buonanno, A., and Damour, T., *Phys. Rev.* D **62**, 064015 (2000).
20. Dimmelmeir, H., Font, J.A., and Müller, E., *Astron. Astrophys.* **393** 523 (2002). Also astro-ph/0204289.
21. New, K.C.B., and Shapiro, S.L., *Astrophys. J.* **548**, 439 (2001).

22. Blanchet, L., *Living Rev. Relativity*, **5**, 3 (2002). [Online article]: cited on 24 July 2003, http://www.livingreviews.org/lrr-2002-3.
23. Ushomirsky, G., Cutler, C., and Bildsten, L., *Mon. Not. Roy. astr. Soc.* **319**, 902 (2000).
24. Cutler, C., *Phys. Rev.* D, **66** 084025. Also gr-qc/0206051.
25. Brady, P.R., Creighton, T., Cutler, C., and Schutz, B.F., *Phys. Rev.* D **57**, 2101 (1998).
26. Schutz, B.F., and Papa, M.A., in *Gravitational Waves and Experimental Gravity*, eds. Trân Thanh Vân, J. Dumarchez, S. Reynaud, C.Salomon, S. Thorsett, J.Y. Vinet (World Publishers, Hanoi, 2000), 199. Also: gr-qc/9905018.
27. Lindblom, L., and Owen, B., *Phys. Rev.* D **65**, 063006 (2002). Also: astro-ph/0110558.
28. Andersson, N., and Kokkotas, K.D., in *Recent developments in general relativity* eds. R. Cianci, R. Collina, M. Francaviglia, P. Fré (Springer, Milano 2002), 121.
29. Finn, L.S., and Thorne, K.S., *Phys. Rev.* D **62**, 124021 (2000).
30. Barack, L., and Ori, A., *Phys. Rev. Lett.* **90** 111101 (2003). Also: gr-qc/0212103.
31. Tinto, M., Armstrong, J.W., and Estabrook, F.B., *Phys. Rev.* D, **63**, 021010 (2000).
32. Schneider, R., Ferrari, S., Matarrese, S., and Portegies Zwart, S.F., *Mon. Not. Roy. astr. Soc.* **324** 797 (2001). Also: astro-ph/0002055.
33. Ungarelli, C., and Vecchio, A., *Phys. Rev.* D, **64**, 121501 (2001).